



Implementing Cisco IOS Network Security (IINS)

Foundation Learning Guide

Second Edition



Implementing Cisco IOS Network Security (IINS) Foundation Learning Guide

Second Edition

Catherine Paquet

Cisco Press
800 East 96th Street
Indianapolis, Indiana 46240 USA

Implementing Cisco IOS Network Security (IINS) Foundation Learning Guide Second Edition

Catherine Paquet

Copyright© 2013 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street
Indianapolis, IN 46240 USA

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

Printed in the United States of America

First Printing November 2012

Library of Congress Cataloging-in-Publication data is on file.

ISBN-13: 978-1-58714-272-7

ISBN-10: 1-58714-272-4

Warning and Disclaimer

This book is designed to provide information about implementing Cisco IOS network security with information necessary to prepare for Cisco exam 640-554, Implementing Cisco IOS Network Security. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

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

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please

contact:

U.S. Corporate and Government Sales

1-800-382-3419

corpsales@pearsontechgroup.com

For sales outside the United States, please contact:

International Sales

international@pearsoned.com

Feedback Information

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

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

We greatly appreciate your assistance.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Manager Global Certification: Erik Ullanderson

Business Operation Manager, Cisco Press: Anand Sundaram

Executive Editor: Brett Bartow

Managing Editor: Sandra Schroeder

Development Editor: Kimberley Debus

Senior Project Editor: Tonya Simpson

Copy Editor: Bill McManus

Technical Editor: Kevin Redmon

Editorial Assistant: Vanessa Evans

Book Designer: Louisa Adair

Cover Designer: Mark Shirar

Composition: Bronkella Publishing

Indexer: Tim Wright

Proofreader: Sheri Cain



Americas Headquarters

Cisco Systems, Inc
San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters

Cisco Systems International BV
Amsterdam, The Netherlands

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.

♻️ CCDE, CCENT Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catelyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo 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. (0812R)

About the Author

Catherine Paquet is a practitioner in the field of internetworking, network security, and security financials. She has authored or contributed to ten books thus far with Cisco Press. Catherine has in-depth knowledge of security systems, remote access, and routing technology. She is a Cisco Certified Network Professional (CCNP) and a CCNP Security. Catherine is also a Cisco IronPort Certified Security Instructor (CICSI) and a Certified Cisco Systems Instructor (CCSI) with Cisco's largest training partner, Global Knowledge, Inc. She also works on IT security projects and implementations for different organizations on a part-time basis. Following her university graduation from the Collège Militaire Royal de St-Jean (Canada), Catherine worked as a system analyst, LAN manager, MAN manager, and eventually as a WAN manager. Later, she received a master's degree in business administration (MBA) with a specialty in management information systems (MIS) from York University.

Catherine has lectured for the Computer Security Institute and for Cisco Systems (Emerging Markets) on the topic of the business case for network security. In 2002 and 2003, she volunteered with the U.N. mission in Kabul, Afghanistan, to train Afghan public servants in the area of networking.

Catherine lives in Toronto with her husband. They have two children, who are both attending college.

About the Technical Reviewer

Kevin Redmon has been an employee of Cisco Systems, Inc. in Research Triangle Park, North Carolina since October 2000. He has a bachelor of science in computer engineering from Case Western Reserve University (Cleveland, Ohio) and a master of science in information security from East Carolina University (Greenville, North Carolina). Kevin was a customer support engineer with the Cisco TAC Firewall Team from September 2007 to March 2011 and now supports the TAC VPN team at Cisco. Kevin enjoys innovating new ideas to keep his mind fresh and currently has a patent listed with the United States Patent and Trade Office. Kevin spends his free time playing mandolin, writing software for home projects, hacking and modding home electronics, and relaxing with his wife and baby girl in Durham, North Carolina.

Dedication

This book is dedicated to my mother, Florence Jacques-Paquet, who became gravely ill during this project. She pulled through—all her life, my mother pulled through. She was ahead of her time, in her thoughts and in her actions, and she made it a priority that H el ene and I developed, like her, self-reliance. Mom, thanks for the gift of education, tenacity, and resiliency. I love you.

Acknowledgments

I'd like to give special recognition to Kevin Redmon for providing his expert technical knowledge in editing this book. Kevin's meticulous and holistic approach to security solutions is unsurpassed. He was not afraid to point out inaccuracies and make recommendations to improve the manuscript. Thank you, Kevin.

A big "thank you" goes out to the production team for this book: Brett Bartow, Drew Cupp, and especially Kimberley Debus, Tonya Simpson, and Bill McManus, who have been incredibly professional and a pleasure to work with. I couldn't have asked for a finer team.

Acknowledgements for this book wouldn't be complete without mentioning my husband of 25 years, Pierre Rivard. Another book, so another year where Pierre spent countless evenings and weekends alone while his wife was working on a manuscript. His understanding, patience, and personal delivery of splendid meals to my "eagle nest" were truly appreciated. Pierre is my rock, my shelter, my soulmate. Pierre, je t'aime.

Contents at a Glance

[Introduction](#)

[Part I Networking Security Fundamentals](#)

[Chapter 1 Network Security Concepts and Policies](#)

[Chapter 2 Security Strategy and Cisco Borderless Network](#)

[Part II Protecting the Network Infrastructure](#)

[Chapter 3 Network Foundation Protection and Cisco Configuration Professional](#)

[Chapter 4 Securing the Management Plane on Cisco IOS Devices and AAA](#)

[Chapter 5 Securing the Data Plane on Cisco Catalyst Switches](#)

[Chapter 6 Securing the Data Plane in IPv6 Environments](#)

[Part III Threat Control and Containment](#)

[Chapter 7 Planning a Threat Control Strategy](#)

[Chapter 8 Access Control Lists for Threat Mitigation](#)

[Chapter 9 Firewall Fundamentals and Network Address Translation](#)

[Chapter 10 Cisco Firewalling Solutions: Cisco IOS Zone-Based Firewall and Cisco ASA](#)

[Chapter 11 Intrusion Prevention Systems](#)

[Part IV Secure Connectivity](#)

[Chapter 12 Fundamentals of Cryptography and VPN Technologies](#)

[Chapter 13 IPsec Fundamentals](#)

[Chapter 14 Site-to-Site IPsec VPNs with Cisco IOS Routers](#)

[Chapter 15 SSL VPNs with Cisco ASA](#)

[Appendix Answers to Chapter Review Questions](#)

[Index](#)

Contents

Introduction

Part I Networking Security Fundamentals

Chapter 1 Network Security Concepts and Policies

Building Blocks of Information Security

Basic Security Assumptions

Basic Security Requirements

Data, Vulnerabilities, and Countermeasures

Data Classification

Vulnerabilities Classifications

Countermeasures Classification

Need for Network Security

Intent Evolution

Threat Evolution

Trends Affecting Network Security

Adversaries, Methodologies, and Classes of Attack

Adversaries

Methodologies

Threats Classification

Man-in-the-Middle Attacks

Overt and Covert Channels

Botnets

DoS and DDoS Attacks

Principles of Secure Network Design

Defense in Depth

Evaluating and Managing the Risk

Levels of Risks

Risk Analysis and Management

Risk Analysis

Building Blocks of Risk Analysis

A Lifecycle Approach to Risk Management

Regulatory Compliance

Security Policies

Security Policy Components

[*Governing Policy*](#)

[*End-User Policies*](#)

[*Technical Policies*](#)

[*Standards, Guidelines, and Procedures*](#)

[*Security Policy Roles and Responsibilities*](#)

[*Security Awareness*](#)

[Secure Network Lifecycle Management](#)

[IT Governance, Risk Management, and Compliance](#)

[Secure Network Life Cycle](#)

[*Initiation Phase*](#)

[*Acquisition and Development Phase*](#)

[*Implementation Phase*](#)

[*Operations and Maintenance Phase*](#)

[*Disposition Phase*](#)

[*Models and Frameworks*](#)

[Network Security Posture](#)

[Network Security Testing](#)

[*Security Testing Techniques*](#)

[*Common Testing Tools*](#)

[Incident Response](#)

[Incident Management](#)

[*Computer Crime Investigations*](#)

[*Laws and Ethics*](#)

[*Liability*](#)

[Disaster Recovery and Business Continuity Planning](#)

[*Business Continuity Concepts*](#)

[Summary](#)

[References](#)

[Publications](#)

[Web Resources](#)

[Review Questions](#)

[**Chapter 2 Security Strategy and Cisco Borderless Network**](#)

[Borderless Networks](#)

[Cisco Borderless Network Security Architecture](#)

[Borderless End Zone](#)

[Borderless Internet](#)

[Borderless Data Center](#)

[Policy Management Layer](#)

[Borderless Network Services](#)

[Borderless Security Products](#)

[SecureX, a Context-Aware Security Approach](#)

[*SecureX Core Components*](#)

[Threat Control and Containment](#)

[Cisco Security Intelligence Operation](#)

[Cloud Security, Content Security, and Data Loss Prevention](#)

[*Content Security*](#)

[*Data Loss Prevention*](#)

[*Cloud-Based Security*](#)

[*Web Security*](#)

[*Email Security*](#)

[Secure Connectivity Through VPNs](#)

[Security Management](#)

[*Cisco Security Manager*](#)

[Summary](#)

[References](#)

[Review Questions](#)

Part II Protecting the Network Infrastructure

Chapter 3 Network Foundation Protection and Cisco Configuration Professional

[Threats Against the Network Infrastructure](#)

[Cisco NFP Framework](#)

[Control Plane Security](#)

[*CoPP*](#)

[*CPPr*](#)

[*Traffic Classes*](#)

[*Routing Protocol Integrity*](#)

[*Cisco AutoSecure*](#)

[Management Plane Security](#)

[*Secure Management and Reporting*](#)

[*Role-Based Access Control*](#)

[*Deploying AAA*](#)

[Data Plane Security](#)

[*Access Control List Filtering*](#)

[Cisco Configuration Professional](#)

[*CCP Initial Configuration*](#)

[Cisco Configuration Professional User Interface and Features](#)

[*Menu Bar*](#)

[*Toolbar*](#)

[*Navigation Pane*](#)

[*Content Pane*](#)

[*Status Bar*](#)

[Cisco Configuration Professional Building Blocks](#)

[*Communities*](#)

[*Creating Communities*](#)

[*Managing Communities*](#)

[*Templates*](#)

[*User Profiles*](#)

[Using CCP to Harden Cisco IOS Devices](#)

[*Security Audit*](#)

[*One-Step Lockdown*](#)

[*Cisco IOS AutoSecure*](#)

[Summary](#)

[References](#)

[Review Questions](#)

[**Chapter 4 Securing the Management Plane on Cisco IOS Devices and AAA**](#)

[Configuring Secure Administration Access](#)

[*Configuring an SSH Daemon for Secure Management Access*](#)

[*Configuring Passwords on Cisco IOS Devices*](#)

[*Setting Timeouts for Router Lines*](#)

[*Configuring the Minimum Length for Router Passwords*](#)

[*Enhanced Username Password Security*](#)

[*Securing ROM Monitor*](#)

[*Securing the Cisco IOS Image and Configuration Files*](#)

[*Configuring Multiple Privilege Levels*](#)

[*Configuring Role-Based Command-Line Interface Access*](#)

[Implementing Secure Management and Reporting](#)

[*Planning Considerations for Secure Management and Reporting*](#)

[*Secure Management and Reporting Architecture*](#)

[*Secure Management and Reporting Guidelines*](#)

[*Enabling Time Features*](#)

[*Network Time Protocol*](#)

[*Using Syslog Logging for Network Security*](#)

[*Implementing Log Messaging for Security*](#)

[*Using SNMP to Manage Network Devices*](#)

[*SNMPv3 Architecture*](#)

[*Enabling SNMP Options Using Cisco CCP*](#)

[*Configuring AAA on a Cisco Router*](#)

[*Authentication, Authorization, and Accounting*](#)

[*Authenticating Router Access*](#)

[*Configuring AAA Authentication and Method Lists*](#)

[*Configuring AAA on a Cisco Router Using the Local Database*](#)

[*Configuring AAA Local Authentication*](#)

[*AAA on a Cisco Router Using Cisco Secure ACS*](#)

[*Cisco Secure ACS Overview*](#)

[*Cisco Identity Services Engine*](#)

[*TACACS+ and RADIUS Protocols*](#)

[*TACACS+*](#)

[*RADIUS*](#)

[*Comparing TACACS+ and RADIUS*](#)

[*AAA on a Cisco Router Using an External Database*](#)

[*Configuration Steps for AAA Using an External Database*](#)

[*AAA Servers and Groups*](#)

[*AAA Authentication Method Lists*](#)

[*AAA Authorization Policies*](#)

[*AAA Accounting Policies*](#)

[*AAA Configuration for TACACS+ Example*](#)

[*Troubleshooting TACACS+*](#)

[*Deploying and Configuring Cisco Secure ACS*](#)

[*Evolution of Authorization*](#)

[*Before: Group-Based Policies*](#)

[*Now: More Than Just Identities*](#)

[*Rule-Based Policies*](#)

[*Configuring Cisco Secure ACS 5.2*](#)

[*Configuring Authorization Policies for Device Administration*](#)

[*Summary*](#)

[References](#)

[Review Questions](#)

Chapter 5 Securing the Data Plane on Cisco Catalyst Switches

[Overview of VLANs and Trunking](#)

[Trunking and 802.1Q](#)

[*802.1Q Tagging*](#)

[*Native VLANs*](#)

[Configuring VLANs and Trunks](#)

[*Step 1: Configuring and Verifying 802.1Q Trunks*](#)

[*Step 2: Creating a VLAN*](#)

[*Step 3: Assigning Switch Ports to a VLAN*](#)

[*Step 4: Configuring Inter-VLAN Routing*](#)

[Spanning Tree Overview](#)

[STP Fundamentals](#)

[Verifying RSTP and PVRST+](#)

[Mitigating Layer 2 Attacks](#)

[Basic Switch Operation](#)

[Layer 2 Best Practices](#)

[Layer 2 Protection Toolkit](#)

[Mitigating VLAN Attacks](#)

[*VLAN Hopping*](#)

[Mitigating Spanning Tree Attacks](#)

[*PortFast*](#)

[Mitigating CAM Table Overflow Attacks](#)

[Mitigating MAC Address Spoofing Attacks](#)

[Using Port Security](#)

[*Errdisable Recovery*](#)

[Summary](#)

[References](#)

[Review Questions](#)

Chapter 6 Securing the Data Plane in IPv6 Environments

[The Need for IPv6](#)

[IPv6 Features and Enhancements](#)

[IPv6 Headers](#)

[Stateless Address Autoconfiguration](#)

[Internet Control Message Protocol Version 6](#)

[IPv6 General Features](#)

[Transition to IPv6](#)

[IPv6 Addressing](#)

[IPv6 Address Representation](#)

[IPv6 Address Types](#)

[*IPv6 Unicast Addressing*](#)

[Assigning IPv6 Global Unicast Addresses](#)

[*Manual Interface Assignment*](#)

[*EUI-64 Interface ID Assignment*](#)

[*Stateless Autoconfiguration*](#)

[*DHCPv6 \(Stateful\)*](#)

[IPv6 EUI-64 Interface Identifier](#)

[IPv6 and Cisco Routers](#)

[IPv6 Address Configuration Example](#)

[Routing Considerations for IPv6](#)

[Revisiting Threats: Considerations for IPv6](#)

[Examples of Possible IPv6 Attacks](#)

[*Recommended Practices*](#)

[Summary](#)

[References](#)

[Review Questions](#)

Part III Threat Control and Containment

Chapter 7 Planning a Threat Control Strategy

[Threats Revisited](#)

[Trends in Network Security Threats](#)

[Threat Mitigation and Containment: Design Fundamentals](#)

[*Threat Control Design Guidelines*](#)

[*Application Layer Visibility*](#)

[*Distributed Security Intelligence*](#)

[*Security Intelligence Analysis*](#)

[Integrated Threat Control Strategy](#)

[Cisco Threat Control and Containment Categories](#)

[*Integrated Approach to Threat Control*](#)

[*Application Awareness*](#)

[*Application-Specific Gateways*](#)

[*Security Management*](#)

[*Cisco Security Intelligence Operations Site*](#)

[*Cisco Threat Control and Containment Solutions Fundamentals*](#)

[*Cisco Security Appliances*](#)

[*Cisco IPSs*](#)

[Summary](#)

[References](#)

[Review Questions](#)

Chapter 8 Access Control Lists for Threat Mitigation

[ACL Fundamentals](#)

[Types of IP ACLs](#)

[ACL Wildcard Masking and VLSM Review](#)

[Subnetting Overview](#)

[*Subnetting Example: Class C*](#)

[*Subnetting Example*](#)

[Variable-Length Subnet Masking](#)

[*A Working VLSM Example*](#)

[ACL Wildcard Bits](#)

[*Example: Wildcard Masking Process for IP Subnets*](#)

[*Example: Wildcard Masking Process with a Single IP Address*](#)

[*Example: Wildcard Masking Process with a Match Any IP Address*](#)

[Using ACLs to Control Traffic](#)

[*Example: Numbered Standard IPv4 ACL—Deny a Specific Subnet*](#)

[*Numbered Extended IPv4 ACL*](#)

[*Displaying ACLs*](#)

[Enhancing ACLs with Object Groups](#)

[ACL Considerations](#)

[Configuring ACLs for Threat Control Using Cisco Configuration Professional](#)

[Rules in Cisco Configuration Professional](#)

[*Working with ACLs in CCP*](#)

[*ACL Editor*](#)

[*Adding Rules*](#)

[*Associating Rules with Interfaces*](#)

[*Enabling Logging with CCP*](#)

[*Monitoring ACLs with CCP*](#)

[*Configuring an Object Group with CCP*](#)

[Using ACLs in IPv6 Environments](#)

[Summary](#)

[References](#)

[Review Questions](#)

Chapter 9 Firewall Fundamentals and Network Address Translation

[Introducing Firewall Technologies](#)

[Firewall Fundamentals](#)

[Firewalls in a Layered Defense Strategy](#)

[Static Packet-Filtering Firewalls](#)

[Application Layer Gateways](#)

[Dynamic or Stateful Packet-Filtering Firewalls](#)

[Other Types of Firewalls](#)

[*Application Inspection Firewalls, aka Deep Packet Inspection*](#)

[*Transparent Firewalls \(Layer 2 Firewalls\)*](#)

[NAT Fundamentals](#)

[Example of Translating an Inside Source Address](#)

[NAT Deployment Choices](#)

[Firewall Designs](#)

[Firewall Policies in a Layered Defense Strategy](#)

[Firewall Rules Design Guidelines](#)

[Summary](#)

[References](#)

[Review Questions](#)

Chapter 10 Cisco Firewalling Solutions: Cisco IOS Zone-Based Firewall and Cisco ASA

[Cisco Firewall Solutions](#)

[Cisco IOS Zone-Based Policy Firewall](#)

[Zone-Based Policy Firewall Overview](#)

[Zones and Zone Pairs](#)

[*Self Zone*](#)

[*Zone-Based Topology Examples*](#)

[Introduction to Cisco Common Classification Policy Language](#)

[Zone-Based Policy Firewall Actions](#)

[Service Policy Zone Pair Assignments](#)

[Zone-Based Policy Firewall: Default Policies, Traffic Flows, and Zone Interaction](#)

[*Zone-Based Policy Firewall: Rules for Router Traffic*](#)

[*Configuring Basic Interzone Policies Using CCP and the CLI*](#)

[*Step 1: Start the Basic Firewall Wizard*](#)

[*Step 2: Select Trusted and Untrusted Interfaces*](#)

[*Step 3: Review and Verify the Resulting Policies*](#)

[*Verifying and Tuning the Configuration*](#)

[*Step 4: Enabling Logging*](#)

[*Step 5: Verifying Firewall Status and Activity*](#)

[*Step 6: Modifying Zone-Based Firewall Configuration Objects*](#)

[*Step 7: Verifying the Configuration Using the CLI*](#)

[*Configuring NAT Services for Zone-Based Firewalls*](#)

[*Step 1: Run the Basic NAT Wizard*](#)

[*Step 2: Select NAT Inside and Outside Interfaces*](#)

[*Step 3: Verify NAT with CCP and the CLI*](#)

[Cisco ASA Firewall](#)

[*Stateful Packet Filtering and Application Awareness*](#)

[*Network Services Offered by the Cisco ASA 5500 Series*](#)

[*Network Address Translation*](#)

[*Additional Network Services*](#)

[*Cisco ASA Security Technologies*](#)

[*Cisco ASA Configuration Fundamentals*](#)

[*Cisco ASA 5505*](#)

[*Cisco ASDM*](#)

[*Preparing the Cisco ASA 5505 for ASDM*](#)

[*Cisco ASDM Features and Menus*](#)

[*Cisco Modular Policy Framework*](#)

[*Class Map: Identifying Traffic on Which a Policy Will Be Enforced*](#)

[*Policy Map: Configuring the Action That Will Be Applied to the Traffic*](#)

[*Service Policy: Activating the Policy*](#)

[*Cisco ASA Modular Policy Framework: Simple Example*](#)

[*Basic Outbound Access Control on Cisco ASA Using Cisco ASDM*](#)

[*Scenario Configuration Steps Using Cisco ASDM*](#)

[Summary](#)

[References](#)

[Cisco.com Resources](#)

[Other Resources](#)

[CCP and ASDM Demo Mode Tutorials](#)

[Chapter 11 Intrusion Prevention Systems](#)

[IPS Fundamentals](#)

[Introducing IDS and IPS](#)

[So, IDS or IPS? Why Not Both?](#)

[Alarm Types](#)

[Intrusion Prevention Technologies](#)

[Signature-Based IDS/IPS](#)

[Policy-Based IDS/IPS](#)

[Anomaly-Based IDS/IPS](#)

[Reputation-Based IPS](#)

[IPS Attack Responses](#)

[IPS Anti-Evasion Techniques](#)

[Risk-Based Intrusion Prevention](#)

[IPv6-Aware IPS](#)

[Alarms](#)

[IPS Alarms: Event Monitoring and Management](#)

[Global Correlation](#)

[IPS Deployment](#)

[Cisco IPS Offerings](#)

[IPS Best Practices](#)

[Cisco IPS Architecture](#)

[Cisco IOS IPS](#)

[Cisco IOS IPS Features](#)

[Scenario: Protecting the Branch Office Against Inside Attack](#)

[Signatures](#)

[Signature Files](#)

[Signature Management](#)

[Examining Signature Microengines](#)

[Signature Tuning](#)

[Optimal Signature Set](#)

[Monitoring IPS Alarms and Event Management](#)

[Configuring Cisco IOS IPS Using Cisco Configuration Professional](#)

[Step 1: Download Cisco IOS IPS Signature Package](#)

[Step 2: Launch IPS Policies Wizard](#)

[Step 3: Verify Configuration and Signature Files](#)

[Step 4: Perform Signature Tuning](#)

[Step 5: Verify Alarms](#)

[Configuring Cisco IOS IPS Using the CLI](#)

[Summary](#)

[References](#)

[Cisco.com Resources](#)

[General IDS/IPS Resource](#)

[Review Questions](#)

[Part IV Secure Connectivity](#)

[Chapter 12 Fundamentals of Cryptography and VPN Technologies](#)

[VPN Overview](#)

[VPN Types](#)

[Site-to-Site VPNs](#)

[Remote-Access VPNs](#)

[Examining Cryptographic Services](#)

[Cryptology Overview](#)

[The History of Cryptography](#)

[Ciphers](#)

[Block and Stream Ciphers](#)

[Block Ciphers](#)

[Stream Ciphers](#)

[The Process of Encryption](#)

[Encryption Application Examples](#)

[Cryptanalysis](#)

[Desirable Encryption Algorithm Features](#)

[Key Management](#)

[Key Management Components](#)

[Keyspaces](#)

[Key Length Issues](#)

[Example of the Impact of Key Length](#)

[Symmetric and Asymmetric Encryption Overview](#)

[Symmetric Encryption Algorithms](#)

[Comparing Symmetric Encryption Algorithms](#)

[DES Modes of Operation](#)

[DES Security Guidelines](#)

[The Rijndael Cipher](#)

[*AES Versus 3DES*](#)

[Asymmetric Encryption Algorithms](#)

[*Public Key Confidentiality*](#)

[Encryption Algorithm Selection](#)

[Cryptographic Hashes and Digital Signatures](#)

[Hashing Algorithms](#)

[*MD5*](#)

[*SHA-1*](#)

[*SHA-2*](#)

[Hashed Message Authentication Codes](#)

[Overview of Digital Signatures](#)

[*Digital Signatures = Encrypted Message Digest*](#)

[Diffie-Hellman](#)

[Diffie-Hellman Example](#)

[Cryptographic Processes in VPNs](#)

[Asymmetric Encryption: Digital Signatures](#)

[Asymmetric Encryption Overview](#)

[*Public Key Authentication*](#)

[RSA and Digital Signatures](#)

[Public Key Infrastructure](#)

[PKI Terminology and Components](#)

[Certificate Classes](#)

[Certificate Authorities](#)

[PKI Standards](#)

[*Certificate Revocation*](#)

[Certificate Use](#)

[*Digital Certificates and CAs*](#)

[Summary](#)

[References](#)

[Books and Articles](#)

[Standards](#)

[Encryption Regulations](#)

[Review Questions](#)

Chapter 13 IPsec Fundamentals

[IPsec Framework](#)

[Suite B Cryptographic Standard](#)

[Encryption Algorithms](#)

[Key Exchange: Diffie-Hellman](#)

[Data Integrity](#)

[Authentication](#)

[IPsec Protocol](#)

[Authentication Header](#)

[Encapsulating Security Payload](#)

[IPsec Modes of Operations](#)

[*Transport Mode*](#)

[*Tunnel Mode*](#)

[IKE Protocol](#)

[IKEv1 Modes](#)

[IKEv1 Phases](#)

[*IKEv1 Phase 1*](#)

[*IKEv1 Phase 1 Example*](#)

[*IKEv1 Phase 2*](#)

[IKE Version 2](#)

[IKEv1 Versus IKEv2](#)

[IPv6 VPNs](#)

[IPsec Services for Transitioning to IPv6](#)

[Summary](#)

[References](#)

[Books](#)

[Cisco.com Resources](#)

[Review Questions](#)

Chapter 14 Site-to-Site IPsec VPNs with Cisco IOS Routers

[Site-to-Site IPsec: Planning and Preparation](#)

[Site-to-Site IPsec VPN Operations](#)

[Planning and Preparation Checklist](#)

[Building Blocks of Site-to-Site IPsec](#)

[*Interesting Traffic and Crypto ACLs*](#)

[*Mirrored Crypto ACLs*](#)

[*Cipher Suite*](#)

[*Crypto Map*](#)

[Configuring a Site-to-Site IPsec VPN Using CCP](#)

[Initiating the VPN Wizard](#)

[*VPN Connection Information*](#)

[*IKE Proposals*](#)

[*Transform Set*](#)

[*Traffic to Protect*](#)

[*Configuration Summary*](#)

[*Creating a Mirror Configuration for the Peer Site*](#)

[*Verifying the IPsec Configuration Using CCP and CLI*](#)

[*Verifying IPsec Configuration Using CLI*](#)

[*Verifying IKE Policy Using the CLI*](#)

[*Verifying IKE Phase 2 Policy Using the CLI*](#)

[*Verifying Crypto Maps Using the CLI*](#)

[*Monitoring Established IPsec VPN Connections*](#)

[*IKE Policy Negotiation*](#)

[*VPN Troubleshooting*](#)

[*Monitoring IKE Security Association*](#)

[*Monitoring IPsec Security Association*](#)

[*Summary*](#)

[*References*](#)

[*Review Questions*](#)

Chapter 15 SSL VPNs with Cisco ASA

[*SSL VPNs in Borderless Networks*](#)

[*Cisco SSL VPN*](#)

[*SSL and TLS Protocol Framework*](#)

[*SSL and TLS*](#)

[*SSL Cryptography*](#)

[*SSL Tunnel Establishment*](#)

[*SSL Tunnel Establishment Example*](#)

[*Cisco SSL VPN Deployment Options and Considerations*](#)

[*Cisco SSL VPN Client: Full Network Access*](#)

[*SSL VPN on Cisco ASA in Clientless Mode*](#)

[*Clientless Configuration Scenario*](#)

[*Task 1: Launch the Clientless SSL VPN Wizard from ASDM*](#)

[*Task 2: Configure the SSL VPN Interface*](#)

[*Task 3: Configure User Authentication*](#)

[*Task 4: Configure User Group Policy*](#)

[*Task 5: Configure a Bookmark List*](#)

[Task 6: Verify the Clientless SSL VPN Wizard Configuration](#)

[Log In to the VPN Portal: Clientless SSL VPN](#)

[SSL VPN on ASA Using the Cisco AnyConnect VPN Client](#)

[Cisco AnyConnect Configuration Scenario](#)

[Phase 1: Configure Cisco ASA for Cisco AnyConnect](#)

[*Task 1: Connection Profile Identification*](#)

[*Task 2: VPN Protocols and Device Certificate*](#)

[*Task 3: Client Image*](#)

[*Task 4: Authentication Methods*](#)

[*Task 5: Client Address Assignment*](#)

[*Task 6: Network Name Resolution Servers*](#)

[*Task 7: Network Address Translation Exemption*](#)

[*Task 8: AnyConnect Client Deployment Summary*](#)

[Phase 2: Configure the Cisco AnyConnect VPN Client](#)

[Phase 3: Verify VPN Connectivity with Cisco AnyConnect VPN Client](#)

[*Verifying VPN Connectivity from Cisco ASA*](#)

[Summary](#)

[References](#)

[Review Questions](#)

[**Appendix A Answers to Chapter Review Questions**](#)

[**Index**](#)

Icons Used in This Book




Router



Voice-Enabled Router



Router with Firewall



Wireless Access Point



NAC Appliance



Multilayer Switch



Switch




ATM/Frame Relay Switch



Secure Catalyst Switch



Cisco ASA



IOS Firewall



PIX Firewall



Firewall Services Module



Firewall



VPN Concentrator



Cisco Mars



Sensor/IDS



Access Server



Cisco Unity Server



Cisco CallManager



IP Phone



Analog Phone



PBX Switch



Phone



PC



Laptop



Security Management



Server



Web Server



Wireless Connection



Ethernet Connection



Serial Connection



Network Cloud

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a show command).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ([{ }]) indicate a required choice within an optional element.

Introduction

Network security is a complex and growing area of IT. As the premier provider of network security devices, Cisco Systems is committed to supporting this growing segment of the industry.

This book teaches you how to design, configure, maintain, and audit network security. It focuses on using Cisco IOS routers for protecting the network by capitalizing on their advanced features as a perimeter router, as a firewall, as an intrusion prevention system, and as a site-to-site VPN device. The book also covers the use of Cisco Catalyst switches for basic network security. While covering the topic of authentication, authorization, and accounting (AAA), this book also introduces Cisco Secure Access Control System (ACS). The final chapter also introduces how to use a Cisco Adaptive Security Appliance (ASA) for both clientless and full client remote-access VPNs. At the end of this book, you will be able to select and implement the appropriate Cisco appliances and services required to build flexible and secure networks.

This book provides you with the knowledge necessary to pass your CCNA Security certification (IINS v2.0) because it provides in-depth information to help you prepare for the IINS exam, which grants the CCNA Security certification. It also starts you on the path toward attaining your Cisco Certified Network Professional (CCNP) Security certification.

The commands and configuration examples presented in this book are based on Cisco IOS Releases 15, Cisco ASA 8.4, and Cisco ACS 5.2.

Goals and Methods

The most important and somewhat obvious goal of this book is to help you pass the IINS v2.0 exam (640-554). In fact, if the primary objective of this book were different, the book's title would be misleading; however, the methods used in this book to help you achieve the CCNA Security are designed to also make you much more knowledgeable about how to do your job.

Although this book has more than enough questions to help you prepare for the actual exam, the method in which they are used is not to simply make you memorize as many questions and answers as you possibly can. One key methodology used in this book is to help you discover the exam topics that you need to review in more depth, to help you fully understand and remember those details, and to help you prove to yourself that you have retained your knowledge of those topics. So, this book does not try to help you pass by memorization, but helps you truly learn and understand the topics. The IINS v2.0 exam, which grants the CCNA Security certification, is just one of the foundation topics in the CCNP Security certification, and mastering the knowledge covered by the exam is vitally important to consider yourself a truly skilled security specialist. This book would do you a disservice if it didn't attempt to help you learn the material. To that end, the book will help you pass the CCNA Security exam by using the following methods:

- Helping you discover which test topics you have not mastered
- Providing explanations and information to fill in your knowledge gaps
- Providing practice questions on the topics

Who Should Read This Book?

This book is not designed to be a general security topics book, although it can be used for that purpose. This book is intended to tremendously increase your chances of passing the CCNA Security exam. Although other objectives can be achieved from using this book, the book is written with three goals in mind: to improve your knowledge of Cisco IOS security, to introduce you to Cisco ASA, and to help you pass the CCNA Security exam.

So why should you want to pass the CCNA Security exam? Because it is one of the milestones toward getting the CCNP Security certification, no small feat in itself. What would getting the CCNP Security certification mean to you? A raise, a promotion, or other recognition? A way to enhance your résumé, and demonstrate that you are serious about continuing the learning process and that you are not content to rest on your laurels? A chance to work in one of the most thrilling and fastest growing sectors of IT, network security? An opportunity to please your reseller-employer, who needs more certified employees for a higher discount from Cisco? These are some of many reasons people pursue the CCNP Security certification.

Strategies for Exam Preparation

The strategy you use for CCNA Security might be slightly different from strategies used by other readers, mainly based on the skills, knowledge, and experience you already have obtained. For instance, if you have attended the IINS course, you might take a different approach than someone who learned firewalling via on-the-job training.

The best way to prepare for this exam is to focus on one chapter at a time and take notes. Some chapters are purely theoretical, such as [Chapter 12](#), which introduces cryptography and VPN technologies. Other chapters are more hands-on where configuration of a Cisco router or a Cisco firewall is demonstrated. Ideally, you will practice the suggested configurations on Cisco equipment to sharpen your hands-on skills prior to attempting the IINS v2.0 exam, in order to achieve the CCNA Security certification.

How This Book Is Organized

Although this book could be read cover to cover, it is designed to be flexible and allow you to move between chapters. However, if you do intend to read every chapter, the order in the book is an excellent sequence to use. [Chapters 1](#) to [15](#), separated in four parts, cover the following topics:

- [Chapter 1](#), “[Network Security Concepts and Policies](#)”: This chapter discusses how to develop a comprehensive network security policy to counter threats against information security. It also teaches you about possible threats and how to describe and implement the process of developing a security policy. It covers the identification of common vulnerabilities and threats, mitigation strategies, and the implementation of a security architecture using a lifecycle approach.
- [Chapter 2](#), “[Security Strategy and Cisco Borderless Network](#)”: This chapter discusses the concept of Borderless Networks. It discusses Cisco Borderless Network Architecture, including the components and underlying technologies. You will learn about the Cisco Security portfolio products that address specifically issues of Borderless Networks, and more precisely about Cisco SecureX. This chapter introduces Cisco threat control and containment products and VPN technologies that will be covered in greater detail in subsequent chapters.

- **[Chapter 3, “Network Foundation Protection and Cisco Configuration Professional”](#)**: This chapter deals with Cisco IOS Network Foundation Protection (NFP) as a framework for infrastructure protection, all its components, and commonly used countermeasures as found in Cisco IOS devices. More precisely, this chapter differentiates the security measures to be implemented on the three conceptual planes of Cisco IOS devices: the control plane, the data plane, and the management plane. This chapter also discusses using Cisco Configuration Professional (CCP) to implement security controls on Cisco IOS routers.
- **[Chapter 4, “Securing the Management Plane on Cisco IOS Devices and AAA”](#)**: This chapter describes how to securely implement the management and reporting features of Cisco IOS devices. It discusses technologies surrounding network management, such as syslog, Network Time Protocol, Secure Shell, and Simple Network Management Protocol. It discusses proper password management, the recovery procedure of the configuration file, and the safeguarding of the IOS. This chapter introduces the subject of authentication, authorization, and accounting (AAA) both locally and on an external database, including the Cisco Secure Access Control System (ACS).
- **[Chapter 5, “Securing the Data Plane on Cisco Catalyst Switches”](#)**: This chapter explains how Cisco IOS routers and switches have their own set of network security requirements. It introduces fundamental switching concepts, such as VLANs, trunking, and Spanning Tree, and shows how attackers can exploit vulnerabilities in the switching infrastructure. It then describes a strategy for protecting the switch data plane using port security.
- **[Chapter 6, “Securing the Data Plane in IPv6 Environments”](#)**: This chapter explains the need for IPv6 and presents its fundamental features, as well as enhancements when compared to IPv4. It covers IPv6 addressing scheme, components, and design principles and how routing functions. The chapter then presents potential threats and develops a strategy for IPv6 security.
- **[Chapter 7, “Planning a Threat Control Strategy”](#)**: This chapter suggests design principles to plan a threat control and containment strategy using firewalls and intrusion prevention systems in Cisco IOS environments. This chapter provides a general evaluation of the current state of enterprise security in the presence of evolving threats. It presents the design considerations for a threat protection strategy as part of a risk management strategy with Cisco threat control and containment solutions.
- **[Chapter 8, “Access Control Lists for Threat Mitigation”](#)**: Cisco provides basic traffic filtering capabilities with access control lists (ACL). This chapter covers the benefits of ACLs and describes their building blocks. The chapter describes summarizable address blocks in the context of CIDR and VLSM environments, demonstrating how ACL wildcard masks allow for threat mitigation in those environments. It also demonstrates how to configure ACLs using both CLI and CCP and how to use object groups. ACLs are examined in the context of IPv4 and IPv6.
- **[Chapter 9, “Firewall Fundamentals and Network Address Translation”](#)**: This chapter explains the operations of the different types of firewall technologies and the role they play in network access control and security architectures. It also describes guidelines for

firewall rule set creation. The chapter then describes the function and building blocks of Network Address Translation.

- **[Chapter 10, “Cisco Firewalling Solutions: Cisco IOS Zone-Based Firewall and Cisco ASA”](#)**: This chapter explains the two Cisco Firewall solutions: Cisco IOS Zone-Based Policy Firewalls and Cisco Adaptive Security Appliance. It describes in detail Cisco IOS Zone-Based Policy Firewall, and how the solution uses the Cisco Common Classification Policy Language (C3PL) for creating firewall policies. The chapter then presents the Cisco ASA firewall, identifying key supported features and the building blocks of its configuration using ASDM. The chapter also briefly describes the deployment of policies using the Cisco Modular Policy Framework.
- **[Chapter 11, “Intrusion Prevention Systems”](#)**: This chapter describes the functions and operations of intrusion detection systems (IDS) and intrusion prevention systems (IPS). It explains the underlying IDS and IPS technology embedded in the Cisco IOS IPS solutions. It describe the use of signatures, the need for IPS alarm monitoring, and the design considerations in deploying IPS.
- **[Chapter 12, “Fundamentals of Cryptography and VPN Technologies”](#)**: This chapter introduces the concepts of cryptography and covers encryption, hashing, and digital signatures and how these techniques provide confidentiality, integrity, authenticity, and nonrepudiation. You will learn about algorithms, symmetric and asymmetric encryption, digital signatures, and Public Key Infrastructure (PKI).
- **[Chapter 13, “IPsec Fundamentals”](#)**: This chapter covers the role and operational impact of IPsec’s main components and its modes of operation in various scenarios. It provides a detailed description of the phases of IPsec connectivity. It also provides an overview of IPv6 VPNs.
- **[Chapter 14, “Site-to-Site IPsec VPNs with Cisco IOS Routers”](#)**: This chapter explains how to configure site-to-site virtual private networks (VPN) using Cisco IOS routers. You will learn how to use both CLI commands and Cisco Configuration Professional to configure, validate, and monitor the VPN configuration. You will also learn site-to-site VPN troubleshooting techniques.
- **[Chapter 15, “SSL VPNs with Cisco ASA”](#)**: This chapter describes the use cases and operational requirements of SSL VPNs and offers a detailed presentation on the operations of SSL. The chapter explains configurations, deployment options, and design considerations. It describes the steps to configure both Cisco VPN clientless mode and Cisco full-tunnel mode on Cisco ASA using the Cisco AnyConnect client. The VPN configuration is demonstrated using Cisco ASDM.

Part I: Networking Security Fundamentals

Chapter 1. Network Security Concepts and Policies

In this chapter, you learn how to develop a comprehensive network security policy to counter threats against information security. You also learn about possible threats and how to describe and implement the process of developing a security policy. In this chapter, you learn about the following topics:

- Fundamental concepts in network security, including identification of common vulnerabilities and threats, and mitigation strategies
- Implementation of a security architecture using a lifecycle approach, including the phases of the process, their dependencies, and the importance of a sound security policy

The open nature of the Internet makes it vital for businesses to pay attention to the security of their networks. As companies move more of their business functions to the public network, they need to take precautions to ensure that the data cannot be compromised and that the data is not accessible to anyone who is not authorized to see it.

Unauthorized network access by an outside hacker or a disgruntled employee can cause damage or destruction to proprietary data, negatively affect company productivity, and impede the capability to compete. The Computer Security Institute reported in its *2010/2011 CSI Computer Crime and Security Survey* (available at <http://gocsi.com/survey>) that on an average day, 41.1 percent of respondents dealt with at least one security incident (see page 11 of the survey). Unauthorized network access can also harm relationships with customers and business partners, who might question the capability of a company to protect its confidential information. The definition of “data location” is being blurred by cloud computing services and other service trends. Individuals and corporations benefit from the elastic deployment of services in the cloud, available at all times from any device, but these dramatic changes in the business services industry exacerbate the risks in protecting data and the entities using it (individuals, businesses, governments, and so on). Security policies and architectures require sound principles and a lifecycle approach, including whether the data is in the server farm, mobile on the employee’s laptop, or stored in the cloud.

To start on our network security quest, this chapter examines the need for security, looks at what you are trying to protect, and examines the different trends for attacks and protection and the principles of secure network design. These concepts are important not only for succeeding with the IINS 640-554 exam, but they are fundamentals at all security endeavors on which you will be embarking.

Building Blocks of Information Security

Establishing and maintaining a secure computing environment is increasingly more difficult as networks become increasingly interconnected and data flows ever more freely. In the commercial world, connectivity is no longer optional, and the possible risks of connectivity do not outweigh the benefits. Therefore, it is very important to enable networks to support security services that provide adequate protection to companies that conduct business in a relatively open environment. This section explains the breadth of assumptions and challenges to establish and maintain a secure network environment.

Basic Security Assumptions

Several new assumptions have to be made about computer networks because of their evolution over the years:

- Modern networks are very large, very interconnected, and run both ubiquitous protocols (such as IP) and proprietary protocols. Therefore, they are often open to access, and a potential attacker can with relative ease attach to, or remotely access, such networks. Widespread IP internetworking increases the probability that more attacks will be carried out over large, heavily interconnected networks, such as the Internet.
- Computer systems and applications that are attached to these networks are becoming increasingly complex. In terms of security, it becomes more difficult to analyze, secure, and properly test the security of the computer systems and applications; it is even more so when virtualization is involved. When these systems and their applications are attached to large networks, the risk to computing dramatically increases.

Basic Security Requirements

To provide adequate protection of network resources, the procedures and technologies that you deploy need to guarantee three things, sometimes referred to as the CIA triad:

- **Confidentiality:** Providing confidentiality of data guarantees that only authorized users can view sensitive information.
- **Integrity:** Providing integrity of data guarantees that only authorized users can change sensitive information and provides a way to detect whether data has been tampered with during transmission; this might also guarantee the authenticity of data.
- **Availability of systems and data:** System and data availability provides uninterrupted access by authorized users to important computing resources and data.

When designing network security, a designer must be aware of the following:

- The threats (possible attacks) that could compromise security
- The associated risks of the threats (that is, how relevant those threats are for a particular system)
- The cost to implement the proper security countermeasures for a threat
- A cost versus benefit analysis to determine whether it is worthwhile to implement the security countermeasures

Data, Vulnerabilities, and Countermeasures

Although viruses, worms, and hackers monopolize the headlines about information security, risk management is the most important aspect of security architecture for administrators. A less exciting and glamorous area, risk management is based on specific principles and concepts that are related to asset protection and security management.

An *asset* is anything of value to an organization. By knowing which assets you are trying to protect, as well as their value, location, and exposure, you can more effectively determine the time, effort, and money to spend in securing those assets.

A *vulnerability* is a weakness in a system or its design that could be exploited by a threat. Vulnerabilities are sometimes found in the protocols themselves, as in the case of some security weaknesses in TCP/IP. Often, the vulnerabilities are in the operating systems and applications.

Written security policies might also be a source of vulnerabilities. This is the case when written policies are too lax or are not thorough enough in providing a specific approach or line of conduct to network administrators and users.

A *threat* is any potential danger to assets. A threat is realized when someone or something identifies a specific vulnerability and exploits it, creating exposure. If the vulnerability exists theoretically but has not yet been exploited, the threat is considered latent. The entity that takes advantage of the vulnerability is known as the threat agent or threat vector.

A *risk* is the likelihood that a particular threat using a specific attack will exploit a particular vulnerability of a system that results in an undesirable consequence. Although the roof of the data center might be vulnerable to being penetrated by a falling meteor, for example, the risk is minimal because the likelihood of that threat being realized is negligible.

Note

If you have a vulnerability but there is no threat toward that vulnerability, technically you have no risk.

An *exploit* happens when computer code is developed to take advantage of a vulnerability. For example, suppose that a vulnerability exists in a piece of software, but nobody knows about this vulnerability. Although the vulnerability exists theoretically, there is no exploit yet developed for it. Because there is no exploit, there really is no problem yet.

A *countermeasure* is a safeguard that mitigates a potential risk. A countermeasure mitigates risk either by eliminating or reducing the vulnerability or by reducing the likelihood that a threat agent will be able to exploit the risk.

Key Concepts

An **asset** is anything of value to an organization.

A **vulnerability** is a weakness in a system or its design that could be exploited by a threat.

A **threat** is a potential danger to information or systems.

A **risk** is the likelihood that a particular vulnerability will be exploited.

An **exploit** is an attack performed against a vulnerability.

A **countermeasure** (safeguard) is the protection that mitigates the potential risk.

Data Classification

To optimally allocate resources and secure assets, it is essential that some form of data classification exists. By identifying which data has the most worth, administrators can put their greatest effort toward securing that data. Without classification, data custodians find it almost impossible to adequately secure the data, and IT management finds it equally difficult to optimally

allocate resources.

Sometimes information classification is a regulatory requirement (required by law), in which case there might be liability issues that relate to the proper care of data. By classifying data correctly, data custodians can apply the appropriate confidentiality, integrity, and availability controls to adequately secure the data, based on regulatory, liability, and ethical requirements. When an organization takes classification seriously, it illustrates to everyone that the company is taking information security seriously.

The methods and labels applied to data differ all around the world, but some patterns do emerge. The following is a common way to classify data that many government organizations, including the military, use:

- **Unclassified:** Data that has little or no confidentiality, integrity, or availability requirements and therefore little effort is made to secure it.
- **Restricted:** Data that if leaked could have undesirable effects on the organization. This classification is common among NATO (North Atlantic Treaty Organization) countries but is not used by all nations.
- **Confidential:** Data that must comply with confidentiality requirements. This is the lowest level of classified data in this scheme.
- **Secret:** Data for which you take significant effort to keep secure because its disclosure could lead to serious damage. The number of individuals who have access to this data is usually considerably fewer than the number of people who are authorized to access confidential data.
- **Top secret:** Data for which you make great effort and sometimes incur considerable cost to guarantee its secrecy since its disclosure could lead to exceptionally grave damage. Usually a small number of individuals have access to top-secret data, on condition that there is a need to know.
- **Sensitive But Unclassified (SBU):** A popular classification by government that designates data that could prove embarrassing if revealed, but no great security breach would occur. SBU is a broad category that also includes the For Official Use Only designation.

It is important to point out that there is no actual standard for private-sector classification. Furthermore, different countries tend to have different approaches and labels. Nevertheless, it can be instructive to examine a common, private sector classification scheme:

- **Public:** Companies often display public data in marketing literature or on publicly accessible websites.
- **Sensitive:** Data in this classification is similar to the SBU classification in the government model. Some embarrassment might occur if this data is revealed, but no serious security breach is involved.
- **Private:** Private data is important to an organization. You make an effort to maintain the secrecy and accuracy of this data.
- **Confidential:** Companies make the greatest effort to secure confidential data. Trade secrets and employee personnel files are examples of what a company would commonly

classify as confidential.

Regardless of the classification labeling used, what is certain is that as the security classification of a document increases, the number of staff that should have access to that document should decrease, as illustrated in [Figure 1-1](#).

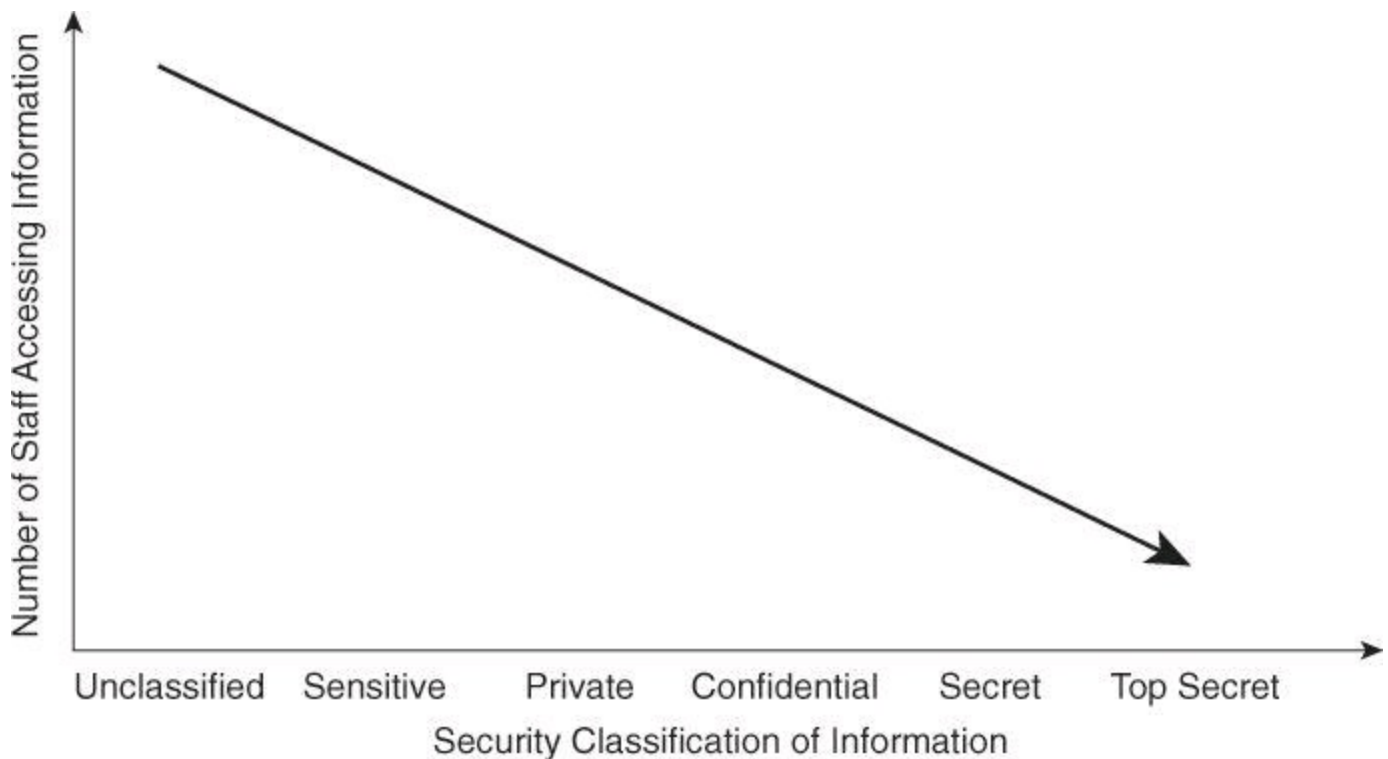


Figure 1-1. Ratio: Staff Access to Information Security Classification

Many factors go into the decision of how to classify certain data. These factors include the following:

- **Value:** Value is the number one criterion. Not all data has the same value. The home address and medical information of an employee is considerably more sensitive (valuable) than the name of the chief executive officer (CEO) and the main telephone number of the company.
- **Age:** For many types of data, its importance changes with time. For example, an army general will go to great lengths to restrict access to military secrets. But after the war is over, the information is gradually less and less useful and eventually is declassified.
- **Useful life:** Often data is valuable for only a set window of time, and after that window has expired, there is no need to keep it classified. An example of this type of data is confidential information about the products of a company. The useful life of the trade secrets of products typically expires when the company no longer sells the product.
- **Personal association:** Data of this type usually involves something of a personal nature. Much of the government data regarding employees is of this nature. Steps are usually taken to protect this data until the person is deceased.

Note

To further understand the value of information, think about the Federal Reserve Bank (commonly called the Fed) and the discount rate it sets. The discount rate is, in essence, the

interest rate charged to commercial banks by the Fed.

Periodically, the Fed announces a new discount rate. Typically, if the rate is higher than the previous rate, the stock market reacts with sell-offs. If the discount rate is lower, the stock market rises.

Therefore, moments before the Fed announces the new discount rate, that information is worth gazillions of dollars. However, the value of this information drops to nothing when it hits the wire, because everyone then has free access to the information.

For a classification system to work, there must be different roles that are fulfilled. The most common of these roles are as follows:

- **Owner:** The owner is the person who is ultimately responsible for the information, usually a senior-level manager who is in charge of a business unit. The owner classifies the data and usually selects custodians of the data and directs their actions. It is important that the owner periodically review the classified data because the owner is ultimately responsible for the data.
- **Custodian:** The custodian is usually a member of the IT staff who has the day-to-day responsibility for data maintenance. Because the owner of the data is not required to have technical knowledge, the owner decides the security controls but the custodian marks the data to enforce these security controls. To maintain the availability of the data, the custodian regularly backs up the data and ensures that the backup media is secure. Custodians also periodically review the security settings of the data as part of their maintenance responsibilities.
- **User:** Users bear no responsibility for the classification of data or even the maintenance of the classified data. However, users do bear responsibility for using the data in accordance with established operational procedures so that they maintain the security of the data while it is in their possession.

Vulnerabilities Classifications

It is also important to understand the weaknesses in security countermeasures and operational procedures. This understanding results in more effective security architectures. When analyzing system vulnerabilities, it helps to categorize them in classes to better understand the reasons for their emergence. You can classify the main vulnerabilities of systems and assets using broad categories:

- Policy flaws
- Design errors
- Protocol weaknesses
- Software vulnerabilities
- Misconfiguration
- Hostile code
- Human factor

This list mentions just a few of the vulnerability categories. For each of these categories, multiple vulnerabilities could be listed.

There are several industry efforts that are aimed at categorizing threats for the public domain. These are some well-known, publicly available catalogs that may be used as templates for vulnerability analysis:

- **Common Vulnerabilities and Exposures (CVE):** A dictionary of publicly known information security vulnerabilities and exposures. It can be found at <http://cve.mitre.org/>. The database provides common identifiers that enable data exchange between security products, providing a baseline index point for evaluating coverage of tools and services.
- **National Vulnerability Database (NVD):** The U.S. government repository of standards-based vulnerability management data. This data enables automation of vulnerability management, security measurement, and compliance. NVD includes databases of security checklists, security-related software flaws, misconfigurations, product names, and impact metrics. The database can be found at <http://nvd.nist.gov>.
- **Common Vulnerability Scoring System (CVSS):** A standard within the computer and networking fields for assessing and classifying security vulnerabilities. This standard is focused on rating a vulnerability compared to others, thus helping the administrator to set priorities. This standard was adopted by significant players in the industry such as McAfee, Qualys, Tenable, and Cisco. More information can be found, including the database and calculator, at <http://www.first.org/cvss>.

Countermeasures Classification

After assets (data) and vulnerabilities, threats are the most important component to understand. Threat classification and analysis, as part of the risk management architecture, will be described later in this chapter.

Once threat vectors are considered, organizations rely on various controls to accomplish in-depth defense as part of their security architecture. There are several ways to classify these security controls; one of them is based on the nature of the control itself. These controls fall into one of three categories:

- **Administrative:** Controls that are largely policies and procedures
- **Technical:** Controls that involve electronics, hardware, software, and so on
- **Physical:** Controls that are mostly mechanical

Later in this chapter, we will discuss models and frameworks from different organizations that can be used to implement network security best practices.

Administrative Controls

Administrative controls are largely policy and procedure driven. You will find many of the administrative controls that help with an enterprise's information security in the human resources department. Some of these controls are as follows:

- Security-awareness training
- Security policies and standards
- Change controls and configuration controls
- Security audits and tests

- Good hiring practices
- Background checks of contractors and employees

For example, if an organization has strict hiring practices that require drug testing and background checks for all employees, the organization will likely hire fewer individuals of questionable character. With fewer people of questionable character working for the company, it is likely that there will be fewer problems with internal security issues. These controls do not single-handedly secure an enterprise, but they are an important part of an information security program.

Technical Controls

Technical controls are extremely important to a good information security program, and proper configuration and maintenance of these controls will significantly improve information security. The following are examples of technical controls:

- Firewalls
- Intrusion prevention systems (IPS)
- Virtual private network (VPN) concentrators and clients
- TACACS+ and RADIUS servers
- One-time password (OTP) solutions
- Smart cards
- Biometric authentication devices
- Network Admission Control (NAC) systems
- Routers with ACLs

Note

This book focuses on technical controls because implementing the Cisco family of security products is the primary topic. However, it is important to remember that a comprehensive security program requires much more than technology.

Physical Controls

While trying to secure an environment with good technical and administrative controls, it is also necessary that you lock the doors in the data center. This is an example of a physical control. Other examples of physical controls include the following:

- Intruder detection systems
- Security guards
- Locks
- Safes
- Racks
- Uninterruptible power supplies (UPS)
- Fire-suppression systems
- Positive air-flow systems

When security professionals examine physical security requirements, life safety (protecting human life) should be their number one concern. Good planning is needed to balance life safety concerns against security concerns. For example, permanently barring a door to prevent unauthorized physical access might prevent individuals from escaping in the event of a fire. By the way, physical security is a field that Cisco entered a few years ago. More information on those products can be found at <http://www.cisco.com/go/physicalsecurity>.

Convergence of Physical and Technical Security

One of the best examples of the convergence of physical and technical security I have witnessed was during a technical visit with a bank in Doha, Qatar, a few weeks before the grand opening of their new head office. They had extensive physical security, using a mix of contactless smart cards and biometrics.

They had cleverly linked the login system for traders to the physical security system. For instance, a trader coming to work in the morning had to use his smart card to enter the building, to activate the turnstile, to call the exact floor where the elevator was to stop, and to be granted access through the glass doors of the trading floors. The movements of the traders were recorded by the physical security systems. Minutes later, upon logging in to perform the first trade of the day, the trading authentication, authorization, and accounting (AAA) system queried the physical security system about the location of the trader. The trader was granted access to the trading system only when the physical security system confirmed to the trading AAA system that the trader was physically on the trading floor.

Controls are also categorized by the type of control they are:

- **Preventive:** The control prevents access.
- **Deterrent:** The control deters access.
- **Detective:** The control detects access.

All three categories of controls can be any one of the three types of controls; for example, a preventive control can be administrative, physical, or technical.

Note

A security control is any mechanism that you put in place to reduce the risk of compromise of any of the three CIA objectives: confidentiality, integrity, and availability.

Preventive controls exist to prevent compromise. This statement is true whether the control is administrative, technical, or physical. The ultimate purpose for these controls is to stop security breaches before they happen.

However, a good security design also prepares for failure, recognizing that prevention will not always work. Therefore, detective controls are also part of a comprehensive security program because they enable you to detect a security breach and to determine how the network was breached. With this knowledge, you should be able to better secure the data the next time.

With effective detective controls in place, the incident response can use the detective controls to

figure out what went wrong, allowing you to immediately make changes to policies to eliminate a repeat of that same breach. Without detective controls, it is extremely difficult to determine what you need to change.

Deterrent controls are designed to scare away a certain percentage of adversaries to reduce the number of incidents. Cameras in bank lobbies are a good example of a deterrent control. The cameras most likely deter at least some potential bank robbers. The cameras also act as a detective control.

Note

To be more concrete, examples of types of physical controls include the following:

- **Preventive:** Locks on doors
 - **Deterrent:** Video surveillance
 - **Detective:** Motion sensor
-
-

Note

It is not always possible to classify a control into only one category or type. Sometimes there is overlap in the definitions, as in the case of the previously mentioned bank lobby cameras. They serve as both deterrent and detective controls.

Need for Network Security

Business goals and risk analysis drive the need for network security. For a while, information security was influenced to some extent by fear, uncertainty, and doubt. Examples of these influences included the fear of a new worm outbreak, the uncertainty of providing web services, or doubts that a particular leading-edge security technology would fail. But we realized that regardless of the security implications, business needs had to come first.

If your business cannot function because of security concerns, you have a problem. The security system design must accommodate the goals of the business, not hinder them. Therefore, risk management involves answering two key questions:

- What does the cost-benefit analysis of your security system tell you?
 - How will the latest attack techniques play out in your network environment?
-

Dealing with Risk

There are actually four ways to deal with risk:

Reduce: This is where we IT managers evolve and it is the main focus of this book. We are responsible for mitigating the risks. Four activities contribute to reducing risks:

- **Limitation/avoidance:** Creating a secure environment by not allowing actions that would cause risks to occur, such as installing a firewall, using encryption systems and strong authentication, and so on
- **Assurance:** Ensuring policies, standards, and practices are followed
- **Detection:** Detecting intrusion attempts and taking appropriate action to terminate the

intrusion

- **Recovery:** Restoring the system to operational state

Ignore: This is not an option for an IT manager. The moment you become aware of a risk, you must acknowledge that risk and decide how to deal with it: accept this risk, transfer this risk, or reduce this risk.

Accept: This means that you document that there is a risk, but take no action to mitigate that risk because the risk is too far-fetched or the mitigation costs are too prohibitive.

Transfer: This is buying insurance against a risk that cannot be eliminated or reduced further.

[Figure 1-2](#) illustrates the key factors you should consider when designing a secure network:

- **Business needs:** What does your organization want to do with the network?
- **Risk analysis:** What is the risk and cost balance?
- **Security policy:** What are the policies, standards, and guidelines that you need to address business needs and risks?
- **Industry best practices:** What are the reliable, well-understood, and recommended security best practices?
- **Security operations:** These operations include incident response, monitoring, maintenance, and auditing the system for compliance.

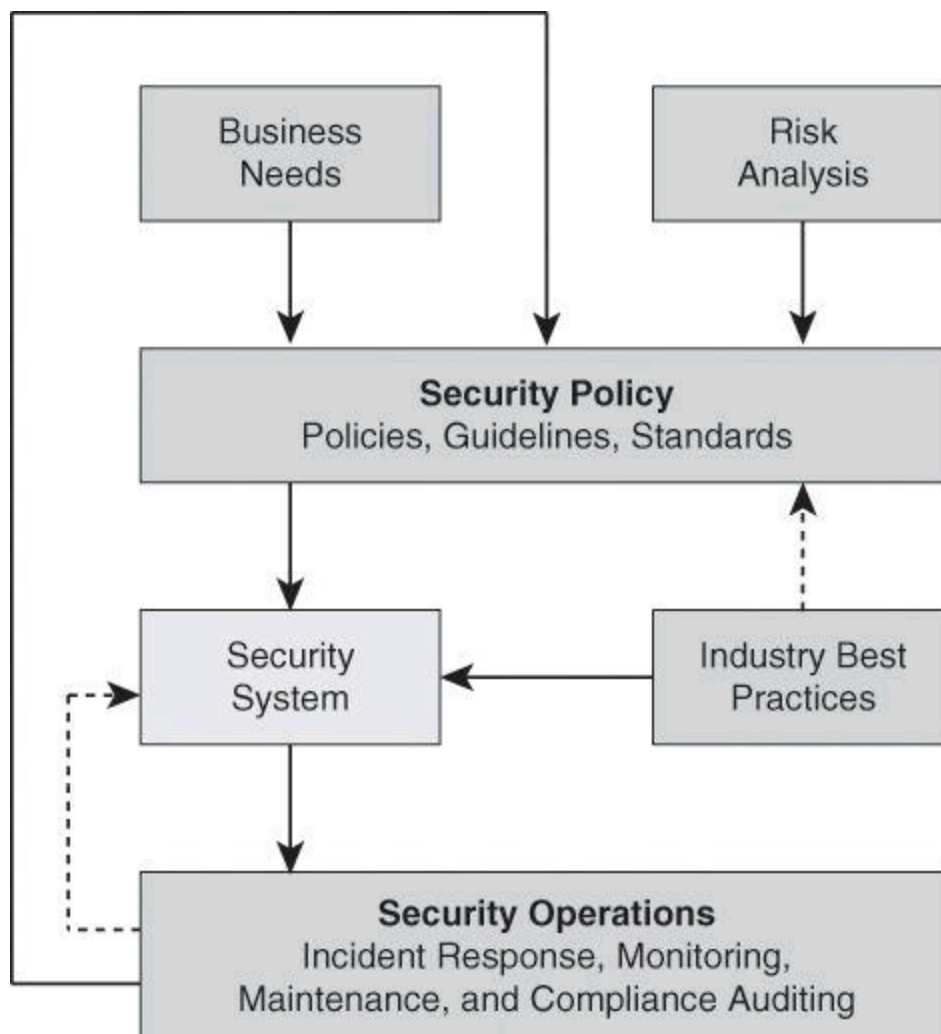


Figure 1-2. Factors Affecting the Design of a Secure Network

Risk management and security policies will be detailed later in this chapter.

Intent Evolution

When viewed from the perspective of motivation intersecting with opportunity, risk management can be driven not only by the techniques or sophistication of the attackers and threat vectors, but also by their motives. Research reveals that hackers are increasingly motivated by profit, where in the past they were motivated by notoriety and fame. In instances of attacks carried out for financial gains, hackers are not looking for attention, which makes their exploits harder to detect. Few signatures exist or will ever be written to capture these “custom” threats. In order to be successful in defending your environments, you must employ a new model to catch threats across the infrastructure.

Attackers are also motivated by government or industrial espionage. The Stuxnet worm, whose earliest versions appear to date to 2009, is an example. This worm differs from its malware “cousins” in that it has a specific, damaging goal: to traverse industrial control systems, such as supervisory control and data acquisition (SCADA) systems, so that it can reprogram the programmable logic controllers, possibly disrupting industrial operations.

This worm was not created to gather credit card numbers to sell off to the highest bidder, or to sell fake pharmaceuticals. This worm appears to have been created solely to invade public or private infrastructure. The cleverness of Stuxnet lies in its ability to traverse non-networked systems, which means that even systems unconnected to networks or the Internet are at risk.

Security experts have called Stuxnet “the smartest malware ever.” This worm breaks the malware mold because it is designed to disrupt industrial control systems in critical infrastructure. This ability should be a concern for every government.

Motivation can also so be political or in the form of vigilantism. Anonymous is currently the best known hacktivist group. As a recent example of its activities, in May 2012, Anonymous attacked the website of the Quebec government after its promulgation of a law imposing new requirements for the right to protest by college and university students.

Threat Evolution

The nature and sophistication of threats, as well as their pervasiveness and global nature, are trends to watch. [Figure 1-3](#) shows how the threats that organizations face have evolved over the past few decades, and how the growth rate of vulnerabilities that are reported in operating systems and applications is rising. The number and variety of viruses and worms that have appeared over the past three years is daunting, and their rate of propagation is frightening. There have been unacceptable levels of business outages and expensive remediation projects that consume staff, time, and funds that were not originally budgeted for such tasks.

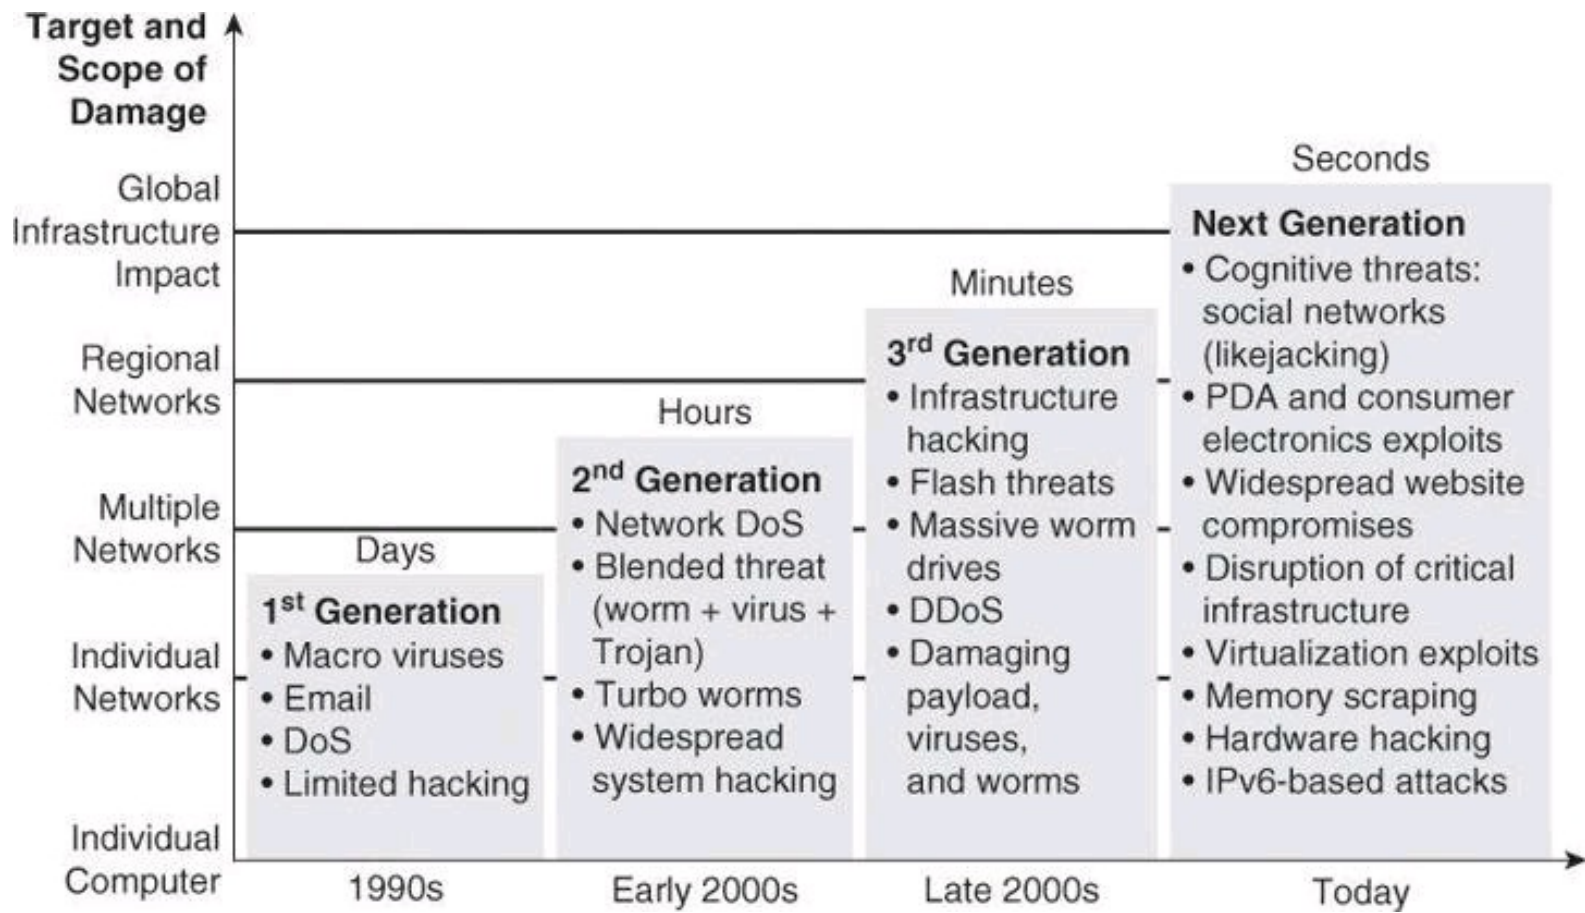


Figure 1-3. Shrinking Time Frame from Knowledge of Vulnerability to Release of Exploits

New exploits are designed to have global impact in minutes. Blended threats, which use multiple means of propagation, are more sophisticated than ever. The trends are becoming regional and global in nature. Early attacks affected single systems or one organization network, while attacks that are more recent are affecting entire regions. For example, attacks have expanded from individual denial of service (DoS) attacks from a single attacker against a single target, to large-scale distributed DoS (DDoS) attacks emanating from networks of compromised systems that are known as botnets.

Threats are also becoming persistent. After an attack starts, attacks may appear in waves as infected systems join the network. Because infections are so complex and have so many end users (employees, vendors, and contractors), multiple types of endpoints (company desktop, home, and server), and multiple types of access (wired, wireless, VPN, and dial-up), infections are difficult to eradicate.

More recent threat vectors are increasingly sophisticated, and the motivation of the attackers is reflected in their impact. Recent threat vectors include the following:

- **Cognitive threats via social networks (likejacking):** Social engineering takes a new meaning in the era of social networking. From phishing attacks that target social network accounts of high-profile individuals, to information exposure due to lack of policy, social networks have become a target of choice for malicious attackers.
- **PDA and consumer electronics exploits:** The operating systems on consumer devices (smartphones, PDAs, and so on) are an option of choice for high-volume attacks. The proliferation of applications for these operating systems, and the nature of the development and certification processes for those applications, augments the problem.

- **Widespread website compromises:** Malicious attackers compromise popular websites, making the sites download malware to connecting users. Attackers typically are not interested in the data on the website, but use it as a springboard to infect the users of the site.
- **Disruption of critical infrastructure:** The Stuxnet malware, which exploits holes in Windows systems and targets a specific Siemens supervisory control and data acquisition (SCADA) program with sabotage, confirmed concerns about an increase in targeted attacks aimed at the power grid, nuclear plants, and other critical infrastructure.
- **Virtualization exploits:** Device and service virtualization add more complexity to the network. Attackers know this and are increasingly targeting virtual servers, virtual switches, and trust relationships at the hypervisor level.
- **Memory scraping:** Increasingly popular, this technique is aimed at fetching information directly from volatile memory. The attack tries to exploit operating systems and applications that leave traces of data in memory. Attacks are particularly aimed at encrypted information that may be processed as unencrypted in volatile memory.
- **Hardware hacking:** These attacks are aimed at exploiting the hardware architecture of specific devices, with consumer devices being increasingly popular. Attack methods include bus sniffing, altering firmware, and memory dumping to find crypto keys.
- **IPv6-based attacks:** These attacks could become more pervasive as the migration to IPv6 becomes widespread. Attackers are focusing initially on covert channels through various tunneling techniques, and man-in-the-middle attacks leverage IPv6 to exploit IPv4 in dual-stack deployments.

Trends Affecting Network Security

Other trends in business, technology, and innovation influence the need for new paradigms in information security. Mobility is one trend. Expect to see billions of new network mobile devices moving into the enterprise worldwide over the next few years. Taking into consideration constant reductions and streamlining in IT budgets, organizations face serious challenges in supporting a growing number of mobile devices at a time when their resources are being reduced.

The second market transition is cloud computing and cloud services. Organizations of all kinds are taking advantage of offerings such as Software as a Service (SaaS) and Infrastructure as a Service (IaaS) to reduce costs and simplify the deployment of new services and applications.

These cloud services add challenges in visibility (how do you identify and mitigate threats that come to and from a trusted network?), control (who controls the physical assets, encryption keys, and so on?), and trust (do you trust cloud partners to ensure that critical application data is still protected when it is off the enterprise network?).

The third market transition is about changes to the workplace experience. Borders are blurring in the organization between consumers and workers and between the various functions within the organization. The borders between the company and its partners, customers, and suppliers, are also fading. As a result, the network is experiencing increasing demand to connect anyone, any device, anywhere, at any time.

These changes represent a challenge to security teams within the organization. These teams now

need to manage noncontrolled consumer devices, such as a personal tablet, coming into the network, and provide seamless and context-aware services to users all over the world. The location of the data and services accessed by the users is almost irrelevant. The data could be internal to the organization or it could be in the cloud. This situation makes protecting data and services a challenging proposition.

Note

Readers interested in staying current with Network Security trends and technologies could subscribe to some of the numerous podcasts available on iTunes, such as:

- Cisco Interactive Network TechWiseTV
 - Security Now!
 - Security Wire Weekly
 - Silver Bullet Security
 - Crypto-Gram Security
-

Attacks are increasingly politically and financially motivated, driven by botnets, and aimed at critical infrastructure; for example:

- Botnets are used for spam, data theft, mail relays, or simply for denial-of-service attacks (ref: <http://en.wikipedia.org/wiki/Botnet>).
- Zeus botnets reached an estimated 3.6 million *bots*, infected workstations, or “zombies” (ref: <http://www.networkworld.com/news/2009/072209-botnets.html>).
- Stuxnet was aimed at industrial systems.
- Malware is downloaded inadvertently from online marketplaces.

One of the trends in threats is the exploitation of trust. Whether they are creating malware that can subvert industrial processes or tricking social network users into handing over login and password information, cybercriminals have a powerful weapon at their disposal: the exploitation of trust. Cybercriminals have become skilled at convincing users that their infected links and URLs are safe to click, and that they are someone the user knows and trusts. Hackers exploit the trust we have in TinyURLs and in security warning banners. With stolen security credentials, cybercriminals can freely interact with legitimate software and systems.

Nowhere is this tactic more widespread than within social networking, where cybercriminals continue to attract victims who are willing to share information with people they believe are known to them, with malware such as Koobface. One noticeable shift in social engineering is that criminals are spending more time figuring out how to assume someone’s identity, perhaps by generating emails from an individual’s computer or social networking account. A malware-laden email or scam sent by a “trusted person” is more likely to elicit a click-through response than the same message sent by a stranger.

Threats originating from countries outside of the United States are rapidly increasing. Global annual spam volumes actually dropped in 2010, the first time this has happened in the history of the Internet. However, spammers are originating in increasingly varied locations and countries.

Money muling is the practice of hiring individuals as “mules,” recruited by handlers or “wranglers” to set up bank accounts, or even use their own bank accounts, to assist in the transfer of money from the account of a fraud victim to another location, usually overseas, via a wire transfer or automated clearing house (ACH) transaction. Money mule operations often involve individuals in multiple countries.

Web malware is definitely on the rise. The number of distinct domains that are compromised to download malware to connecting users is increasing dramatically. The most dangerous aspect of this type of attack is the fact that users do not need to do much to get infected. Many times, the combination of malware on the website and vulnerabilities on web browsers is enough to provoke infection just by connecting to the website. The more popular the site, the higher the volume of potential infection.

Recently there have been major shifts in the compliance landscape. Although enforcement of existing regulations has been weak in many jurisdictions worldwide, regulators and standards bodies are now tightening enforcement through expanded powers, higher penalties, and harsh enforcement actions. In the future it will be more difficult to hide failures in information security wherever organizations do business. Legislators are forcing transparency through the introduction of breach notification laws in Europe, Asia, and North America as data breach disclosure becomes a global principle.

As more regulations are introduced, there is a trend toward increasingly prescriptive rules. For example, recent amendments introduced in the United Kingdom in 2011 bring arguably more prescriptive information protection regulations to the Privacy and Electronic Communications Directive. Such laws are discussed in more detail later in this chapter. Any global enterprise that does business in the United Kingdom today will likely be covered by these regulations. Lately, regulators are also making it clear that enterprises are responsible for ensuring the protection of their data when it is being processed by a business partner, including cloud service providers. The new era of compliance creates formidable challenges for organizations worldwide.

For many organizations, stricter compliance could help focus management attention on security, but if managers take a “check-list approach” to compliance, it will detract from actually managing risk and may not improve security. The new compliance landscape will increase costs and risks. For example, it takes time and resources to substantiate compliance. Increased requirements for service providers give rise to more third-party risks.

With more transparency, there are now greater consequences for data breaches. For example, expect to see more litigation as customers and business partners seek compensation for compromised data. But the harshest judgments will likely come from the court of public opinion, with the potential to permanently damage an enterprise’s reputation.

The following are some of the U.S. and international regulations that many companies are subject to:

- Sarbanes-Oxley (SOX)
- Federal Information Security Management Act (FISMA)
- Gramm-Leach-Bliley Act (GLBA)
- Payment Card Industry Data Security Standard (PCI DSS)

- Health Insurance Portability and Accountability Act (HIPAA)
- Digital Millennium Copyright Act (DMCA)
- Personal Information Protection and Electronic Documents Act (PIPEDA) in Canada
- European Union Data Protection Directive (EU 95/46/EC)
- Safe Harbor Act - European Union and United States
- International Convergence of Capital Measurement and Capital Standards (Basel II)

The challenge becomes to comply with these regulations and, at the same time, make that compliance translate into an effective security posture.

Adversaries, Methodologies, and Classes of Attack

Who are hackers? What motivates them? How do they conduct their attacks? How do they manage to breach the measures we have in place to ensure confidentiality, integrity, and availability? Which best practices can we adopt to defeat hackers? These are some of the questions we try to answer in this section.

People are social beings, and it is quite common for systems to be compromised through social engineering. Harm can be caused by people just trying to be “helpful.” For example, in an attempt to be helpful, people have been known to give their passwords over the phone to attackers who have a convincing manner and say they are troubleshooting a problem and need to test access using a real user password. End users must be trained, and reminded, that the ultimate security of a system depends on their behavior.

Of course, people often cause harm within organizations intentionally: most security incidents are caused by insiders. Thus, strong internal controls on security are required, and special organizational practices might need to be implemented.

An example of a special organizational practice that helps to provide security is the separation of duty, where critical tasks require two or more persons to complete them, thereby reducing the risk of insider threat. People are less likely to attack or misbehave if they are required to cooperate with others.

Unfortunately, users frequently consider security too difficult to understand. Software often does not make security options or decisions easy for end users. Also, users typically prefer “whatever” functionality to no functionality. Implementation of security measures should not create an internally generated DoS, meaning, if security is too stringent or too cumbersome for users, either they will not have access to all the resources needed to perform their work or their performance will be hindered by the security operations.

Adversaries

To defend against attacks on information and information systems, organizations must begin to define the threat by identifying potential adversaries. These adversaries can include the following:

- Nations or states
- Terrorists
- Criminals
- Hackers

- Corporate competitors
- Disgruntled employees
- Government agencies, such as the National Security Agency (NSA) and the Federal Bureau of Investigations (FBI)

Hackers comprise the most well-known outside threat to information systems. They are not necessarily geniuses, but they are persistent people who have taken a lot of time to learn their craft.

Many titles are assigned to hackers:

- **Hackers:** Hackers are computer enthusiasts who break into networks and systems to learn more about them. Some hackers generally mean no harm and do not expect financial gain. Unfortunately, hackers may unintentionally pass valuable information on to people who do intend to harm the system. Hackers are subdivided into the following categories:
 - White hat (ethical hacker)
 - Blue hat (bug tester)
 - Gray hat (ethically questionable hacker)
 - Black hat (unethical hacker)
- **Crackers (criminal hackers):** Crackers are hackers with a criminal intent to harm information systems. Crackers are generally working for financial gain and are sometimes called black hat hackers.
- **Phreakers (phone breakers):** Phreakers pride themselves on compromising telephone systems. Phreakers reroute and disconnect telephone lines, sell wiretaps, and steal long-distance services.

Note

When describing individuals whose intent is to exploit a network maliciously, these individuals are often incorrectly referred to as hackers. In this section, the term hacker is used, but might refer to someone more correctly referred to as a cracker, or black hat hacker.

-
- **Script kiddies:** Script kiddies think of themselves as hackers, but have very low skill levels. They do not write their own code; instead, they run scripts written by other, more skilled attackers.
 - **Hactivists:** Hactivists are individuals who have a political agenda in doing their work. When government websites are defaced, this is usually the work of a hactivist.

Methodologies

The goal of any hacker is to compromise the intended target or application. Hackers begin with little or no information about the intended target, but by the end of their analysis, they have accessed the network and have begun to compromise their target. Their approach is usually careful and methodical, not rushed and reckless. The seven-step process that follows is a good representation of the methods that hackers use:

Step 1. Perform footprint analysis (reconnaissance).

Step 2. Enumerate applications and operating systems.

Step 3. Manipulate users to gain access.

Step 4. Escalate privileges.

Step 5. Gather additional passwords and secrets.

Step 6. Install back doors.

Step 7. Leverage the compromised system.

Caution

Hackers have become successful by thinking “outside the box.” This methodology is meant to illustrate the steps that a structured attack might take. Not all hackers will follow these steps in this order.

To successfully hack into a system, hackers generally first want to know as much as they can about the system. Hackers can build a complete profile, or “footprint,” of the company security posture. Using a range of tools and techniques, an attacker can discover the company domain names, network blocks, IP addresses of systems, ports and services that are used, and many other details that pertain to the company security posture as it relates to the Internet, an intranet, remote access, and an extranet. By following some simple advice, network administrators can make footprinting more difficult.

After hackers have completed a profile, or footprint, of your organization, they use tools such as those in the list that follows to enumerate additional information about your systems and networks. All these tools are readily available to download, and the security staff should know how these tools work. Additional tools (introduced later in the “[Security Testing Techniques](#)” section) can also be used to gather information and therefore hack.

- **Netcat:** Netcat is a featured networking utility that reads and writes data across network connections.
- **Microsoft EPDump and Microsoft Remote Procedure Call (RPC) Dump:** These tools provide information about Microsoft RPC services on a server.
- **GetMAC:** This application provides a quick way to find the MAC (Ethernet) layer address and binding order for a computer running Microsoft Windows locally or across a network.
- **Software development kits (SDK):** SDKs provide hackers with the basic tools that they need to learn more about systems.

Another common technique that hackers use is to manipulate users of an organization to gain access to that organization. There are countless cases of unsuspecting employees providing information to unauthorized people simply because the requesters appear innocent or to be in a position of authority. Hackers find names and telephone numbers on websites or domain registration records by footprinting. Hackers then directly contact these people by phone and convince them to reveal passwords. Hackers gather information without raising any concern or suspicion. This form of attack is called *social engineering*. One form of a social engineering attack is for the hacker to pose as a

visitor to the company, a delivery person, a service technician, or some other person who might have a legitimate reason to be on the premises and, after gaining entrance, walk by cubicles and look under keyboards to see whether anyone has put a note there containing the current password.

The next thing the hacker typically does is review all the information that they have collected about the host, searching for usernames, passwords, and Registry keys that contain application or user passwords. This information can help hackers escalate their privileges on the host or network. If reviewing the information from the host does not reveal useful information, hackers may launch a Trojan horse attack in an attempt to escalate their privileges on the host. This type of attack usually means copying malicious code to the user system and giving it the same name as a frequently used piece of software.

After the hacker has obtained higher privileges, the next task is to gather additional passwords and other sensitive data. The targets now include such things as the local security accounts manager database or the Active Directory of a domain controller. Hackers use legitimate tools such as pwdump and lsadump applications to gather passwords from machines running Windows, which then can be cracked with the very popular Cain & Abel software tool. By cross-referencing username and password combinations, the hacker is able to obtain administrative access to all the computers in the network.

If hackers are detected trying to enter through the “front door,” or if they want to enter the system without being detected, they try to use “back doors” into the system. A back door is a method of bypassing normal authentication to secure remote access to a computer while attempting to remain undetected. The most common backdoor point is a listening port that provides remote access to the system for users (hackers) who do not have, or do not want to use, access or administrative privileges.

After hackers gain administrative access, they enjoy hacking other systems on the network. As each new system is hacked, the attacker performs the steps that were outlined previously to gather additional system and password information. Hackers try to scan and exploit a single system or a whole set of networks and usually automate the whole process.

In addition, hackers will cover their tracks either by deleting log entries or falsifying them.

Thinking Outside the Box

In 2005, David Sternberg hacked the Postal Bank in Israel by physically breaking into one of the bank’s branches in Haifa and connecting a wireless access point in the branch’s IT infrastructure. Sternberg rented office space about 100 feet from the bank and proceeded to transfer funds to bank accounts in his name or in friends’ names.

So instead of trying for months to break into the IT security of the bank, Sternberg thought outside of the box and broke through physical security to gain access to the IT system.

Sternberg was discovered when bank auditors noticed regular transfers from the main bank account to the same individual accounts.

I guess that Sternberg had not heard about the security axiom that says “predictability is the enemy of security.”

A common thread in infosec forums is that information security specialists must patch all

security holes in a network—a hacker only has to find the one that wasn't patched. Security is like a chain. It is only as strong as its weakest link.

Threats Classification

In classifying security threats, it is common to find general categories that resemble the perspective of the attacker and the approaches that are used to exploit software. Attack patterns are a powerful mechanism to capture and communicate the perspective of the attacker. These patterns are descriptions of common methods for exploiting vulnerabilities. The patterns derive from the concept of design patterns that are applied in a destructive rather than constructive context and are generated from in-depth analysis of specific, real-world exploit examples. The following list illustrates examples of threat categories that are based on this criterion. Notice that some threats are not malicious attacks. Examples of nonmalicious threats include forces of nature such as hurricanes and earthquakes.

Later in this chapter, you learn about some of the general categories under which threats can be regrouped, such as:

- Enumeration and fingerprinting
- Spoofing and impersonation
- Man-in-the-middle
- Overt and covert channels
- Blended threats and malware
- Exploitation of privilege and trust
- Confidentiality
- Password attacks
- Availability attacks
 - Denial of service (DoS)
 - Botnet
- Physical security attacks
- Forces of nature

To assist in enhancing security throughout the security lifecycle, there are many publicly available classification databases that provide a catalog of attack patterns and classification taxonomies. They are aimed at providing a consistent view and method for identifying, collecting, refining, and sharing attack patterns for specific communities of interest. The following are four of the most prominent databases:

- **Common Attack Pattern Enumeration and Classification (CAPEC):** Sponsored by the U.S. Department of Homeland Security as part of the software assurance strategic initiative of the National Cyber Security Division, the objective of this effort is to provide a publicly available catalog of attack patterns along with a comprehensive schema and classification taxonomy. More information can be found at <http://capec.mitre.org>.
- **Open Web Application Security Project (OWASP) Application Security Verification**

Standard (ASVS): OWASP is a not-for-profit worldwide charitable organization focused on improving the security of application software. The primary objective of ASVS is to normalize the range in the coverage and level of rigor available in the market when it comes to performing web application security verification using a commercially workable open standard. More information can be found at <https://www.owasp.org>.

- **Web Application Security Consortium Threat Classification (WASC TC):** Sponsored by the WASC, this is a cooperative effort to clarify and organize the threats to the security of a website. The project is aimed at developing and promoting industry-standard terminology for describing these issues. Application developers, security professionals, software vendors, and compliance auditors have the ability to access a consistent language and definitions for web security-related issues. More information can be found at <http://www.webappsec.org>.

- **Malware Attribute Enumeration and Characterization (MAEC):** Created by MITRE, this effort is international in scope and free for public use. MAEC is a standardized language for encoding and communicating high-fidelity information about malware based on attributes such as behaviors, artifacts, and attack patterns. More information can be found at <http://maec.mitre.org>.

Enumeration and Fingerprinting with Ping Sweeps and Port Scans

Enumeration and fingerprinting are types of attacks that use legitimate tools for illegitimate purposes. Some of the tools, such as port-scan and ping-sweep applications, run a series of tests against hosts and devices to identify vulnerable services that need attention. IP addresses and port or banner data from both TCP and UDP ports are examined to gather information.

In an illegitimate situation, a port scan is a series of messages sent by someone attempting to break into a computer to learn which computer network services (each service is associated with a well-known port number) the computer provides. Port scanning can be automated to scan a range of TCP or UDP port numbers on a host to detect listening services. Port scanning, a favorite computer hacker approach, provides information to the hacker about where to probe for weaknesses. Essentially, a port scan consists of sending a message to each port, one at a time. The kind of response received indicates whether the port is being used and needs further probing.

A ping sweep, also known as an Internet Control Message Protocol (ICMP) sweep, is a basic network-scanning technique that is used to determine which IP addresses map to live hosts (computers). A ping sweep consists of ICMP echo-requests (pings) sent to multiple hosts, whereas a single ping consists of ICMP echo-requests that are sent to one specific host computer. If a given address is live, that host returns an ICMP echo-reply. The goal of the ping sweep is to find hosts available on the network to probe for vulnerabilities. Ping sweeps are among the oldest and slowest methods that are used to scan a network.

IP Spoofing Attacks

The prime goal of an IP spoofing attack is to establish a connection that allows the attacker to gain root access to the host and to create a backdoor entry path into the target system.

IP spoofing is a technique used to gain unauthorized access to computers whereby the intruder sends messages to a computer with an IP address that indicates the message is coming from a trusted

host. The attacker learns the IP address of a trusted host and modifies the packet headers so that it appears that the packets are coming from that trusted host.

At a high level, the concept of IP spoofing is easy to comprehend. Routers determine the best route between distant computers by examining the destination address, and ignore the source address. In a spoofing attack, an attacker outside your network pretends to be a trusted computer by using a trusted internal or external IP address.

If an attacker manages to change the routing tables to divert network packets to the spoofed IP address, the attacker can receive all the network packets addressed to the spoofed address and reply just as any trusted user can.

IP spoofing can also provide access to user accounts and passwords. For example, an attacker can emulate one of your internal users in ways that prove embarrassing for your organization. The attacker could send email messages to business partners that appear to have originated from someone within your organization. Such attacks are easier to perpetrate when an attacker has a user account and password, but they are also possible when attackers combine simple spoofing attacks with their knowledge of messaging protocols.

A rudimentary use of IP spoofing also involves bombarding a site with IP packets or ping requests, spoofing a source, a third-party registered public address. When the destination host receives the requests, it responds to what appears to be a legitimate request. If multiple hosts are attacked with spoofed requests, their collective replies to the third-party spoofed IP address create an unsupportable flood of packets, thus creating a DoS attack.

Technical Discussion of IP Spoofing

TCP/IP works at Layer 3 and Layer 4 of the Open Systems Interconnection (OSI) model, IP at Layer 3 and TCP at Layer 4. IP is a connectionless model, which means that packet headers do not contain information about the transaction state that is used to route packets on a network. There is no method in place to ensure proper delivery of a packet to the destination, since at Layer 3, there is no acknowledgement sent back to the source by the destination once it has received the packet.

The IP header contains the source and destination IP addresses. Using one of several tools, an attacker can easily modify the source address field. Note that in IP each datagram is independent of all others because of the stateless nature of IP. To engage in IP spoofing, hackers find the IP address of a trusted host and modify their own packet headers to appear as though packets are coming from that trusted host (source address).

TCP uses a connection-oriented design. This design means that the participants in a TCP session must first build a connection using the three-way handshake, as shown in [Figure 1-4](#).




Figure 1-4. TCP Three-Way Handshake

After the connection is established, TCP ensures data reliability by applying the same process to every packet as the two machines update one another on progress. The sequence and acknowledgments take place as follows:

1. The client selects and transmits an initial sequence number.
2. The server acknowledges the initial sequence number and sends its own sequence number.
3. The client acknowledges the server sequence number, and the connection is open to data transmission.

Sequence Prediction

The basis of IP spoofing during a TCP communication lies in an inherent security weakness known as sequence prediction. Hackers can guess or predict the TCP sequence numbers that are used to construct a TCP packet without receiving any responses from the server. Their prediction allows them to spoof a trusted host on a local network. To mount an IP spoofing attack, the hacker listens to communications between two systems. The hacker sends packets to the target system with the source IP address of the trusted system, as shown in [Figure 1-5](#).

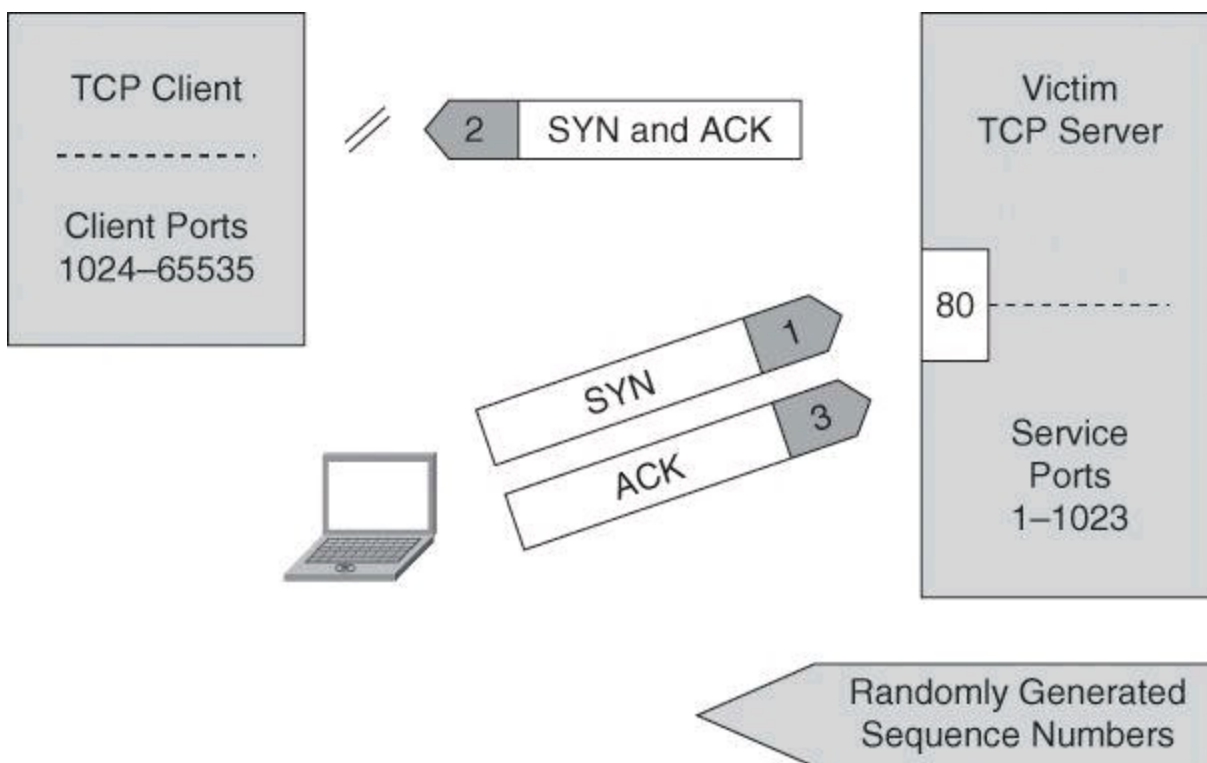


Figure 1-5. Sequence Number Prediction

If the packets from the hacker have the sequence numbers that the target system is expecting, and if these packets arrive before the packets from the real, trusted system, the hacker becomes the trusted host.

To engage in IP spoofing, hackers must first use a variety of techniques to find an IP address of a trusted host and then modify their packet headers to appear as though packets are coming from that trusted host. Further, the attacker can engage other unsuspecting hosts to generate traffic that appears as though it too is coming from the trusted host, thus flooding the network.

Trust Exploitation

Trust exploitation refers to an individual taking advantage of a trust relationship within a network.

As an example of trust exploitation, consider the network shown in [Figure 1-6](#), where system A is in the demilitarized zone (DMZ) of a firewall. System B, located in the inside of the firewall, trusts System A. When a hacker on the outside network compromises System A in the DMZ, the attacker can leverage the trust relationship it has to gain access to System A.




Figure 1-6. Trust Exploitation

A DMZ can be seen as a semi-secure segment of your network. A DMZ is typically used to provide to outside users access to corporate resources, because these users are not allowed to reach inside servers directly. However, a DMZ server might be allowed to reach inside resources directly. In a trust exploitation attack, a hacker could hack a DMZ server and use it as a springboard to reach the inside network.

Several trust models may exist in a network:

- Windows
 - Domains
 - Active Directory
- Linux and UNIX
 - Network File System (NFS)
 - Network Information Services Plus (NIS+)

Password Attacks

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

To execute a brute-force attack, an attacker can use a program that runs across the network and attempts to log in to a shared resource, such as a server. When an attacker gains access to a resource, the attacker has the same access rights as the rightful user. If this account has sufficient privileges, the attacker can create a back door for future access, without concern for any status and password changes to the compromised user account.

Just as with packet sniffers and IP spoofing attacks, a brute-force password attack can provide access to accounts that attackers then use to modify critical network files and services. For example, an attacker compromises your network integrity by modifying your network routing tables. This trick reroutes all network packets to the attacker before transmitting them to their final destination. In such

a case, an attacker can monitor all network traffic, effectively becoming a man in the middle.

Passwords present a security risk if they are stored as plain text. Thus, passwords must be encrypted in order to avoid risks. On most systems, passwords are processed through an encryption algorithm that generates a one-way hash on passwords. You cannot reverse a one-way hash back to its original text. Most systems do not decrypt the stored password during authentication; they store the one-way hash. During the login process, you supply an account and password, and the password encryption algorithm generates a one-way hash. The algorithm compares this hash to the hash stored on the system. If the hashes are the same, the algorithm assumes that the user supplied the proper password.

Remember that passing the password through an algorithm results in a password hash. The hash is not the encrypted password, but rather a result of the algorithm. The strength of the hash is such that the hash value can be re-created only by using the original user and password information, and that it is impossible to retrieve the original information from the hash. This strength makes hashes perfect for encoding passwords for storage. In granting authorization, the hashes, rather than the plain-text password, are calculated and compared.

Hackers use many tools and techniques to crack passwords:

- **Word lists:** These programs use lists of words, phrases, or other combinations of letters, numbers, and symbols that computer users often use as passwords. Hackers enter word after word at high speed (called a *dictionary attack*) until they find a match.
- **Brute force:** This approach relies on power and repetition. It compares every possible combination and permutation of characters until it finds a match. Brute force eventually cracks any password, but it might take a long, long time. Brute force is an extremely slow process because it uses every conceivable character combination.
- **Hybrid crackers:** Some password crackers mix the two techniques. This combines the best of both methods and is highly effective against poorly constructed passwords.

Password cracking attacks any application or service that accepts user authentication, including the following:

- NetBIOS over TCP (TCP 139)
- Direct host (TCP 445)
- FTP (TCP 21)
- Telnet (TCP 23)
- Simple Network Management Protocol (SNMP) (UDP 161)
- Point-to-Point Tunneling Protocol (PPTP) (TCP 1723)
- Terminal services (TCP 3389)

Note

RainbowCrack is a compilation of hashes that provides crackers with a list that they can use to attempt to match hashes that they capture with sniffers.

Confidentiality breaches can occur when an attacker attempts to obtain access to read-sensitive data. These attacks can be extremely difficult to detect because the attacker can copy sensitive data without the knowledge of the owner and without leaving a trace.

A confidentiality breach can occur simply because of incorrect file protections. For instance, a sensitive file could mistakenly be given global read access. Unauthorized copying or examination of the file would probably be difficult to track without having some type of audit mechanism running that logs every file operation. If a user had no reason to suspect unwanted access, however, the audit file would probably never be examined.

In [Figure 1-7](#), the attacker is able to compromise an exposed web server. Using this server as a beachhead, the attacker then gains full access to the database server from which customer data is downloaded. The attacker then uses information from the database, such as a username, password, and email address, to intercept and read sensitive email messages destined for a user in the branch office. This attack is difficult to detect because the attacker did not modify or delete any data. The data was only read and downloaded. Without some kind of auditing mechanism on the server, it is unlikely that this attack will be discovered.

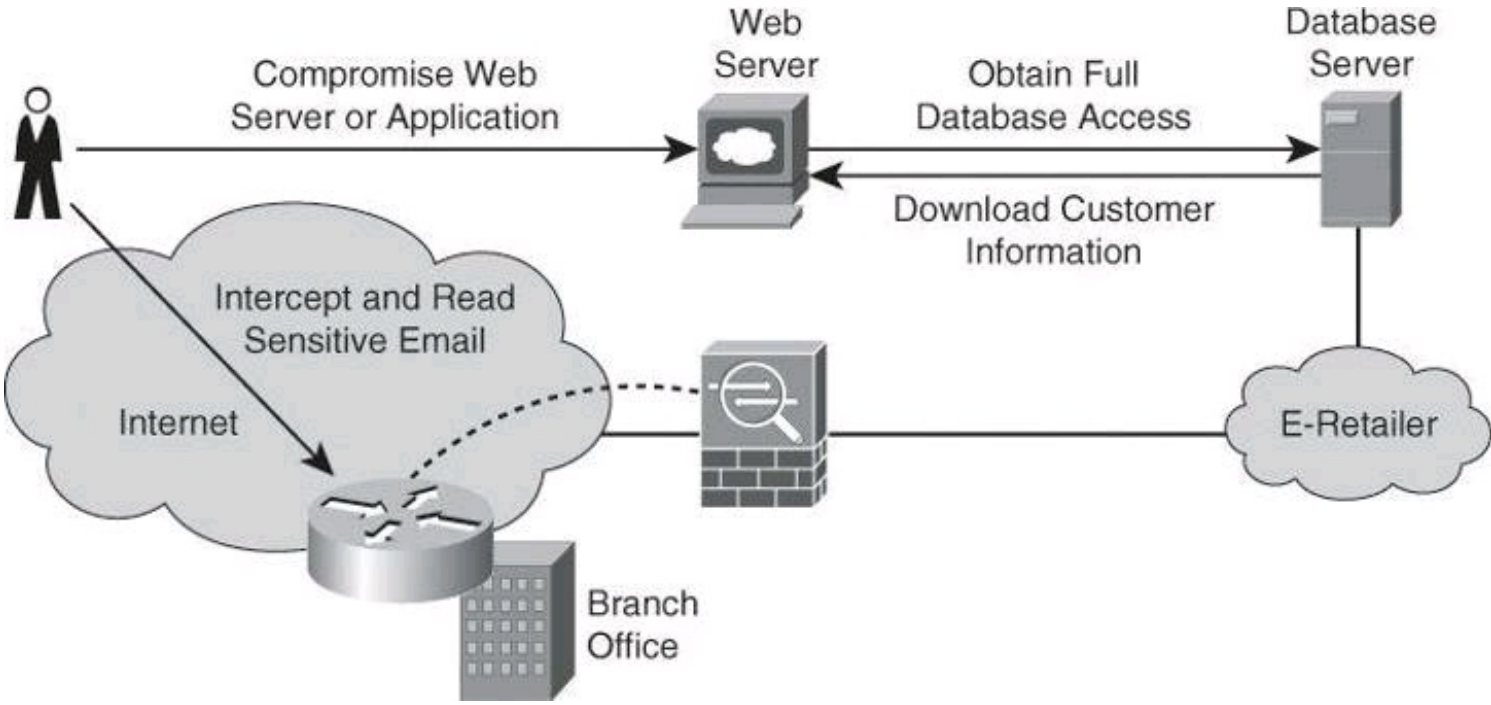


Figure 1-7. Breach of Confidentiality

Attackers can use many methods to compromise confidentiality, the most common of which are as follows:

- **Ping sweeps and port scanning:** Searching a network host for open ports.
- **Packet sniffing:** Intercepting and logging traffic that passes over a digital network or part of a network.
- **Emanations capturing:** Capturing electrical transmissions from the equipment of an organization to deduce information regarding the organization.
- **Overt channels:** Listening on obvious and visible communications. Overt channels can be used for covert communication.
- **Covert channels:** Hiding information within a transmission channel that is based on

encoding data using another set of events.

- **Wiretapping:** Monitoring the telephone or Internet conversations of a third party, often covertly.
- **Social engineering:** Using social skills or relationships to manipulate people inside the network to provide the information needed to access the network.
- **Dumpster diving:** Searching through company dumpsters or trash cans looking for information, such as phone books, organization charts, manuals, memos, charts, and other documentation that can provide a valuable source of information for hackers.
- **Phishing:** Attempting to criminally acquire sensitive information, such as usernames and passwords, by masquerading as trustworthy entities.
- **Pharming:** Redirecting the traffic of a website to another, rogue website.

Many of these methods are used to compromise more than confidentiality. They are often elements of attacks on integrity and availability.

Man-in-the-Middle Attacks

A complex form of IP spoofing is called man-in-the-middle attack, where the hacker monitors the traffic that comes across the network and introduces himself as a stealth intermediary between the sender and the receiver, as shown in [Figure 1-8](#).

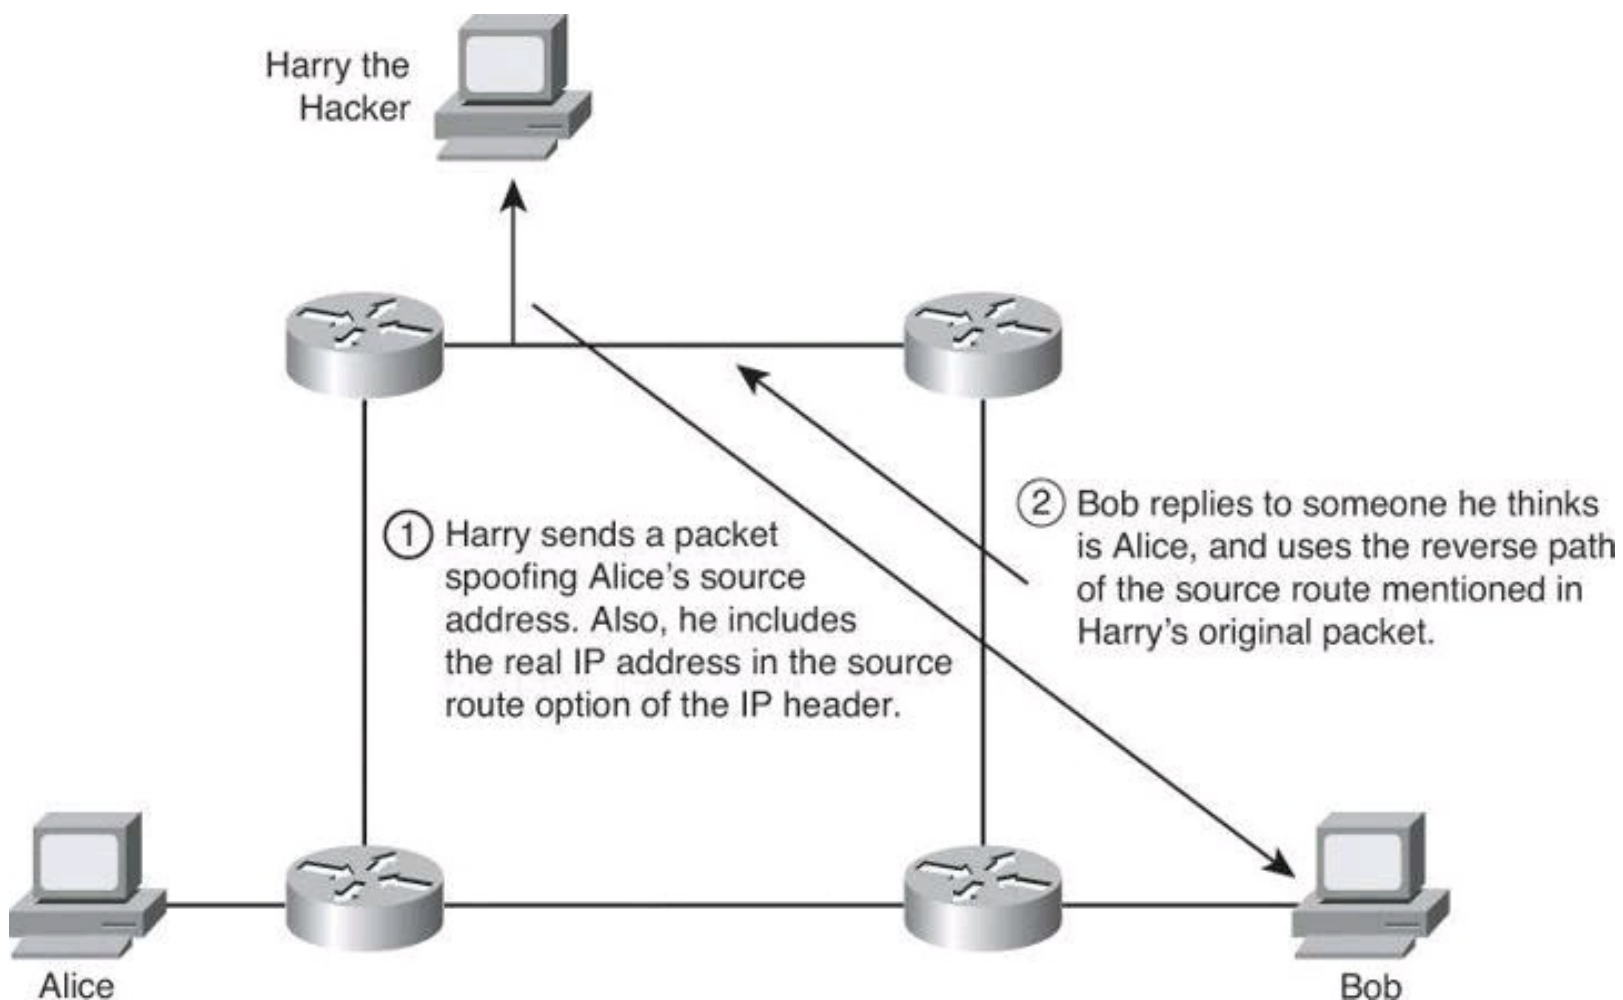


Figure 1-8. IP Source Routing Attack

Hackers use man-in-the-middle attacks to perform many security violations:

- Theft of information
- Hijacking of an ongoing session to gain access to your internal network resources
- Analysis of traffic to derive information about your network and its users
- DoS
- Corruption of transmitted data
- Introduction of new information into network sessions

Attacks are blind or nonblind. A blind attack interferes with a connection that takes place from outside, where sequence and acknowledgment numbers are unreachable. A nonblind attack interferes with connections that cross wiring used by the hacker. A good example of a blind attack can be found at http://wiki.cas.mcmaster.ca/index.php/The_Mitnick_attack.

TCP session hijacking is a common variant of the man-in-the-middle attack. The attacker sniffs to identify the client and server IP addresses and relative port numbers. The attacker modifies his or her packet headers to spoof TCP/IP packets from the client, and then waits to receive an ACK packet from the client communicating with the server. The ACK packet contains the sequence number of the next packet that the client is expecting. The attacker replies to the client using a modified packet with the source address of the server and the destination address of the client. This packet results in a reset that disconnects the legitimate client. The attacker takes over communications with the server by spoofing the expected sequence number from the ACK that was previously sent from the legitimate client to the server. (This could also be an attack against confidentiality.)

Another clever man-in-the-middle attack is for the hacker to successfully introduce himself as the DHCP server on the network, providing its own IP address as the default gateway during the DHCP offer.

Note

At this point, having read about many different attacks, you might be concerned that the security of your network is insufficient. Do not despair: many of the attacks described here are mitigated by techniques explained in this book or in other Cisco Press security books, such as *CCNP Security SECURE 642-637 Official Cert Guide*.

Overt and Covert Channels

Overt and covert channels refer to the capability to hide information within or using other information:

- **Overt channel:** A transmission channel that is based on tunneling one protocol inside of another. It could be a clear-text transmission inserted inside another clear-text protocol header.
- **Covert channel:** A transmission channel that is based on encoding data using another set of events. The data is concealed.

There are numerous ways that Internet protocols and the data that is transferred over them can provide overt and covert channels. The bad news is that firewalls generally cannot detect these channels; therefore, attackers can use them to receive confidential information in an unauthorized

manner.

With an overt channel, one protocol is tunneled within another to bypass the security policy; for example, Telnet over FTP, instant messaging over HTTP, and IP over Post Office Protocol version 3 (POP3). Another example of an overt channel is using water-marks in JPEG images to leak confidential information.

One common use of overt channel is for instant messaging (IM). Most organization firewalls allow outbound HTTP but block IM. A user on the inside of the network can leak confidential information using IM over an HTTP session.

In [Figure 1-9](#), the firewall allows outbound HTTP while a user on the inside of the network is leaking confidential information using instant messaging over HTTP.

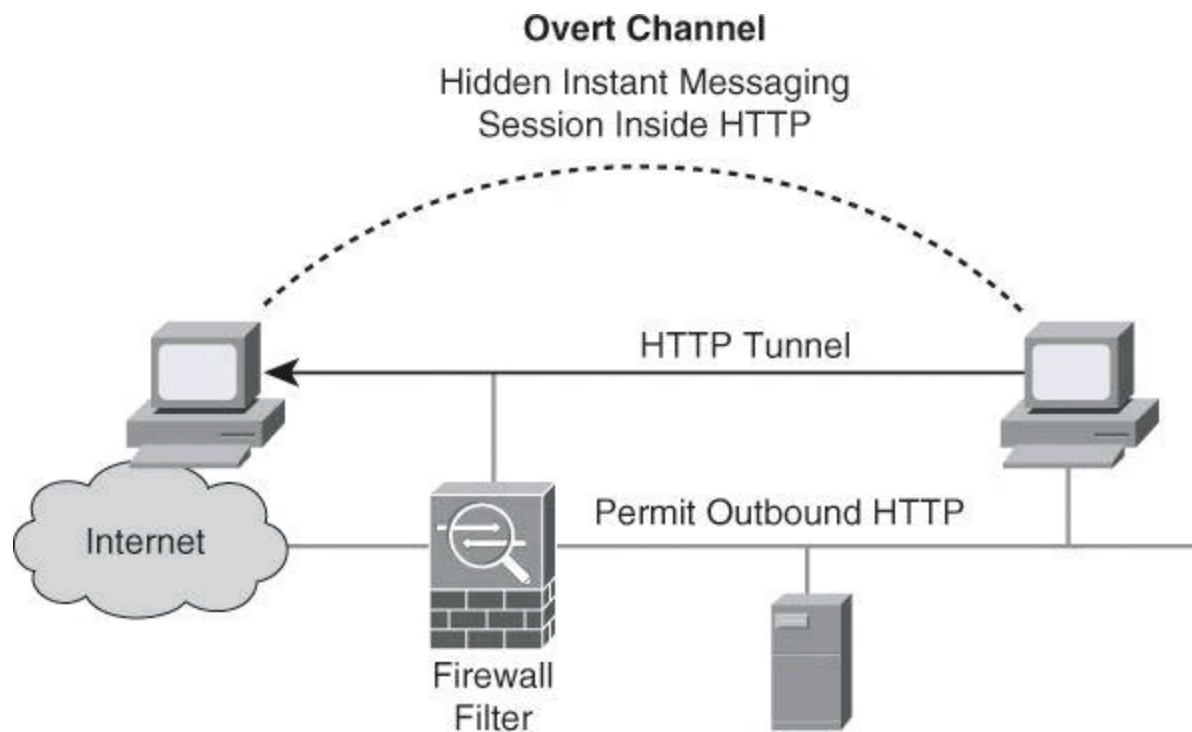


Figure 1-9. Overt Channel

Note

You can use the advanced protocol inspection in the Cisco IPS products and Cisco ASA 5500 series appliances to counter attacks such as a hidden IM session being sent inside HTTP.

Steganography is another example of an overt channel. Steganography (from the Greek word *steganos*, meaning “covered” or “secret”) literally means covered or secret writing. The combination of CPU power and interest in privacy has led to the development of techniques for hiding messages in digital pictures and digitized audio.

For example, certain bits of a digital graphic can be used to hide messages. The key to knowing which bits are special is shared between two parties that want to communicate privately. The private message typically has so few bits relative to the total number of bits in the image that changing them is not visually noticeable. Without a direct comparison of the original and the processed image, it is practically impossible to tell that anything has been changed. Still, it might be detected by statistical analysis that detects non-randomness. This non-randomness in a file indicates that information is being passed inside of the file.

Note

Steganography is very difficult to detect or prevent.

With a covert channel, information is encoded as another set of events. For example, an attacker could install a Trojan horse on a target host. The Trojan horse could be written to send binary information back to the server of the attacker. The client, infected with the Trojan horse, could return to the hacker’s server a ping status report in a binary format, where a 0 would represent a successful ping over a one-minute period, and a 1 would represent two successful pings over a one-minute period. The hacker could keep connectivity statistics for all the compromised clients he has around the world.

If ICMP is not permitted through a firewall, another tactic is to have the client visit the web page of the attacker. The Trojan horse software, now installed on the client, has a “call home” feature that automatically opens a connection to TCP port 80 at a specific IP address, the address of the hacker’s web server. All of this work is done so that the hacker can keep precise statistics of how many compromised workstations he possesses around the world. One visit per day would be represented by a 1, and no visits would be represented by a 0. As you might imagine, this technique is usually quite limited in bandwidth.

Note

Covert channels are very difficult to detect or prevent.

Phishing, Pharming, and Identity Theft

Identity theft continues to be a problem. In computing, phishing is an attempt to criminally acquire sensitive information, such as usernames, passwords, and credit card details, by masquerading as a trustworthy entity. Phishing is typically carried out by email or instant message (IM), although sometimes phone contact is attempted; the phisher often directs users to enter details at a website, as shown on the left in [Figure 1-10](#). Phishing is an example of social engineering.

Phishing

Pharming

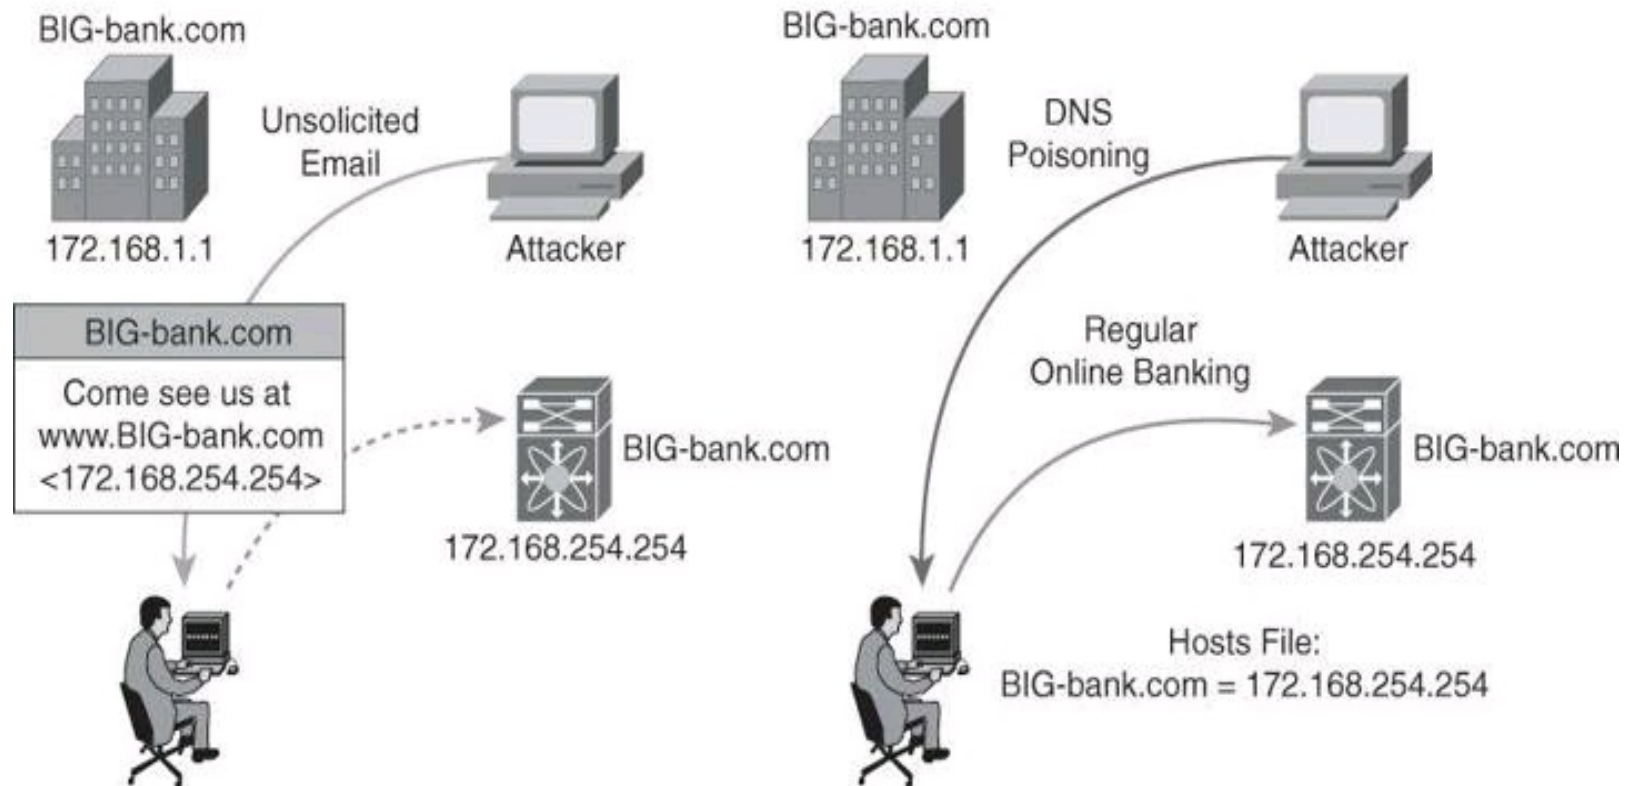


Figure 1-10. Phishing and Pharming Attacks

Note

A variation on phishing is spear phishing. In this case, a hacker sends an email that appears genuine to all the employees of an organization and hopes that a few get hooked. As an example, the email could say: “This is Christina, your HR director. The Automatic Payment organization which processes your pay is unable to do so this week. Please email me directly your banking information, and I will ensure that your pay is directly deposited in your bank account for Thursday morning.”

Pharming, also illustrated in [Figure 1-10](#), is an attack aimed at redirecting the traffic of a website to another website. Pharming is conducted either by changing the hosts file on a victim computer or by exploiting a vulnerable Domain Name System (DNS) server. Pharming has become a major concern to businesses hosting e-commerce and online banking websites.

Note

Antivirus software and spyware-removal software cannot protect against pharming. Additional methods are needed such as server-side software, DNS protection, and web browser protection.

To protect against pharming, organizations implement “personalization” technologies, such as user-chosen images on the login page. Consider also supporting identified email initiatives such as DomainKeys Identified Mail (DKIM); these initiatives are beyond the scope of this book.

Availability Attacks

DoS attacks attempt to compromise the availability of a network, host, or application. They are considered a major risk because they can easily interrupt a business process and cause significant loss. These attacks are relatively simple to conduct, even by an unskilled attacker.

DoS attacks are usually the consequence of one of the following:

- The failure of a host or application to handle an unexpected condition, such as maliciously formatted input data or an unexpected interaction of system components.
- The inability of a network, host, or application to handle an enormous quantity of data, which crashes the system or brings it to a halt. Even if the firewall protects the corporate web server sitting on the DMZ from receiving a large amount of data and thus from crashing, the link connecting the corporation with its service provider will be totally clogged, and this bandwidth starvation will itself be a DoS.

Hackers can use many types of attacks to compromise availability:

- Botnets
- DoS
- DDoS
- SYN floods
- ICMP floods
- Electrical power
- Computer environment

Note

Many availability attacks can be used against confidentiality and integrity.

Botnets

Botnet is a term for a collection of software robots, or bots, that run autonomously and automatically. They run on groups of “zombie” computers controlled by crackers.

Although the term *botnet* can be used to refer to any group of bots, it is generally used to refer to a collection of compromised systems running worms, Trojan horses, or back doors, under a common command and control infrastructure. The originator of a botnet controls the group of computers remotely, usually through a means such as Internet Relay Chat (IRC).

Often, the command and control takes place via an IRC server or a specific channel on a public IRC network. A bot typically runs hidden. Generally, the attacker has compromised a large number of systems using various methods, such as exploits, buffer overflows, and so on. Newer bots automatically scan their environment and propagate using detected vulnerabilities and weak passwords. Sometimes a controller will hide an IRC server installation on an educational or corporate site, where high-speed connections can support a large number of other bots.

Several botnets have been found and removed from the Internet. The Dutch police found a 1.5-million node botnet (<http://www.wisegeek.com/what-is-a-botnet.htm>), and the Norwegian ISP Telenor disbanded a 10,000-node botnet. Large, coordinated international efforts to shut down

botnets have also been initiated. Some estimates indicate that up to 25 percent of all personal computers are part of a botnet (<http://everythingexplained.at/Botnet/>).

DoS and DDoS Attacks

DoS attacks are the most publicized form of attack. They are also among the most difficult to eliminate. A DoS attack on a server sends an extremely large volume of requests over a network or the Internet. These large volumes of requests cause the attacked server to slow down dramatically. Consequently, the attacked server becomes unavailable for legitimate access and use.

DoS attacks differ from most other attacks because DoS attacks do not try to gain access to your network or the information on your network. These attacks focus on making a service unavailable for normal use. Attackers typically accomplish this by exhausting some resource limitation on the network or within an operating system or application. These attacks typically require little effort to execute because they either take advantage of protocol weaknesses or use traffic normally allowed into a network. DoS attacks are among the most difficult to completely eliminate because of the way they use protocol weaknesses and accepted traffic to attack a network. Some hackers regard DoS attacks as trivial and in bad form because they require so little effort to execute. Still, because of their ease of implementation and potentially significant damage, DoS attacks deserve special attention from security administrators.

System administrators can install software fixes to limit the damage caused by all known DoS attacks. However, as with viruses, hackers constantly develop new DoS attacks.

A DDoS attack generates much higher levels of flooding traffic by using the combined bandwidth of multiple machines to target a single machine or network. The DDoS attack enlists a network of compromised machines that contain a remotely controlled agent, or zombie, attack program. A master control mechanism provides direction and control. When the zombies receive instructions from the master agent, they each begin generating malicious traffic aimed at the victim.

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

[Figure 1-11](#) shows the process of a DDoS attack:

1. The hacker uses a host to scan for systems to hack.
2. After the hacker accesses handler systems, the hacker installs zombie software on them to scan, compromise, and infect agent systems.
3. Remote control attack software is loaded on agent systems.
4. When the hacker issues instructions to handlers on how to carry out the DDoS attack.

1. Scan for systems to hack.

Client System

4. The client issues commands to handlers that control agents in a mass attack.

2. Install software to scan, compromise, and infect agents.

Handler Systems

3. Agents are loaded with remote control attack software.

Agent Systems

Figure 1-11. DDoS Attack

Note

Stacheldracht, which means “barbed-wire” in German, is a well-known tool used to conduct DDoS.

Blended Threats

The actual breach and vulnerability exploit is often accomplished using a combination of malware that infects, propagates, and delivers its payload following different techniques associated with traditional malware. Known as blended threats, these attack mechanisms combine the characteristics of viruses, worms, Trojan horses, spyware, and other malware.

A blended threat will exploit a vulnerability such as a buffer overflow or lack of HTTP input validation. Such attacks can spread without human intervention by scanning for other hosts to infect, embedding code in HTML, or by spamming, to name a few methods.

Blended threats plant Trojans and back doors. They are often part of botnet attacks, which try to raise privilege levels, create network shares, and steal data.

Most blended attacks are considered “zero day,” meaning that they have not been previously identified. Blended attacks are ever-evolving and pretested by cybercriminals on common antivirus products before they are released. These threats easily breach firewalls and open channels, and they represent a challenge to detect and mitigate.

Offline Versus Online Password Cracking

Password cracking techniques can be classified as offline or online. Offline password cracking involves having the hashed result of the original password. At its own pace, the hacker could try hashing different combinations of characters until one of the hash results matches the hash of the original password. Online password cracking involves, as an example, different combinations of password on a live system. It is more difficult to achieve success with this method because most login pages lock after a certain number of unsuccessful login attempts.

Principles of Secure Network Design

In planning an overall strategy for security architecture design, sound principles are needed to accomplish an effective security posture. The selective combination of these principles provides the fundamentals for threat mitigation within the context of a security policy and risk management.

- **Defense in depth:** This is an umbrella term that encompasses many of the other guidelines in this list. It is defined by architectures based on end-to-end security, using a layered approach. The objective is to create security domains and separate them by different types of security controls. The concept also defines redundancy of controls, where the failure of one layer is mitigated by the existence of other layers of controls.
- **Compartmentalization:** Creating security domains is crucial. Different assets with different values should reside in different security domains, be it physically or logically. Granular trust relationships between compartments would mitigate attacks that try to gain a foothold in lower-security domains to exploit high-value assets in higher-security domains.
- **Least privilege:** This principle applies a need-to-know approach to trust relationships between security domains. The idea, which originated in military and intelligence operations, is that if fewer people know about certain information, the risk of unauthorized access is diminished. In network security, this results in restrictive policies, where access to and from a security domain is allowed only for the required users, application, or network traffic. Everything else is denied by default.
- **Weakest link:** This is a fundamental concept—a security system is as effective as its weakest link. A layered approach to security, with weaker or less protected assets residing in separated security domains, mitigates the necessary existence of these weakest links. Humans are often considered to be the weakest link in information security architectures.
- **Separation and rotation of duties:** This is the concept of developing systems where more than one individual is required to complete a certain task. The principle is that this requirement can mitigate fraud and error. This applies to information security controls, and it applies to both technical controls and human procedures to manage those controls.
- **Hierarchically trusted components and protection:** This principle applies a hierarchical approach to the compartmentalization and least privilege ideas, aiming at providing a more structured approach to data classification and security controls. The concept assumes that the hierarchy will be easier to implement and manage, resulting in similarly manageable and compartmentalized security controls.

- **Mediated access:** This principle is based on centralizing security controls to protect groups of assets or security domains. In that sense, firewalls, proxies, and other security controls act on behalf of the assets they are designed to protect, and mediate the trust relationships between security domains. Special considerations should be in place to prevent the mediation component from becoming a single point of failure.

- **Accountability and traceability:** This concept implies the existence of risk and the ability to manage and mitigate it, and not necessarily avoid or remove it. Information security architectures should provide mechanisms to track activity of users, attackers, and even security administrators. They should include provisions for accountability and nonrepudiation. This principle translates into specific functions, such as security audits, event management and monitoring, forensics, and others.

Cisco has always been a proponent of defense in depth. This was made clear in 2000 when it released its Cisco SAFE Blueprint for enterprise (SAFE is not an acronym), where it laid out its vision for defense in depth.

Defense in Depth

Addressing the fact that a security system is only as strong as its weakest link is often difficult when designing a system's security. The complexity of modern systems makes it hard to identify each individual weak link, let alone the weakest one. Thus, it is often most desirable to eliminate possible weaknesses by instituting several concurrent security methods.

Securing information and systems against all threats requires multiple, overlapping protection approaches that address the human, technological, and operational aspects of information technology. Using multiple, overlapping protection approaches ensures that the system is never unprotected from the failure or circumvention of any individual protection approach.

When a system is designed and implemented, its quality should always be questioned through design reviews and testing. Identification of various failure modes might help a designer evaluate the probability of element failure, and identify the links that are the most critical for the security of the whole system. Many systems have a security-based single point of failure, an element of functionality or protection that, if compromised, would cause the compromise of the whole system. It is desirable to eliminate or at least harden such single points of failure in a high-assurance system.

Defense in depth is a philosophy that provides layered security to a system by using multiple security mechanisms:

- Security mechanisms should back each other up and provide diversity and redundancy of protection.
- Security mechanisms should not depend on each other, so that their security does not depend on other factors outside their control.
- Using defense in depth, you can eliminate single points of failure and augment weak links in the system to provide stronger protection with multiple layers.

The defense-in-depth strategy recommends several principles:

- **Defend in multiple places:** Given that insiders or outsiders can attack a target from multiple points, an organization must deploy protection mechanisms at multiple locations to

resist all classes of attacks. At a minimum, you should include three defensive focus areas:

- **Defend the networks and infrastructure:** Protect the local- and wide-area communications networks from attacks, such as DoS attacks. Provide confidentiality and integrity protection for data that is transmitted over the networks; for example, use encryption and traffic flow security measures to resist passive monitoring.
 - **Defend the enclave boundaries:** Deploy firewalls and intrusion detection systems (IDS) or intrusion prevention systems (IPS) or both to resist active network attacks.
 - **Defend the computing environment:** Provide access controls and host intrusion prevention systems (HIPS) on hosts and servers to resist insider, close-in, and distribution attacks.
- **Build layered defenses:** Even the best available information assurance products have inherent weaknesses. Therefore, it is only a matter of time before an adversary finds an exploitable vulnerability. An effective countermeasure is to deploy multiple defense mechanisms between the adversary and the target. Each of these mechanisms must present unique obstacles to the adversary. Further, each mechanism should include both protection and detection measures. These measures increase the risk of detection for adversaries while reducing their chances of success, or make successful penetrations unaffordable. One example of a layered defense is to have nested firewalls (each coupled with IDS or IPS) that are deployed at outer and inner network boundaries. The inner firewalls may support more granular access control and data filtering.
- **Use robust components:** Specify the security robustness (that is, strength and assurance) of each information assurance component as a function of the value of what it is protecting and the threat at the point of application. For example, it is often more effective and operationally suitable to deploy stronger mechanisms at the network boundaries than at the user desktop.
- **Employ robust key management:** Deploy robust encryption key management and public key infrastructures that support all the incorporated information assurance technologies and that are highly resistant to attack.
- **Deploy an IDS or IPS:** Deploy infrastructures to detect and prevent intrusions and to analyze and correlate the results and react accordingly. These infrastructures should help the operations staff answer the following questions:
- Am I under attack?
 - Who is the source?
 - What is the target?
 - Who else is under attack?
 - What are my options?

Evaluating and Managing the Risk

The security policy developed in your organization drives all the steps taken to secure network resources. The development of a comprehensive security policy prepares you for the rest of your

security implementation. To create an effective security policy, it is necessary to do a risk analysis, which will be used to maximize the effectiveness of the policy and procedures that will be put in place. Also, it is essential that everyone be aware of the policy; otherwise, it is doomed to fail.

All design guidelines and principles, and the resulting security architecture, should be aimed at managing risk. Risk is, or should be, the building block of information security.

Levels of Risks

By its very nature, risk management is a tradeoff between the effort (cost) to protect organizational assets and the resulting level of exposure of those assets. This simple rule is a good starting point: the cost to protect an asset will likely not be greater than the value of the asset itself. There are obviously exceptions to the rule; for instance, cases that involve national security, or instances where the value of the asset is incalculable, such as cases where human life is involved.

The tradeoffs in risk management are based on its building blocks: assets and vulnerabilities, threats and countermeasures. Different values and scenarios for these components move the risk indicators up and down. Understanding these values and scenarios is critical in defining a risk management strategy.

For instance, would you use old, worn tires at high speed on a highway? The answer is obviously no. The asset that you are trying to protect (your life) is too valuable, and the countermeasure to mitigate the risk of navigating the highway, driving at a slow speed, is not good enough. It is inexpensive but not effective.

However, using a worn-down tire as a swing does not result in life-threatening risk in the majority of situations. The asset (your life) remains the same, but the threats that are able to exploit the vulnerabilities of the tire are mitigated or nonexistent. The premise changes again if you think that this worn-down tire will be used to swing your child. You may or may not risk using the old tire, but the value of the asset may prevent you from facing risk even if it is minimal.

The previous example is a simplistic view of information security risk. Imagine an organizational risk management effort, considering thousands of assets with different (and often subjective) valuation criteria, different (and often unknown) levels of vulnerability, and potentially exposed to an avalanche of threats that change by the minute. Risk management becomes a delicate balance and involves constant tuning of countermeasures in the face of sophisticated threat vectors, exploiting assets that are often located outside of corporate control.

Information security risk management is a comprehensive process that requires organizations to frame risk (in other words, establish the context for risk-based decisions), assess risk, respond to risk, and monitor risk on an ongoing basis. The result is a dynamic process in nature, evolving along with internal factors (assets, vulnerabilities, security policies, and architectures) and external factors (threats, and business, legal, and compliance forces).

Other sections in this chapter will expand on these concepts and present commonly used risk management strategies, within the context of a security policy and a security lifecycle process.

Risk Analysis and Management

Every process of security should first address the following questions:

- Which are the threats the system is facing?
- Which are the probable threats and what would be their consequence, if exploited?

The threat-identification process provides an organization with a list of threats to which a system is subject in a particular environment.

Note

An interesting method of modeling security threats is the attack trees method developed by Bruce Schneier. You can find more information about this method at http://en.wikipedia.org/wiki/Attack_tree.

Risk Analysis

Risk analysis is the systematic study of uncertainties and risks. Risk analysts seek to identify the risks that a company faces, understand how and when they arise, and estimate the impact (financial or otherwise) of adverse outcomes. Risk managers start with risk analysis, and then seek to take actions that will mitigate these risks. Risk analysis tries to estimate the probability and severity of threats faced by an organization's system that needs protection, and then provides to the organization a prioritized list of risks that the organization must mitigate. This allows the organization to focus on the most important threats first.

Two types of risk analysis are of interest in information security:

- **Quantitative:** Quantitative risk analysis uses a mathematical model that assigns monetary values to assets, the cost of threats being realized, and so on. Quantitative risk analysis provides an actual monetary figure of expected losses, which is typically based on an annual cost. You can then use this number to justify proposed countermeasures. For example, if you can establish that you will lose \$1,000,000 by doing nothing, you can justify spending \$300,000 to reduce that risk by 50 percent to 75 percent.
- **Qualitative:** Qualitative risk analysis uses a scenario model. This approach is best for large cities, states, and countries to use because it is impractical for such entities to try to list all their assets, which is the starting point for any quantitative risk analysis. By the time a typical national government could list all of its assets, the list would have hundreds or thousands of changes and would no longer be accurate.

Qualitative risk analysis is straightforward provided you have the resources to document all the assets. However, quantitative risk analysis is more tricky, so we will take a closer look at it.

Quantitative Risk Analysis Formula

Quantitative risk analysis relies on specific formulas to determine the value of the risk decision variables. These include formulas that calculate the asset value (AV), exposure factor (EF), single loss expectancy (SLE), annualized rate of occurrence (ARO), and annualized loss expectancy (ALE). The ALE formula is as follows: $ALE = (AV * EF) * ARO$.

The AV is the value of an asset. This would include the purchase price, the cost of deployment, and the cost of maintenance. In the case of a database or a web server, the AV should also include the cost of development. AV is not an easy number to calculate.

The EF is an estimate of the degree of destruction that will occur. For example, suppose that you consider flood a threat. Could it destroy your data center? Would the destruction be 60 percent, 80 percent, or 100 percent? The risk-assessment team would have to make a determination that evaluates everything possible, and then make a judgment call. For this example, assume that a flood will have a 60 percent destruction factor, because you store a backup copy of all media and data offsite. Your only losses would be the hardware and productivity.

As another example of EF, consider data entry errors, which are much less damaging than a flood. A single data entry error would hardly be more than a fraction of a percent in exposure. The exposure factor of a data entry error might be as small as .001 percent.

Caution

One of the ironies of risk analysis is how much estimating (guessing) is involved.

The SLE calculation is a number that represents the expected loss from a single occurrence of the threat. The SLE is defined as $AV * EF$.

To use our previous examples, you would come up with the following results for the SLE calculations:

- Flood threat
 - Exposure factor: 60 percent
 - AV of the enterprise: US\$10,000,000
 - $\$10,000,000 * .60 = \$6,000,000$
- Data entry error
 - Exposure factor: .001 percent
 - AV of data and databases: \$1,000,000
 - $\$1,000,000 * .000001 = \10 SLE

The ARO is a value that estimates the frequency of an event and is used to calculate the ALE.

Continuing the preceding example, the type of flood that you expect could reach your data center would be a “flood of the century” type of event. Therefore, you give it a 1/100 chance of occurring this year, making the ARO for the flood 1/100.

Furthermore, you expect the data entry error to occur 500 times a day. Because the organization is open for business 250 days per year, you estimate the ARO for the data entry error to be $500 * 250$, or 125,000 times.

Risk analysts calculate the ALE in annualized terms to address the cost to the organization if the organization does nothing to counter existing threats. The ALE is derived from multiplying the SLE by the ARO. The following ALE calculations continue with the two previous examples:

- Flood threat
 - SLE: \$6,000,000
 - ARO: .01
 - $\$6,000,000 * .01 = \$60,000$ ALE

- Data input error
 - SLE: \$10
 - ARO: 125,000
 - $\$10 * 125,000 = \$1,250,000$ ALE

A decision to spend \$50,000 to enhance the security of our database applications to reduce data entry errors by 90 percent is now an easy decision. It is equally easy to reject a proposal to enhance our defenses against floods that costs \$3,000,000.

When you perform a quantitative risk analysis, you identify clear costs as long as the existing conditions remain the same. You compile a list of expected issues, the relative cost of those events, and the total cost if all expected threats are realized. These numbers are put into annual terms to coincide with the annual budgets of most organizations.

You then use these numbers in decision making. If an organization has a list of 10 expected threats, it can then prioritize the threats and address the most serious threats first. This prioritization enables management to focus their resources where it will do the most good.

For example, suppose an organization has the following list of threats and costs as the product of performing a quantitative risk analysis:

- **Insider network abuse:** \$1,000,000 in lost productivity
- **Data input error:** \$500,000
- **Worm outbreak:** \$100,000
- **Viruses:** \$10,000
- **Laptop theft:** \$10,000

Decision makers could easily decide that it is of greatest benefit to address insider network abuse and leave the antivirus solution alone. They could also find it easy to support a \$200,000 URL filtering solution to address insider network abuse and reject a \$40,000 solution designed to enhance laptop safety. Without these numbers from a risk analysis, the decisions made would likely differ.

Building Blocks of Risk Analysis

Conducting a risk analysis starts with the gathering of pertinent information. The building blocks of the process follow the definition of risk used in this book: the organizational impact of threat vectors exploiting vulnerabilities of the assets you are trying to protect.

In that sense, the initial information gathering, in preparation for the risk calculations described in the previous example, should collect and define the following:

- **Assets and their value:** This information, shown in [Table 1-1](#), is typically obtained from data classification, inventories of assets, and other sources. A general principle is to use discrete numerical values for the exposure factor (EF) based on discrete values that reflect the impact of losing the asset. These values are generally based on data classification techniques (confidential, secret, top secret, and so on), and the impact is based on organizationally relevant criteria (replacement cost, liability, and so on).

Table 1-1. List of Assets and Their Value

	Confidentiality	Integrity	Availability
Low Value	Limited effect	Limited effect	Limited effect
Moderate Value	Serious effect	Serious effect	Serious effect
High Value	Severe effect	Severed effect	Severe effect

• **Vulnerabilities:** This information is typically gathered from vulnerability assessments, which will be discussed further later in this chapter. Several tools are available, like Nessus and other commercial vulnerability assessment products. The use of public- or platform-specific vulnerability classification databases is commonplace. They include the Common Vulnerabilities and Exposures (CVE) effort by MITRE, <http://cve.mitre.org>, and the National Vulnerability Database (NVD) sponsored by the National Institute of Standards and Technology (NIST), <http://nvd.nist.gov>. An example of vulnerability categorization is shown in [Table 1-2](#).

Table 1-2. Example of Vulnerability Categorization Headings

Categorization	Procedures	Processes	Systems	Network

• **Threats, their impact, and rate or probability of occurrence:** This information is commonly obtained from publicly available databases, such as the MITRE Common Attack Pattern Enumeration and Classification (CAPEC), <http://capec.mitre.org>. Calculating the rate of occurrence is a probabilistic exercise and is often considered subjective and specific for individual organizations or industries. [Table 1-3](#) shows an example of this information gathering.

Table 1-3. Example of Threats, Impact, and Probability of Occurrence

Impact Category	Critical	Serious	Moderate	Minor	Negligible
Definition	Inability to achieve minimum requirements	Major cost and schedule increases	Moderate cost and schedule increases	Small cost and schedule increases	No effect

Risk Scores

With asset, vulnerability, and threat components defined, risk scores are obtained by applying formulas of quantitative risk analysis. [Figure 1-12](#) illustrates the process.

Categorize threats, their severity and probability of occurrence.

Categorize assets, their value and vulnerabilities.



A risk matrix produces risk scores based on asset, vulnerability, threat, and countermeasure landscape



Figure 1-12. Obtaining a Risk Score

A risk matrix is then calculated, including risk scores for assets and groups of assets and, ideally, an organization risk score that can be used in security monitoring, incident response, and policy reviews. These risk scores provide an idea of the landscape of assets, threats, vulnerabilities, and countermeasures, the components of risk, at a given point in time.


A Lifecycle Approach to Risk Management

Managing risk is a complex, multifaceted activity that requires the involvement of the entire organization, including the following:

- Senior leaders and executives who provide the strategic vision and top-level goals and objectives for the organization
- Midlevel leaders who plan, execute, and manage projects
- Individuals who operate the information systems supporting the organization's mission and business functions

[Figure 1-13](#) shows that risk management is a comprehensive process that requires organizations to do the following:

- Frame risk (that is, establish the context for risk-based decisions)
- Assess risk
- Respond to risk once determined
- Monitor risk on an ongoing basis using effective organizational communications and a feedback loop for continuous improvement in the risk-related activities of organizations



Source: NIST 800-39, 2011

Figure 1-13. Lifecycle Approach to Risk Management According to NIST 800-39

Risk management is carried out as a holistic, organization-wide activity that addresses risk from the strategic level to the tactical level. Approaching risk management in this way ensures that risk-based decision making is integrated into every aspect of the organization.

Regulatory Compliance

Compliance regulations have been a major driver for security in organizations of all kinds, and the following trends have emerged over the past decade:

- Strengthened enforcement
- Global spread of data breach notification laws
- More prescriptive regulations
- Growing requirements regarding third parties (business partners)
- Risk-based compliance on the rise
- Compliance process streamlined and automated

The compliance regulation defines not only the scope and parameters for the risk and security architectures of an organization, but also the liability for those who do not comply. Recently there

have been major shifts in the compliance landscape:

- Although enforcement of existing regulations has been weak in many jurisdictions worldwide, regulators and standards bodies are now tightening enforcement through expanded powers, higher penalties, and harsh enforcement actions.
- In the future, it will be more difficult to hide information security failings wherever organizations do business. Legislators are forcing transparency through the introduction of breach notification laws in Europe, Asia, and North America as data breach disclosure becomes a global principle.
- As more regulations are introduced, there is a trend toward increasingly prescriptive rules. For example, laws in the states of Massachusetts and Nevada, which went into effect in 2010, apply not only to companies based in these states but also to all external organizations that manage the personal information of these states' residents.
- Regulators are also making it clear that enterprises are responsible for ensuring the protection of their data when it is being processed by a business partner, including cloud service providers.
- For many organizations, stricter compliance could help focus management attention on security; but if they take a “check-list approach” to compliance, it will detract from actually managing risk and may not improve security.
- The new compliance landscape will increase costs and risks. For example, it takes time and resources to substantiate compliance. Increased requirements for service providers give rise to more third-party risks.
- With more transparency, there are now greater consequences for data breaches. For example, expect to see more litigation as customers and business partners seek compensation for compromised data. But the harshest judgments will likely come from the court of public opinion—with the potential to permanently damage the reputation of an enterprise.

[Table 1-4](#) illustrates some examples of relevant compliance regulations (most of which were introduced earlier in the chapter) that affect organizations all over the world. Geographic boundaries are blurring as globalization makes organizations subject to regulations in several countries. Industry scope boundaries are also blurring. For instance, many service organizations providing services to the U.S. government have to comply with U.S. federal regulations related to information security.

Table 1-4. Examples of Compliance Regulations

Regulation	Geographic Scope	Applies To
EU Data Protection Directive (EU 95/46/EC)	European Union	All organizations operating in the 27 EU member countries
Sarbanes-Oxley	United States	All publicly traded companies in the U.S. (exemption for smaller reporting companies)
PIPEDA	Canada	All organizations in Canada
PCI DSS	Global	All organizations processing credit card data
HIPAA	United States	All healthcare organizations in the U.S.
FISMA	United States	Federal agencies and service organizations
Basel II	Global	All internationally active banks with assets of \$250 billion or more
DMCA	United States	Individuals and organizations in the U.S.
NERC	North America	North America users, owners, and operators of the bulk electric power system
GLBA	United States	All financial institutions in the U.S.
Safe Harbor Act	European Union	U.S. companies doing business in the EU

The following are descriptions of some of the regulations listed in [Table 1-4](#):

- The Gramm-Leach-Bliley Act (GLBA) of 1999 erased long-standing antitrust laws that prohibited banks, insurance companies, and securities firms from merging and sharing information with one another. The idea was that smaller firms would then be able to pursue acquisitions or alliances, or both, that would help encourage competition against many of the larger financial institutions. Included in the GLBA were several consumer privacy protections. Namely, companies must tell their customers what kinds of data they plan to share and with whom, and they must give their customers a chance to opt out of that data sharing.
- On the healthcare side, the Health Insurance Portability and Accountability Act (HIPAA) of 2000 requires the U.S. Department of Health and Human Services to develop a set of national standards for healthcare transactions. These standards provide assurance that the electronic transfer of confidential patient information will be as safe as, or safer than, paper-based patient records.
- The Sarbanes-Oxley (SOX) Act of 2002 is a U.S. law that was created in response to a number of major corporate and accounting scandals, including those affecting Enron, Tyco International, Peregrine Systems, and WorldCom. These scandals resulted in a decline of public trust in accounting and reporting practices.
- The Federal Information Security Management Act (FISMA) of 2002 was intended to

bolster computer and network security within the U.S. government and affiliated parties by requiring yearly audits. FISMA also brought attention within the U.S. government to cyber security, which the U.S. government had largely neglected previously.

Globalization, as with any other context, is changing the face of regulatory compliance. Regulators are not just looking at ways to strengthen existing laws. Regulators are also introducing new laws that are aimed at forcing more transparency, in a way that affects organizations on a global basis.

Data breach disclosure is becoming a global principle as jurisdictions worldwide adopt privacy and data protection laws that include a general obligation to notify government agencies, individuals, and other authorities such as law enforcement of unauthorized access or use of personal data. Requirements vary, including who must be notified, the type of data that triggers notification, and if there is a risk-of-harm threshold.

California's landmark legislation SB-1386 set off a wave of state breach notification laws that now cover almost the entire United States. Recently, this trend has spread to the European Union. The Privacy and Electronic Communications Directive (e-Privacy Directive) was amended in late 2009 to include data breach notification. It is now mandatory for telephone companies and ISPs in the EU to inform national regulatory authorities of any data security breach. Depending on the effects of the breach, they may also be required to inform subscribers. The upcoming overhaul of the EU Data Protection Directive is expected to include data breach notification requirements, which would broaden breach disclosure to cover all industries in all 27 member countries in the EU.

[Table 1-5](#) shows how regulations are becoming the norm around the world.

Table 1-5. Acceleration of Compliance Regulation Around the World

Year	Country	Data Breach Notification Law
2003	U.S.	California's landmark SB-1386 starts wave of state laws.
2003–2010	U.S.	46 states enact notification laws.
2008	U.K.	Information Commissioner's Office issues best practice guidance requiring notification.
2009	EU	e-Privacy Directive amended to include notification requirements for electronic communications sector.
	Germany	National privacy law amended to include notification.
2010	Austria	National privacy law amended to include notification.
	France	Draft legislation passed in senate would make notification mandatory.
	Canada	National privacy law amended to include notification.
	Mexico	New privacy law enacted that includes notification.
	Ireland	Code of Practice issued regarding notification.
	Hong Kong	Privacy Commissioner issues guidance note on breach notification.
	EU	Data Protection Directive under review for revision; proposed law expected by 2011 to include notification requirements for all industries; to be implemented in all 27 EU member countries.

Security Policies

Every organization has something that someone else wants. Someone might want that something for himself, or he might want the satisfaction of denying something to its rightful owner. Your assets are what need the protection of a security policy.

Determine what your assets are by asking (and answering) the following questions:

- What do you have that others want?
- What processes, data, or information systems are critical to you, your company, or your organization?
- What would stop your company or organization from doing business or fulfilling its mission?

The answers identify assets in a wide range, including critical databases, vital applications, vital company customer and employee information, classified commercial information, shared drives, email servers, and web servers.

A security policy comprises a set of objectives for the company, rules of behavior for users and administrators, and requirements for system and management that collectively ensure the security of network and computer systems in an organization. A security policy is a “living document,” meaning that the document is never finished and is continuously updated as technology and employee

requirements change.

The security policy translates, clarifies, and communicates the management position on security as defined in high-level security principles. The security policy acts as a bridge between these management objectives and specific security requirements. It informs users, staff, and managers of their obligatory requirements for protecting technology and information assets. It should specify the mechanisms that you need to meet these requirements. It also provides a baseline from which to acquire, configure, and audit computer systems and networks for compliance with the security policy. Therefore, an attempt to use a set of security tools in the absence of at least an implied security policy is meaningless.

The three reasons for having a security policy are as follows:

- To inform users, staff, and managers
- To specify mechanisms for security
- To provide a baseline



One of the most common security policy components is an acceptable use policy (AUP). This component defines what users are allowed and not allowed to do on the various components of the system, including the type of traffic that is allowed on the networks. The AUP should be as explicit as possible to avoid ambiguity or misunderstanding. For example, an AUP might list the prohibited website categories.

Note

Some sites refer to an acceptable use policy as an *appropriate use policy*.

A properly defined security policy does the following:

- Protects people and information
- Sets the rules for expected behavior
- Authorizes staff to monitor, probe, and investigate
- Defines the consequences of violations



The audience for the security policy is anyone who might have access to your network, including employees, contractors, suppliers, and customers. However, the security policy should treat each of these groups differently.

The audience determines the content of the policy. For example, you probably do not need to include a description of *why* something is necessary in a policy that is intended for the technical staff. You can assume that the technical staff already knows why a particular requirement is included. Managers are also not likely to be interested in the technical aspects of why a particular requirement

is needed. However, they might want the high-level overview or the principles supporting the requirement. When end users know why a particular security control has been included, they are more likely to comply with the policy.

In the policy, users can be organized into two audiences:

- Internal audience
 - Managers and executives
 - Departments and business units
 - Technical staff
 - End users
- External audience
 - Partners
 - Customers
 - Suppliers
 - Consultants and contractors

One document will not likely meet the needs of the entire audience of a large organization. The goal is to ensure that the information security policy documents are coherent with its audience needs.

Security Policy Components

[Figure 1-14](#) shows the hierarchy of a corporate policy structure that is aimed at effectively meeting the needs of all audiences.




Figure 1-14. Components of a Comprehensive Security Policy

Most corporations should use a suite of policy documents to meet their wide and varied needs:

- **Governing policy:** This policy is a high-level treatment of security concepts that are important to the company. Managers and technical custodians are the intended audience. The governing policy controls all security-related interaction among business units and supporting departments in the company. In terms of detail, the governing policy answers the “what” security policy questions.
- **End-user policies:** This document covers all security topics important to end users. In terms of detail level, end-user policies answer the “what,” “who,” “when,” and “where” security policy questions at an appropriate level of detail for an end user.
- **Technical policies:** Security staff members use technical policies as they carry out their

security responsibilities for the system. These policies are more detailed than the governing policy and are system or issue specific (for example, access control or physical security issues). In terms of detail, technical policies answer the “what,” “who,” “when,” and “where” security policy questions. The “why” is left to the owner of the information.

Note

To assist you at drafting your security policies, consider the SANS security policies repository at <http://www.sans.org/resources/policies>.

For readers interested in security policies for academic institutions, visit the University of Toronto’s Computer Security Administration website for a comprehensive example of a network security policy for a higher education institution: http://www.cns.utoronto.ca/newsite/documentation/policies/policy_5.htm

Governing Policy

The governing policy outlines the security concepts that are important to the company for managers and technical custodians:

- It controls all security-related interactions among business units and supporting departments in the company.
- It aligns closely with not only existing company policies, especially human resource policies, but also any other policy that mentions security-related issues, such as issues concerning email, computer use, or related IT subjects.
- It is placed at the same level as all companywide policies.
- It supports the technical and end-user policies.
- It includes the following key components:
 - A statement of the issue that the policy addresses
 - A statement about your position as IT manager on the policy
 - How the policy applies in the environment
 - The roles and responsibilities of those affected by the policy
 - What level of compliance to the policy is necessary
 - Which actions, activities, and processes are allowed and which are not
 - What the consequences of noncompliance are

End-User Policies

End-user policies are compiled into a single policy document that covers all the topics pertaining to information security that end users should know about, comply with, and implement. This policy may overlap with the technical policies and is at the same level as a technical policy. Grouping all the end-user policies together means that users have to go to only one place and read one document to learn everything that they need to do to ensure compliance with the company security policy.

Technical Policies

Security staff members use the technical policies in the conduct of their daily security

responsibilities. These policies are more detailed than the governing policy and are system or issue specific (for example, router security issues or physical security issues). These policies are essentially security handbooks that describe what the security staff does, but not how the security staff performs its functions.

The following are typical policy categories for technical policies:

- General policies
 - **Acceptable use policy (AUP):** Defines the acceptable use of equipment and computing services, and the appropriate security measures that employees should take to protect the corporate resources and proprietary information.
 - **Account access request policy:** Formalizes the account and access request process within the organization. Users and system administrators who bypass the standard processes for account and access requests may cause legal action against the organization.
 - **Acquisition assessment policy:** Defines the responsibilities regarding corporate acquisitions and defines the minimum requirements that the information security group must complete for an acquisition assessment.
 - **Audit policy:** Use to conduct audits and risk assessments to ensure integrity of information and resources, investigate incidents, ensure conformance to security policies, or monitor user and system activity where appropriate.
 - **Information sensitivity policy:** Defines the requirements for classifying and securing information in a manner appropriate to its sensitivity level.
 - **Password policy:** Defines the standards for creating, protecting, and changing strong passwords.
 - **Risk-assessment policy:** Defines the requirements and provides the authority for the information security team to identify, assess, and remediate risks to the information infrastructure that is associated with conducting business.
 - **Global web server policy:** Defines the standards that are required by all web hosts.
- Email policies
 - **Automatically forwarded email policy:** Documents the policy restricting automatic email forwarding to an external destination without prior approval from the appropriate manager or director.
 - **Email policy:** Defines the standards to prevent tarnishing the public image of the organization.
 - **Spam policy:** The AUP covers spam.
- Remote-access policies
 - **Dial-in access policy:** Defines the appropriate dial-in access and its use by authorized personnel.
 - **Remote-access policy:** Defines the standards for connecting to the organization network from any host or network external to the organization.

- **VPN security policy:** Defines the requirements for remote-access IP Security (IPsec) or Layer 2 Tunneling Protocol (L2TP) VPN connections to the organization network.
- Personal device and phone policies
 - **Analog and ISDN line policy:** Defines the standards to use analog and ISDN lines for sending and receiving faxes and for connection to computers.
 - **Personal communication device policy:** Defines the information security's requirements for personal communication devices, such as voicemail, smartphones, tablets, and so on.
- Application policies
 - **Acceptable encryption policy:** Defines the requirements for encryption algorithms that are used within the organization.
 - **Application service provider (ASP) policy:** Defines the minimum security criteria that an ASP must execute before the organization uses the ASP's services on a project.
 - **Database credentials coding policy:** Defines the requirements for securely storing and retrieving database usernames and passwords.
 - **Interprocess communications policy:** Defines the security requirements that any two or more processes must meet when they communicate with each other using a network socket or operating system socket.
 - **Project security policy:** Defines requirements for project managers to review all projects for possible security requirements.
 - **Source code protection policy:** Establishes minimum information security requirements for managing product source code.
- Network policies
 - **Extranet policy:** Defines the requirement that third-party organizations that need access to the organization networks must sign a third-party connection agreement.
 - **Minimum requirements for network access policy:** Defines the standards and requirements for any device that requires connectivity to the internal network.
 - **Network access standards:** Defines the standards for secure physical port access for all wired and wireless network data ports.
 - **Router and switch security policy:** Defines the minimal security configuration standards for routers and switches inside a company production network or used in a production capacity.
 - **Server security policy:** Defines the minimal security configuration standards for servers inside a company production network or used in a production capacity.
- **Wireless communication policy:** Defines standards for wireless systems that are used to connect to the organization networks.

• **Document retention policy:** Defines the minimal systematic review, retention, and destruction of documents received or created during the course of business. The categories of retention policy are, among others:

- **Electronic communication retention policy:** Defines standards for the retention of email and instant messaging.
- **Financial retention policy:** Defines standards for the retention of bank statements, annual reports, pay records, accounts payable and receivable, and so on.
- **Employee records retention policy:** Defines standards for the retention of employee personal records.
- **Operation records retention policy:** Defines standards for the retention of past inventories information, training manuals, suppliers lists, and so forth.

Standards, Guidelines, and Procedures

Security policies establish a framework within which to work, but they are too general to be of much use to individuals responsible for implementing these policies. Because of this, other, more-detailed documents exist. Among the more important of these detailed documents are the standards, guidelines, and procedures documents.

Whereas policy documents are very much high-level overview documents, the standards, guidelines, and procedures documents are documents that the security staff will use regularly to implement the security policies.

Standards

Standards enable an IT staff to be consistent. They specify the use of specific technologies so that IT staff members can narrow the focus of their expertise to those technologies instead of trying to know everything about all sorts of technologies. Standards also try to provide consistency in the network, because supporting multiple versions of hardware and software is unreasonable unless it is necessary. The most successful IT organizations have standards to improve efficiency and to keep things as simple as possible.

Standardization also applies to security. One of the most important security principles is consistency. If you support 100 routers, it is important that you configure all 100 routers as similarly as possible. If you do not do this, it is difficult to maintain security. When you do not strive for the simplest of solutions, you usually fail in being secure.

Guidelines

Guidelines help provide a list of suggestions on how you can do things better. Guidelines are similar to standards, but are more flexible and are not usually mandatory. You will find some of the best guidelines available in repositories known as “best practices.” The following is a list of widely available guidelines:

- National Institute of Standards and Technology (NIST) Computer Security Resource Center; <http://csrc.nist.gov/>
- National Security Agency (NSA) Security Configuration Guides;

http://www.nsa.gov/ia/mitigation_guidance/security_configuration_guides/index.shtml

- The Common Criteria for Information Technology Security Evaluation; <http://www.commoncriteriaportal.org/>
- Defense Information Systems Agency (DISA) Field Security Operations Office – Security Technical Information Guides (STIG); <http://iase.disa.mil/stigs/>

Note

Note that the Rainbow Series from NIST was historically a reliable source for InfoSec guidelines but is now outdated.

Procedures

Procedure documents are longer and more detailed than the standards and guidelines documents. Procedure documents include the details of implementation, usually with step-by-step instructions and graphics. Procedure documents are extremely important for large organizations to enable them to have the consistency of deployment that is necessary to have a secure environment. Inconsistency is the enemy of security.

[Table 1-6](#) provides a comparative chart for standards, guidelines, and procedures, which accompany security policies.

Table 1-6. Comparison Between Standards, Guidelines, Procedures

	Characteristics
Standards	Specify the use of specific technologies in a uniform way
	Improve efficiency
	Are usually mandatory
	Accomplish consistency and uniformity
Guidelines	Are similar to standards, but more flexible and not usually mandatory
	Can be used to define how standards should be developed or to guarantee adherence to general security policies
	Include NIST Computer Security Resource Center, NSA Security Configuration Guides, Common Criteria, and others
Procedures	Are usually required
	Are the lowest level of the policy chain
	Provide detailed steps used to perform specific tasks
	Provide the steps required to implement the policies, standards, and guidelines
	Are also known as practices

Security Policy Roles and Responsibilities

In any organization, it is senior management, such as the CEO, that is always ultimately responsible for everything. Typically, senior management only oversees the development of a security policy. The creation and maintenance of a security policy is usually delegated to the people in charge of IT or security operations.

Sometimes the senior security or IT management personnel, such as the chief security officer (CSO), the chief information officer (CIO), or the chief information security officer (CISO), will have the expertise to create the policy, sometimes they will delegate it, and sometimes it will be a bit of both strategies. But the senior security person is always intimately involved in the development and maintenance of security policy. Guidelines can provide a framework for policy decision making.

Senior security staff is often consulted for input on a proposed policy project. They might even be responsible for the development and maintenance of portions of the policy. It is more likely that senior staff will be responsible for the development of standards and procedures.

Everyone else who is involved in the security policy has the duty to abide by it. Many policy statements will include language that refers to a potential loss of employment for violation of the policy. IT staff and end users alike are responsible to know the policy and follow it.

Security Awareness

Technical, administrative, and physical controls can all be defeated without the participation of the end-user community. To get accountants, administrative assistants, and other end users to think about information security, you must regularly remind them about security. The technical staff also needs regular reminders because their jobs tend to emphasize performance, such as introducing new technologies, increasing throughput, and the like, rather than secure performance, such as how many attacks they repelled. Therefore, leadership must develop a nonintrusive program that keeps everyone aware of security and how to work together to maintain the security of their data. The three key components used to implement this type of program are awareness, training, and education.

An effective computer security-awareness and training program requires proper planning, implementation, maintenance, and periodic evaluation. In general, a computer security-awareness and training program should encompass the following seven steps:

Step 1. Identify program scope, goals, and objectives.

The scope of the program should provide training to all types of people who interact with IT systems. Because users need training that relates directly to their use of particular systems, you need to supplement a large, organization-wide program with more system-specific programs.

Step 2. Identify training staff.

It is important that trainers have sufficient knowledge of computer security issues, principles, and techniques. It is also vital that they know how to communicate information and ideas effectively.

Step 3. Identify target audiences.

Not everyone needs the same degree or type of computer security information to do his or her job. A computer security-awareness and

training program that distinguishes between groups of people, presents only the information that is needed by the particular audience, and omits irrelevant information will have the best results.

Step 4. Motivate management and employees.

To successfully implement an awareness and training program, it is important to gain the support of management and employees. Consider using motivational techniques to show management and employees how their participation in a computer security and awareness program will benefit the organization.

Step 5. Administer the program.

Several important considerations for administering the program include visibility, selection of appropriate training methods, topics, and materials, and presentation techniques.

Step 6. Maintain the program.

You should make an effort to keep abreast of changes in computer technology and security requirements. A training program that meets the needs of an organization today may become ineffective when the organization starts to use a new application or changes its environment, such as by connecting to the Internet.

Step 7. Evaluate the program.

An evaluation should attempt to ascertain how much information is retained, to what extent computer security procedures are being followed, and the general attitudes toward computer security.

A successful IT security program consists of the following:

1. Developing IT security policy that reflects business needs tempered by known risks.
2. Informing users of their IT security responsibilities, as documented in agency security policy and procedures.
3. Establishing processes for monitoring and reviewing the program.

You should focus security awareness and training on the entire user population of the organization. Management should set the example for proper IT security behavior within an organization. An awareness program should begin with an effort that you can deploy and implement in various ways and be aimed at all levels of the organization, including senior and executive managers. The effectiveness of this effort usually determines the effectiveness of the awareness and training program and how successful the IT security program will be.

Secure Network Lifecycle Management

The lifecycle approach looks at the different phases of security, such as assessment, testing, implementation, monitoring and so forth, to provide methodology in securing our networks. The roles of risk, regulatory compliance, and security policies in designing and building effective security architectures have been described. How are these three components related?

IT Governance, Risk Management, and Compliance

Organizational efforts for IT governance, risk management, and compliance (sometimes known as IT GRC) are often separated by department or regulation type within organizations. This can create many problems, including unidentified risks, redundancies, and higher costs, requiring more resources, time, and effort to achieve a secure IT environment that meets regulatory compliance requirements. Moreover, while business processes and business process improvements are common practices in most organizations, this approach is often missing in the area of security.

Today, organizations of all kinds are making a conscious effort to simplify the process, given the multiple places in which these three areas operate concurrently. The result is a more effective process of defining risk within the context of existing organizational rules and business objectives, and within the framework of compliance regulations, as shown in [Figure 1-15](#). The IT governance component creates stringent requirements for information security architectures, within the goal of adding business value, in addition to mitigating risk.

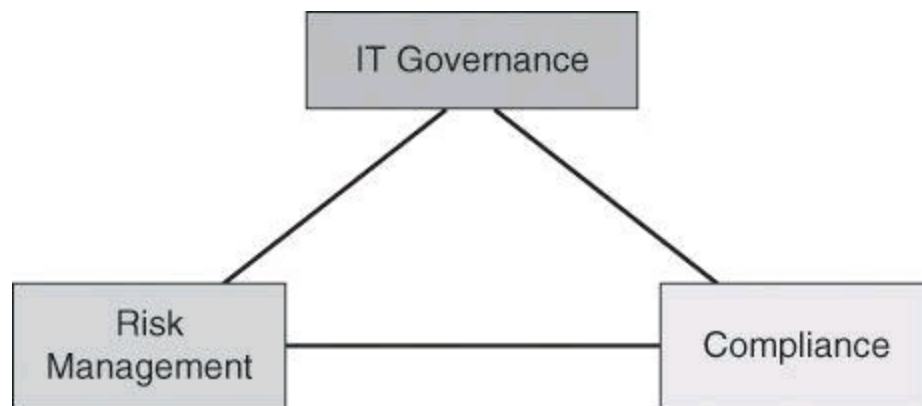


Figure 1-15. Organization-wide Integration of IT Governance, Risk Management, Compliance

This convergence results in an ideal framework and context to create a lifecycle approach to information security.

Secure Network Life Cycle

By framing security within the context of IT governance, compliance, and risk management, and by building it with a sound security architecture at its core, the result is usually a less expensive and more effective process. Including security early in the information process within the system design life cycle (SDLC) usually results in less-expensive and more-effective security when compared to adding it to an operational system.

A general SDLC includes five phases:

1. Initiation
2. Acquisition and development
3. Implementation
4. Operations and maintenance
5. Disposition

Each of these five phases includes a minimum set of security steps that you need to follow to effectively incorporate security into a system during its development. An organization either uses the general SDLC or develops a tailored SDLC that meets its specific needs. In either case, the National Institute of Standards and Technology (NIST) recommends that organizations incorporate the

associated IT security steps of this general SDLC into their development process.

Initiation Phase

The initiation phase of the SDLC includes the following:

- **Security categorization:** This step defines three levels (low, moderate, and high) of potential impact on organizations or individuals should a breach of security occur (a loss of confidentiality, integrity, or availability). Security categorization standards help organizations make the appropriate selection of security controls for their information systems.
- **Preliminary risk assessment:** This step results in an initial description of the basic security needs of the system. A preliminary risk assessment should define the threat environment in which the system will operate.

Acquisition and Development Phase

The acquisition and development phase of the SDLC includes the following:

- **Risk assessment:** This step is an analysis that identifies the protection requirements for the system through a formal risk-assessment process. This analysis builds on the initial risk assessment that was performed during the initiation phase, but is more in depth and specific.
- **Security functional requirements analysis:** This step is an analysis of requirements and can include the following components: system security environment, such as the enterprise information security policy and enterprise security architecture, and security functional requirements.
- **Security assurance requirements analysis:** This step is an analysis of the requirements that address the developmental activities required and the assurance evidence needed to produce the desired level of confidence that the information security will work correctly and effectively. The analysis, based on legal and functional security requirements, is used as the basis for determining how much and what kinds of assurance are required.
- **Cost considerations and reporting:** This step determines how much of the development cost you can attribute to information security over the life cycle of the system. These costs include hardware, software, personnel, and training.
- **Security planning:** This step ensures that you fully document any agreed upon security controls, whether they are just planned or in place. The security plan also provides a complete characterization or description of the information system and attachments of or references to key documents that support the information security program of the agency. Examples of documents that support the information security program include a configuration management plan, a contingency plan, an incident response plan, a security awareness and training plan, rules of behavior, a risk assessment, a security test and evaluation results, system interconnection agreements, security authorizations and accreditations, and a plan of action and milestones.
- **Security control development:** This step ensures that the security controls that the respective security plans describe are designed, developed, and implemented. The security

plans for information systems that are currently in operation may call for the development of additional security controls to supplement the controls that are already in place or the modification of selected controls that are deemed less than effective.

- **Developmental security test and evaluation:** This step ensures that security controls that you develop for a new information system are working properly and are effective. Some types of security controls, primarily those controls of a nontechnical nature, cannot be tested and evaluated until the information system is deployed. These controls are typically management and operational controls.
- **Other planning components:** This step ensures that you consider all the necessary components of the development process when you incorporate security into the network life cycle. These components include the selection of the appropriate contract type, the participation by all the necessary functional groups within an organization, the participation by the certifier and accreditor, and the development and execution of the necessary contracting plans and processes.

Implementation Phase

The implementation phase of the SDLC includes the following:

- **Inspection and acceptance:** This step ensures that the organization validates and verifies that the functionality that the specification describes is included in the deliverables.
- **System integration:** This step ensures that the system is integrated at the operational site where you will deploy the information system for operation. You enable the security control settings and switches in accordance with the vendor instructions and the available security implementation guidance.
- **Security certification:** This step ensures that you effectively implement the controls through established verification techniques and procedures. This step gives organization officials confidence that the appropriate safeguards and countermeasures are in place to protect the information system of the organization. Security certification also uncovers and describes the known vulnerabilities in the information system.
- **Security accreditation:** This step provides the necessary security authorization of an information system to process, store, or transmit information that is required. This authorization is granted by a senior organization official and is based on the verified effectiveness of security controls to some agreed upon level of assurance and an identified residual risk to agency assets or operations.

Operations and Maintenance Phase

The operations and maintenance phase of the SDLC includes the following:

- **Configuration management and control:** This step ensures that there is adequate consideration of the potential security impacts due to specific changes to an information system or its surrounding environment. Configuration management and configuration control procedures are critical to establishing an initial baseline of hardware, software, and firmware components for the information system and subsequently controlling and maintaining an accurate inventory of any changes to the system.

- **Continuous monitoring:** This step ensures that controls continue to be effective in their application through periodic testing and evaluation. Security control monitoring, such as verifying the continued effectiveness of those controls over time, and reporting the security status of the information system to appropriate agency officials are essential activities of a comprehensive information security program.

Disposition Phase

The disposition phase of the SDLC includes the following:

- **Information preservation:** This step ensures that you retain information, as necessary, to conform to current legal requirements and to accommodate future technology changes that can render the retrieval method of the information obsolete.
- **Media sanitization:** This step ensures that you delete, erase, and write over data as necessary.
- **Hardware and software disposal:** This step ensures that you dispose of hardware and software as directed by the information system security officer.

Models and Frameworks

The five-phase approach of the SDLC gives context to the process of designing, creating, and maintaining security architectures. It is based on NIST Publication 800-64 revision 2. Other frameworks and models exist, providing similar guidance to your security architecture:

- The ISO 27000 series is a comprehensive set of controls comprising best practices in information security. It is about information security, not IT security. It is also an internationally recognized information security standard, broad in scope and generic in applicability. It focuses on risk identification, assessment, and management. It is aligned with common business goals:

- Ensure business continuity
- Minimize business damage
- Maximize return on investments

ISO 27000 standards are much more commonly applied in commercial organizations than in government. Originally created as BS17799, this framework was first submitted in 1995, and revised in 1998, but was not adopted by the ISO until 1999. Significantly revised in 2005, it was formally converted to two related ISO/International Electrotechnical Commission (ISO/IEC) standards, 27001 and 27002.

- Control Objectives for Information and Related Technology (COBIT) provides good practices across a domain and process framework and presents activities in a manageable and logical structure. The good practices provided by COBIT represent the consensus of experts. These good practices are strongly focused more on control and less on execution.

These practices will help optimize IT-enabled investments, ensure service delivery, and provide a measure against which to judge when things do go wrong. COBIT is generally considered complementary to ISO/IEC 27001 and 27002.

- The Information Technology Infrastructure Library (ITIL) was developed under the

supervision of the Central Computer and Telecommunications Agency in the UK. ITIL is a set of eight practice guidebooks covering most aspects of IT service management. The fourth service management set is Security Management. ITIL Security Management is based on the code of practice in ISO 27002.

[Table 1-7](#) provides a summary of the different frameworks.

Table 1-7. Comparison of Frameworks

Framework	Strengths	Focus
COBIT	IT controls IT metrics	IT governance Audit
ISO 27000 series	Global acceptance Certification Security control	Information security Management system
ITIL	Processes Certification	IT service management
NIST 800 series	Detailed, granular Tiered controls Available for free	Information systems FISMA (federal government)

Network Security Posture

By assessing all aspects of the networked business environment, it is possible to determine the ability of the organization to detect, defend against, and respond to network attacks. The following are the key activities:

- **Security posture assessment (also known as vulnerability assessment):** The first step in planning network security requires an evaluation of the network security posture of the organization. The security posture assessment provides a snapshot of the security state of the network by conducting a thorough assessment of the network devices, servers, desktops, and databases. The effectiveness of the network security is analyzed against recognized industry best practices to identify the relative strengths and weaknesses of the environment and document specific vulnerabilities that could threaten the business. Because network security involves all aspects of the business, it is necessary to assess security from various perspectives, including the internal, external, dial-up, and wireless networks, and to provide recommendations on how to improve overall network security.
- **Internal assessment:** With so much attention devoted to threats and incidents by hackers, administrators may overlook the security of the internal, trusted network. The internal assessment is a controlled network attack simulation that is used to gauge the exposure present on internal systems, applications, and network devices. The assessment identifies the steps that are needed to thwart intentional attacks or unintentional mistakes from trusted insiders to effectively secure valuable information assets. To go beyond automated detection of vulnerabilities, you could simulate a real intruder in a controlled, safe manner

to confirm vulnerabilities manually. The assessment provides a more structured approach to identifying vulnerabilities that may go undetected. This secondary exploitation may include attempting to exploit trusted relationships between hosts, exploiting password weakness, or gaining administrative access to systems.

- **External assessment:** The goal of an external assessment is to quantify the security risk that is associated with Internet-connected systems. After researching and confirming the registration of Internet devices, assessors scan the device for external visibility. Because most services have inherent and well-known vulnerabilities, it must be determined whether the services offered are potentially vulnerable.
- **Wireless assessment:** The wireless assessment provides an evaluation of the security posture of the wireless network within the organization and identifies risks and exposures that are associated with a wireless deployment. Assessors analyze the wireless technology architecture and configurations to identify authorized and unauthorized access points and to recommend solutions to strengthen the security of the wireless infrastructure. Assessors also check outside customer buildings to find wireless network traffic leaking from the buildings.
- **Security posture assessment analysis and documentation:** This assessment quantifies the security posture of the organization network by using metrics and graphs. The report should also provide technical details, including analysis of each IP address, an explanation of methods that are used to compromise network devices and systems, and a description of the likelihood that an attacker will use that same approach. The report then prioritizes the vulnerabilities, recommends actions to correct the security risks, and details remediation steps that will prevent future exploitation.

Network Security Testing

Security testing provides insight into the other SDLC activities, such as risk analysis and contingency planning. You should document security testing and make the documentation available for staff involved in other IT and security-related areas. Typically, you conduct network security testing during the implementation and operational stages, after the system has been developed, installed, and integrated.

During the implementation stage, you should conduct security testing and evaluation on specific parts of the system and on the entire system as a whole. Security test and evaluation (ST&E) is an examination or analysis of the protective measures that are placed on an information system after it is fully integrated and operational. The following are the objectives of the ST&E:

- Uncover design, implementation, and operational flaws that could allow the violation of the security policy
- Determine the adequacy of security mechanisms, assurances, and other properties to enforce the security policy
- Assess the degree of consistency between the system documentation and its implementation

Once a system is operational, it is important to ascertain its operational status. You can conduct many tests to assess the operational status of the system. The types of tests you use and the frequency

in which you conduct them depend on the importance of the system and the resources available for testing. You should repeat these tests periodically and whenever you make a major change to the system. For systems that are exposed to constant threat, such as web servers, or systems that protect critical information, such as firewalls, you should conduct tests more frequently.

Security Testing Techniques

You can use security testing results in the following ways:

- As a reference point for corrective action
- To define mitigation activities to address identified vulnerabilities
- As a benchmark to trace the progress of an organization in meeting security requirements
- To assess the implementation status of system security requirements
- To conduct cost and benefit analysis for improvements to system security
- To enhance other lifecycle activities, such as risk assessments, certification and authorization (C&A), and performance-improvement efforts

There are several different types of security testing. Some testing techniques are predominantly manual, and other tests are highly automated. Regardless of the type of testing, the staff that sets up and conducts the security testing should have significant security and networking knowledge, including significant expertise in the following areas: network security, firewalls, IPSs, operating systems, programming, and networking protocols, such as TCP/IP.

Many testing techniques are available, including the following:

- Network scanning
- Vulnerability scanning
- Password cracking
- Log review
- Integrity checkers
- Virus detection
- War dialing
- War driving (802.11 or wireless LAN testing)
- Penetration testing

Common Testing Tools

Many testing tools are available in the modern marketplace that you can use to test the security of your systems and networks. The following list is a collection of tools that are quite popular; some of the tools are freeware, some are not:

- Nmap
- GFI LanGuard
- Tripwire
- Nessus
- Metasploit

- SuperScan by Foundstone, a division of McAfee

Many other excellent tools exist. This list is only a representative sampling.

Note

Visit <http://www.backtrack-linux.org> to download BackTrack 5, released in August 2011. BackTrack 5 is packed with hundreds of security tools to test and secure your network. Use BackTrack 5 responsibly and legally, which entails getting written permission from the organization where you would like to use BackTrack prior to using it.

Incident Response

Risk cannot be completely eliminated in some business environments.

Security Diminishing Returns and Residual Risk

Earlier I mentioned that a way to deal with risk is to reduce it by investing in security measures. The concept of diminishing returns applies to those security investments. Looking at [Figure 1-16](#), you will notice that each additional security investment reduces risk (at least in theory). However, also notice that each additional security investment yields a lower additional risk reduction than the previous investment. In economics, this is what is called *diminishing returns*. Also, notice that regardless of how many resources you dedicate toward mitigating a risk, you can never reduce it to zero. There will always be residual risk. If that residual risk is unacceptable for your organization, you could consider buying insurance against it. Buying insurance against a risk would be considered *transferring the risk*.




Figure 1-16. Security Investment: Diminishing Returns and Residual Risk

One way to eliminate risk is to simply withdraw from doing business at all, an unlikely scenario. For this reason, incident response has become an important component of the secure network life

cycle. The breadth and sophistication of threat vectors in information security has increased exponentially. Every day new techniques emerge, and the motivation of the attackers becomes increasingly aggressive, driven by political reasons, industrial espionage, and terrorism. Preventative measures help, but not all incidents can be prevented. Risk avoidance is unlikely; risk mitigation is more realistic.

It is, then, almost required to implement an incident response capability to streamline the incident detection capabilities, contain the impact of those incidents to minimize loss and destruction, reduce the scope of weaknesses, and restore services within the parameters of the organization.

Implementing an incident response plan effectively can be challenging because of the amount and scope of the resources needed. The first critical step is to deploy an effective intrusion detection and prevention capability. Even if the incident response plan is not in place, incident detection and prevention can provide a first line of response. However, incident response is not completely effective without framing it within an incident response plan. Assessing the current and potential business impact of incidents is critical. Other crucial factors include the implementation of effective methods of collecting, analyzing, and reporting data. Also, it is important to define the framework of communication between the teams involved (for example, technical teams, human resources, legal) and between the organization and external entities (such as other incident response teams and law enforcement).

Incident Management

The incident response process has several phases:

- **Preparation:** As with any other activity, preparation is the building block of incident response methodologies. Preparation creates the foundation for a sound incident response plan and lays the groundwork for an incident prevention culture within the organization. These are some examples of the tasks typically implemented during the preparation phase:
 - Prepare the facilities (such as a central coordination room and storage facilities for collected evidence) and the communication mechanisms (cell phones, contact and on-call information, and others).
 - Define the incident analysis hardware and software tools, such as protocol analyzers and forensics software.
 - Define prevention procedures, such as patch management and user awareness and training methods.
- **Detection and analysis:** With any luck, this is where the incident response team will spend most of its time. This phase starts with the definition of a threat vector classification scheme, in order to define detection and analysis capabilities more effectively per type of threat. Clearly defining the difference between events and incidents is critical. The incident response team should analyze and implement tools for log and event correlation, in order to facilitate the navigation across eventually thousands of security-related events. Efficiently and effectively identifying the business- and risk-relevant incidents out of thousands of events is a key component of the detection strategy. The best way to start is to define a sound framework to prioritize, document, and provide notice about incidents.
- **Containment, eradication, and recovery:** When an incident has been detected and

analyzed, it is important to contain it before the spread of the incident overwhelms resources or the damage increases. The containment strategy could start with a clear definition of tools to identify the attacker through IP addresses, usernames, and other means, followed by a clear definition of the context and time to perform this function (need for evidence preservation, time and resources to implement the strategy, sustainable service availability, and others). All containment strategies should also include steps to eradicate the threat and vulnerabilities, or at least mitigate them, and steps to recover operating systems, hardware components, and productive time. In light of this, ensure that the security policies are adapted to let remediation take place in a timely and effective manner if an attack is detected.

- **Post-incident activity:** This phase is crucial. The more the incident response team learns from past experience and (specially) mistakes, the more prepared it will be for future incidents. Focusing on how to collect and use data is a good first step. How to document what happened, especially the symptoms and fingerprint of the attack, should follow, leading to a full root-cause analysis. At this point, the incident response team should have a clear understanding of the options to go after the attacker (involve law enforcement, prosecution, and others).

Computer Crime Investigations

If you intend to successfully prosecute an individual who breaches your security, it is necessary to establish three things in most countries (in addition to evidence, the collection of which is covered next):

- **Motive:** Motive is concerned with why an individual performed the illegal act. As you investigate a computer crime, it is important to start with individuals who might have been motivated to commit the crime.
- **Opportunity:** Having identified a list of suspects, the next thing to consider is whether they had the opportunity to commit the crime. For example, if you can establish that three of the suspects were all participating in a wedding at the time of the security breach, they may have been motivated, but they did not have the opportunity. They were busy doing something else.
- **Means:** The means is an important thing to prove as well. Do not accuse someone who does not have the technical knowledge to accomplish the deed. Means is the ability to perform the crime. However, keep in mind that hacking tools have become easy for even a novice to use.

If you do not establish these three things, it is difficult to prove that the perpetrator is guilty of the offense should you decide to prosecute. When you can establish motive, opportunity, and means, and offer evidence, you are closer to a list of possible guilty parties.

Note

Different countries have different legal standards. Most countries and courts in the world accept this particular standard.

When working with computer data as part of a forensics case, you must maintain the integrity of the data if you will rely on the data in a court of law. It is difficult to maintain the integrity of the data in the virtual world of computers where it is trivial to change time stamps or any item of data. The flipping of a single bit can sometimes be all that is required to falsely establish an alibi.

Collection of Evidence and Forensics

Data collection is a volatile thing in the virtual world of computers. For this reason, a common procedure in response to security breaches is the immediate isolation of the infected system. Dumping the memory to disk is required because the system flushes the memory every time a device is powered off. Multiple copies of the hard drive are usually made after the device is powered down, to establish master copies. These master copies are usually locked up in a safe, and investigators use working copies for both the prosecution and the defense. You can answer any charges of tampering with data by comparing working copies to the master copy that has been secured and untouched since the beginning of the investigation.

It is important to note that when making copies of hard drives, a hardware write blocker must be used to ensure that the data on the source drive has not been modified by the copy. EnCase Forensic suite from Guidance Software is a product that uses hardware write blocker.

Laws and Ethics

This section describes key laws and codes of ethics that are binding on information systems security (infosec) professionals.

For many businesses today, one of the biggest considerations for setting security policies is compliance with the law. For that reason, it is important for infosec professionals to be at least conversant in the basics of law.

In most countries, there are three types of laws:

- **Criminal:** Concerned with crimes, and its penalties usually involve the risk of fines or imprisonment, or both. If fines are paid, they are usually to the court and are used to defray court costs.
- **Civil (also called tort):** Focuses on correcting wrongs that are not crimes. An example of a civil law case is if one company sues another company for infringing on a patent. The penalty in civil law is usually monetary, although there can also be performance requirements such as ceasing to infringe on the patent. If money is awarded, it is given to the party who won the lawsuit. Imprisonment is not possible in civil law.
- **Administrative:** Involves government agencies enforcing regulations. For example, a company may owe its employees vacation pay. An administrative court could force the company to pay and would probably also levy a fine that is payable to the agency. Therefore, in administrative law cases, monetary awards are often split between the government agency and the victim whose wrongs have been righted.

Ethics involves a standard that is higher than the law. It is a set of moral principles that adherents follow to be considered ethical. These ethics are often formalized in codes appropriately entitled “codes of ethics” by the professions formalizing the code.

The information security profession has a number of codes that have been formalized:

- International Information Systems Security Certification Consortium, Inc. (ISC)2 Code of Ethics
- The Computer Ethics Institute's Ten Commandments of Computer Ethics
- RFC 1087, "Ethics and the Internet," by the Internet Activities Board (IAB)
- Generally Accepted System Security Principles (GASSP)

Liability

Companies must take into account the legal liability for the country in which they reside. Take, for example, an Internet service provider (ISP) that has hundreds of e-businesses that rely on the ISP to run their websites with 100 percent uptime. If a hacker or a virus takes down this ISP, there is a chance for the ISP to be found liable, if it is discovered that the ISP did not take enough precautions or did not secure the network against internal or external threats.

In such cases, legal liability is likely to depend on what prevention technologies and practices are available and whether these technologies and practices are reasonably cost-effective to implement. While developing and implementing our security procedures, we must demonstrate due diligence and due care.

Showing due diligence includes everything from implementing technologies such as firewalls, intrusion-detection tools, content filters, traffic analyzers, and VPNs, to having best practices for continuous risk-assessment and vulnerability testing.

Due care is concerned with the operations and maintenance of the secure mechanisms put in place by practicing due diligence.

Lack of due care can lead to downstream liability. This is the case when a network is used by hackers as a springboard to conduct an attack against a third party. The victim of the attack could prosecute not only the hackers, but also the organization whose security was lax enough that its network was used as the launching pad for the attack.

Disaster Recovery and Business Continuity Planning

Business continuity planning and disaster recovery procedures address the continuing operations of an organization in the event of a disaster or prolonged service interruption that affects the mission of the organization. Such plans should address an emergency response phase, a recovery phase, and a return to normal operation phase. You should identify the responsibilities of personnel during an incident and the resources that are available to them.

In reality, contingency and disaster recovery plans do not address every possible scenario or assumption. Rather, they focus on the events most likely to occur and they identify an acceptable method of recovery. Periodically, you should exercise the plans and procedures to ensure that they are effective and well understood.

Business continuity planning provides a short- to medium-term framework to continue the organizational operations. The following are objectives of business continuity planning:

- Moving or relocating critical business components and people to a remote location while the original location is being repaired
- Using different channels of communication to deal with customers, shareholders, and

partners until operations return to normal

Disaster recovery is the process of regaining access to the data, hardware, and software necessary to resume critical business operations after a natural or human-induced disaster. A disaster recovery plan should also include plans for coping with the unexpected or sudden loss of key personnel. A disaster recovery plan is part of a larger process known as business continuity planning.

After the events of September 11, 2001, when many companies lost irreplaceable data, the effort put into protecting such data has changed. It is believed that some companies spend up to 25 percent of their IT budget on disaster recovery planning to avoid larger losses. Research indicates that of companies that had a major loss of computerized records, 43 percent never reopened, 51 percent closed within two years, and only 6 percent survived long term (<http://searchenterprisewan.techtarget.com/definition/disaster-recovery-plan> and http://en.wikipedia.org/wiki/Disaster_recovery).

Not all disruptions to business operations are equal. Whether the disruption is natural or human, intentional or unintentional, the effect is the same. A good disaster recovery plan takes into account the magnitude of the disruption, recognizing that there are differences between catastrophes, disasters, and nondisasters. In each case, a disruption occurs, but the scale of that disruption can dramatically differ.

- **Nondisaster:** A situation where a business process is unavailable for a given period of time
- **Disaster:** A situation that makes a facility unusable for an entire day or more
- **Catastrophe:** A situation that destroys the facility

A circular icon with a dotted border containing the text "Key Topic".

Key
Topic

Business Continuity Concepts

Building a business continuity plan requires extensive planning, with knowledge of the business requirements, budgets, and levels of risk the organization is willing to take. Some of the building block components, however, are more easily defined. The goal, from a rather simplified point of view, is to define objectives for the recovery of host computing systems that run the applications that support the business processes. These objectives are stated as the Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

RTO is the number of hours or days that management has set as the objective for resuming a business process or a system. RPO describes the age of the data you want the ability to restore to in event of a disaster. For example, if the RPO is 8 hours, systems should be restored in the state they were in no longer than 8 hours ago. The technical disaster recovery strategy depends upon meeting RTO and RPO specifications. The RTO and RPO requirements determine which option of the disaster recovery plan to implement. Recovery time and how current data is are key components in determining the level of service a business process requires in the event of a major disruption. To properly implement a disaster recovery plan, one must know the RTO and RPO that the organization is willing to accept in a disaster. The technical disaster recovery strategy of different options of recovery is based upon a combination of these requirements.

Key Concepts

Maximum Tolerable Downtime (MTD)

The total amount of time the system owner or authorizing official is willing to accept for a mission or business process outage or disruption, and includes all impact considerations.

Recovery Time Objective (RTO)

The maximum amount of time that a system resource can remain unavailable before there is an unacceptable impact on other system resources, supported mission, or business processes.

Recovery Point Objective (RPO)

The point in time, prior to a disruption or system outage, to which mission or business process data can be recovered (given the most recent backup copy of the data) after an outage.

Summary

To have a comprehensive security solution, it is important to cover all aspects of the operation of an organization. Comprehensive security requires suitable reliance on technical, physical, and administrative controls; implementing defense in depth; and developing an all-inclusive security policy. You will also be required to demonstrate forward thinking, taking into consideration the threats of tomorrow.

In this chapter you have learned that

- The confidentiality, integrity, and availability of the data need to be protected.
- Assets, vulnerabilities, and countermeasures can be classified to assist in developing a comprehensive set of security policies.
- New trends and threats are appearing frequently in the borderless network environment where we are evolving.
- To provide a comprehensive security solution, it is essential that there be a combination of technical, physical, and administrative controls in place.
- Defense in depth is a philosophy used to provide layered security to a system by using multiple security mechanisms.
- A security policy is a set of objectives for the company, rules of behavior for users and administrators, and requirements for system and management that collectively ensures the security of network and computer systems in an organization.

References

For additional information, refer to these resources.

Publications

Harris, S. *CISSP All-in-One Exam Guide, Fifth Edition* (McGraw-Hill Professional, 2010).

McClure, S., Scambray, J., and Kurtz, G. *Hacking Exposed, Sixth Edition* (McGraw-Hill

Professional, 2009).

McClure, S., Scambray, J., and Kurtz, G. *Hacking Exposed, Seventh Edition* (McGraw-Hill Professional, 2012).

NIST SP 800-27 Rev A, *Engineering Principles for Information Technology Security*.

NIST SP 800-42, *Guidelines on Network Security Testing*.

NIST SP 800-64 Rev. A, *Security Considerations in the Information System Development Life Cycle*.

Richardson, R. *2010-2011 CSI Computer Crime and Security Survey* (<http://gocsi.com/survey>).

Wood, C. *Information Security Policies Made Easy, Version 11* (Information Shield, 2009).

Web Resources

Insecure.org, <http://www.insecure.org/nmap/>

SecurityFocus, <http://www.securityfocus.com/>

Security-Solutions.net, <http://www.security-solutions.net/download/index.html>

The GNU Netcat Project (G. Giacobbi), <http://netcat.sourceforge.net/>

The Jargon File, <http://www.catb.org/~esr/jargon/html/index.html>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the appendix, “[Answers to Chapter Review Questions](#).”

1. Which are the three primary objectives of security?

- a. Integrity
- b. Confidentiality
- c. Antireplay functionality
- d. Authentication
- e. Availability

2. Which are the three categories of controls?

- a. Administrative
- b. Executive
- c. Managerial
- d. Technical
- e. Physical

3. Show that you understand the different types of controls by matching them with their related technology.

Type of controls

- a. Preventative
- b. Deterrent

- c. Detective Technologies
- d. Motion sensor
- e. Video surveillance
- f. Lock

4. Match the different types of hackers and the like with their appropriate description.

Hacker types

- a. White hat
- b. Black hat
- c. Gray hat
- d. Blue hat
- e. Cracker
- f. Phreaker
- g. Script kiddy
- h. Hacktivist

Hacker descriptions

- i. Bug tester
- j. Hacker with little skill
- k. Unethical hacker
- l. Hacker of telecommunication systems
- m. Ethically questionable hacker
- n. Hacker with a political agenda
- o. Synonymous with black hat hacker
- p. Breaks security for nonmalicious reasons

5. Organize the following steps in the order in which they are used to compromise targets and applications.

- a. Escalate privilege
- b. Leverage the compromised system
- c. Perform footprint analysis
- d. Install back doors
- e. Enumerate applications and operating systems
- f. Gather additional passwords and secrets
- g. Manipulate users to gain access

6. Which of the following is (are) not part of the technical policies. (Select all that apply.)

- a. End-user policy
- b. Acceptable usage policy

- c. Email policy
- d. Governing policy
- e. Rainbow Series
- f. Network policy
- g. Common Criteria Standard
- h. Wireless policy

7. Reorder the classification levels of the private sector, from the least secure document to the most secure document.

- a. Confidential
- b. Private
- c. Public
- d. Sensitive

8. Which of the following is not a criterion used to classify data?

- a. Value
- b. Age
- c. Useful life
- d. Copyright
- e. Personal association

9. Match each of the following information classification roles with its definition.

Roles

- a. Owner
- b. Custodian
- c. User

Definitions

- d. Responsible for using the data
- e. Responsible on a day-to-day basis for the classified data
- f. Ultimately responsible for the data

10. Which of the following is a technical control?

- a. Network Admission Control system
- b. Security policies and standards
- c. Security audits
- d. Security awareness training
- e. Change and configuration management

11. Which of the following is not a characteristic of defense in depth?

- a. Security mechanisms back each other up.
- b. Security mechanisms do not depend on each other.

- c. Does not require IDS or IPS.
- d. The weakest links can be augmented so that single points of failure can be eliminated.

12. Match the definition with the appropriate attack method.

Definitions

- a. Searching a network host and open ports
- b. Capturing electrical transmission
- c. Hiding information within a transmission
- d. Intercepting traffic that passes over a physical network

Attack methods

- e. Packet sniffing
- f. Man-in-the-middle
- g. Emanation capturing
- h. Covert channel
- i. Impersonation
- j. Port scanning

13. Reorder the phases of a system development life cycle.

- a. Operations and maintenance
- b. Initiation
- c. Disposition
- d. Acquisition and development
- e. Implementation

14. Which of the following security concepts limits a user's rights to the lowest possible level needed to perform his tasks?

- a. Need to know
- b. Least privilege
- c. Universal participation
- d. Diversity of defense

Chapter 2. Security Strategy and Cisco Borderless Network

In this chapter, you learn about the Cisco vision of a security strategy for Cisco Borderless Networks. You read about the architecture, components, and underlying technologies, and see Cisco products and solutions within the architecture.

In this chapter, you learn about the following Cisco Borderless Networks topics:

- Cisco Borderless Network Architecture
- Cisco security portfolio of products solving issues of Borderless Networks
- Cisco SecureX Architecture presenting its features and benefits
- Cisco threat control and containment products and technologies
- Cisco content security products and technologies
- Cisco VPN solutions and technologies
- Security management products and technologies

Borderless Networks

Borders are dissolving in multiple dimensions in current networks.

The device border is being dissolved by the phenomenon of mobility. Expect to see billions of new network mobile devices moving into enterprises worldwide over the next few years. If you consider constant reductions and streamlining in IT budgets, organizations face serious challenges to support a growing number of mobile devices at a time when resources are being reduced.

The application border is also dissolving. This is driven by market forces in cloud computing and cloud services. Organizations of all kinds are taking advantage of offerings such as Software as a Service (SaaS) and Infrastructure as a Service (IaaS) to reduce costs and simplify the deployment of new services and applications. These cloud services add challenges in visibility (how do you identify and mitigate threats that come to and from a trusted network?), control (who controls the physical assets, encryption keys, and so on?), and trust (do you trust cloud partners to ensure that critical application data is still protected when it is off the enterprise network?).

The location border is being blurred by constant changes to the workplace experience and millions of professionals that engage in work-related activities equally from home, the office, or while traveling. Borders are blurring in the organization as well, between consumers and workers and between the various functions within the organization. The borders between the company and its partners, customers, and suppliers are also fading. As a result, the network is having an increasing demand to connect anyone, anything, anywhere, at any time.

Taken together, these trends are driving consistent innovation of network-based business processes as the Internet morphed from a network to a platform for integrating supply chains, exchanging information for business intelligence, outsourcing entire business functions, and establishing new relationships. This trend will only accelerate as cloud computing matures.

The market changes also represent a challenge to security teams within the organization. They now need to manage uncontrolled consumer devices coming into the network, and provide seamless and

context-aware services to users all over the world. The location of the data and services they access is almost irrelevant. Their location could be internal to the organization, or it could be in the cloud. This makes protecting data and services a challenging proposition.

The paradigm of access anytime, anywhere, using any device changes the network management approach that has traditionally been used to design and maintain networks. Scalability, availability, performance, and other metrics become much more complex as the IT organization moves beyond the infrastructure it can directly control. These metrics are no longer linear, but multidimensional. IT organizations need to manage the same metrics, but across the device, the application, and the location border.

Cisco Borderless Network Security Architecture

The Cisco Borderless Network Architecture, shown in [Figure 2-1](#), addresses these issues using a comprehensive set of integrated components:

- **Borderless end zone:** Implemented by security controls aimed at protecting end devices and preventing intrusion to and from those end devices. Malware, viruses, and other types of malicious code are commonplace. The borderless end zone is built to detect these threats and eradicate them, or at least mitigate them, even before users connect to the network. As users connect from multiple locations, the network reacts to suspicious behavior by controlling access based on the security posture of the connecting endpoint. This reaction mitigates the impact of endpoint-based threats and helps contain their effect in a cost-effective manner.
- **Borderless Internet:** Implemented by scanning and enforcement engines, which are deployed onsite and offsite and managed by the organization itself but also by cloud providers. These scanning and enforcement engines inspect network content from Layer 2 through Layer 7, assuming the role of firewalls, intrusion detection systems (IDS), and intrusion prevention systems (IPS), network proxies, web gateways, and so on. These engines are located throughout the corporate network, but also rely on cloud-based services for security intelligence, real-time updates, and scanning support for mobile users, remote offices, and public network-resident devices such as network sensors or supervisory control and data acquisition (SCADA) systems.
- **Borderless data center:** Data centers deserve special attention in a cloud-driven business environment. The Cisco Borderless Network Architecture leverages existing infrastructure components such as Cisco Adaptive Security Appliances (ASA), IPS sensors, and switching components such as Cisco Catalyst and Cisco Nexus switches, layering virtualized components on top to provide security solutions that are built for “in the cloud” deployment, “by the cloud” deployment, and “for the cloud” deployment.
- **Policy management layer:** Security policy is managed in central locations and then enforced throughout the network based on context-specific variables. When an employee on an untrusted network in Moscow checks email via Outlook Web Access (OWA), that employee can view messages but is precluded from downloading attachments.
- **Borderless Network Services:** Professional and technical services are provided by Cisco and its partners. Cisco and its partners help an organization achieve the full benefits of the deployment of the Cisco solutions by providing consulting services in areas such as

wireless, unified communication, data center, and, of course, network security, to name a few.

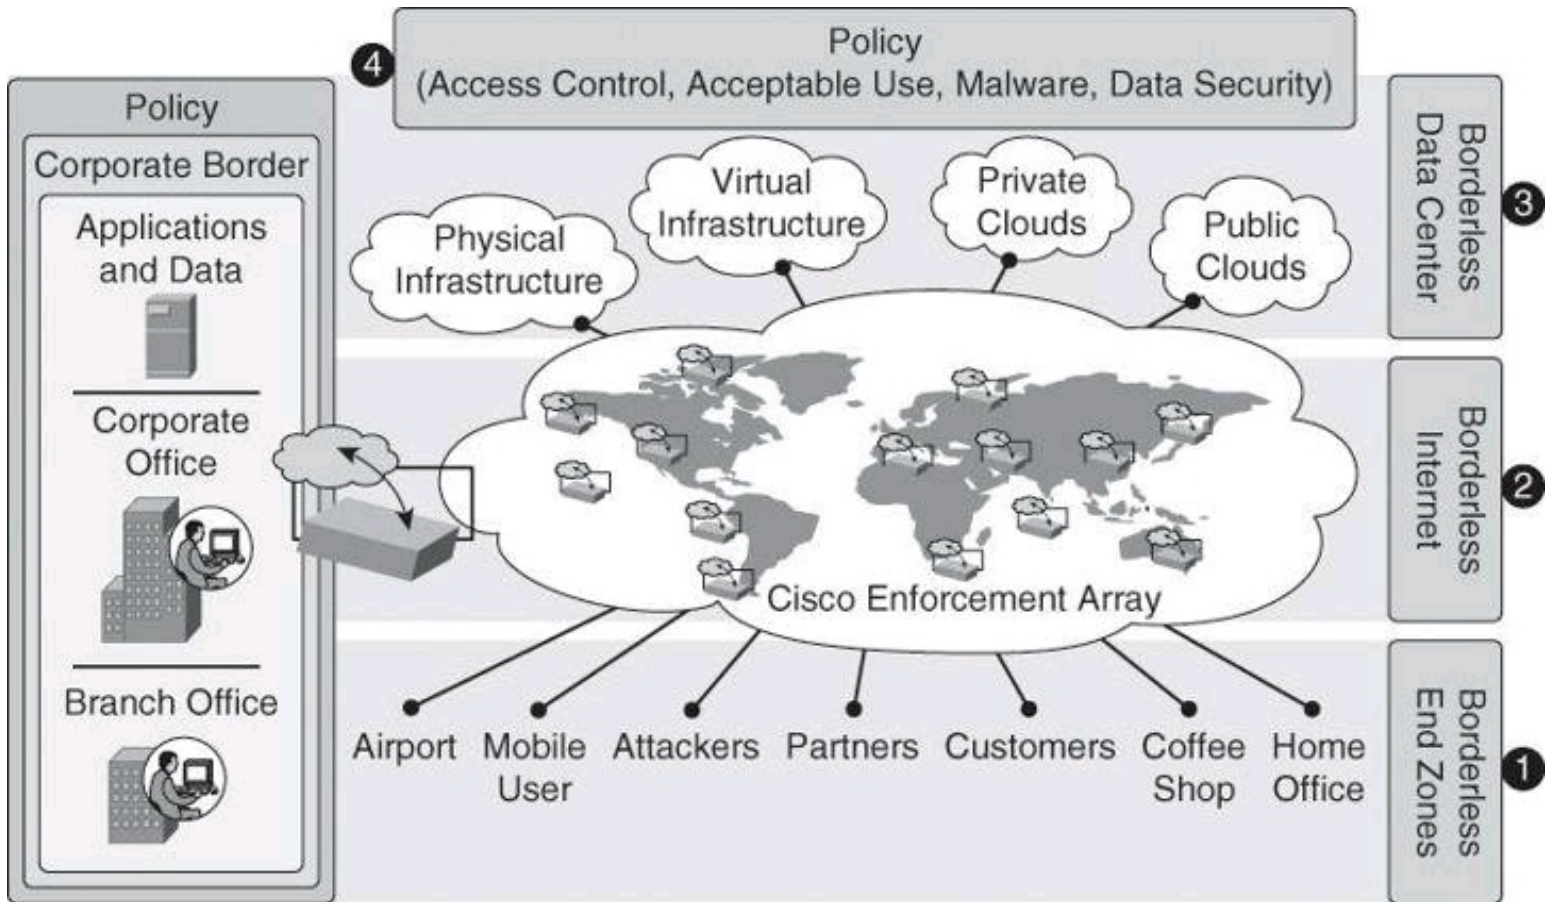


Figure 2-1. Cisco Borderless Network Security Architecture

The type of distributed, shared security intelligence, in the form of log files, vulnerability scanning data, system configuration, and event data, is available on a “publish and subscribe” basis in industry-standard formats. Additionally, onsite security information is supplemented by security telemetry in the cloud. All of this data will be available to analysts for event detection and forensics, but the network will also consume this data itself in order to react to anomalies, enforce policies, or block attacks in real time.

Borderless End Zone

Each major component of the architecture results in specific benefits. The borderless end zone offers deployment flexibility and strong security services in multiple dimensions as users connect to the network.

PCs and mobile devices are active participants in borderless network security in several ways. In borderless security, all endpoints authenticate themselves to the organization’s network and submit to a health and configuration check before gaining access. This is true for employees and third-party users. Once admitted to the network, endpoints and users are given access to specific applications and data based upon parameters such as user role, location, time of day, and so on. This does not require multiple authentications. Rather, the network understands whether a user is working from an office desk, from a home PC, or from an untrusted network around the world and enforces policies for access controls, data confidentiality, and encrypted transport accordingly. Finally, endpoints continue to have onboard security defenses, but this will evolve from current fat client software to a cloud-

based model—especially for mobile devices with limited resources.

The Cisco AnyConnect client is one of the building blocks of this architectural component, providing multiple layers of defense even before users connect to the network, leveraged around Secure Sockets Layer (SSL) VPNs.

[Figure 2-2](#) provides a summary of the advantages of the borderless end zone. Many of the concepts presented in this figure are beyond the scope of this book but are explored in greater details in other Cisco literature and courses.

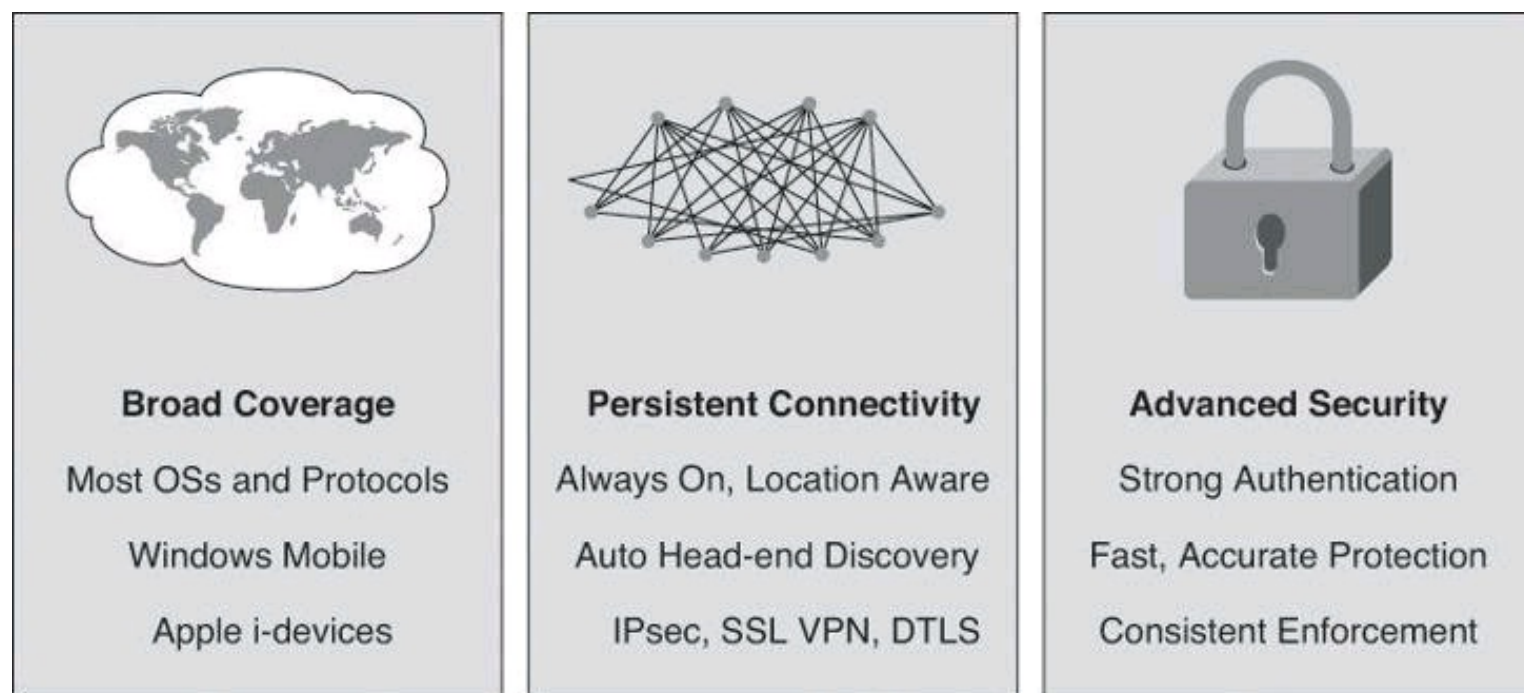


Figure 2-2. Borderless End Zone Advantages and Technologies

The other side of the equation is represented by enforcement engines that make up the security array. Cisco has a strong security product line that includes firewalls, IPSs, web security, and email security. These detection and enforcement engines are evolving in two ways:

- They are built to be highly scalable and interpret and control traffic at the application level. This capability allows the network to catch up with malicious attackers, who are increasingly utilizing application layer threats to perpetrate their plans.
- They are available in multiple forms: appliance, module in a switch, or router and hybrid hosted.

Borderless Internet

The combination of intelligent endpoints and an array of scanning and enforcement engines makes the network work as a system with integrated, collaborative, and adaptive security. This is a sharp contrast with a point product approach, where function-specific products create “islands of security” constrained by independent threat, systems, and policy management. These disparate point tools are antithetical to the dynamic and integrated security that is required for borderless networks to fulfill the “any user, any content, any location, any time” nature.

The Cisco Borderless Network Architecture demands teamwork and cooperation akin to that of a symphony orchestra, where all of the individual “instruments” come together to achieve a common

goal, as shown in [Figure 2-3](#).

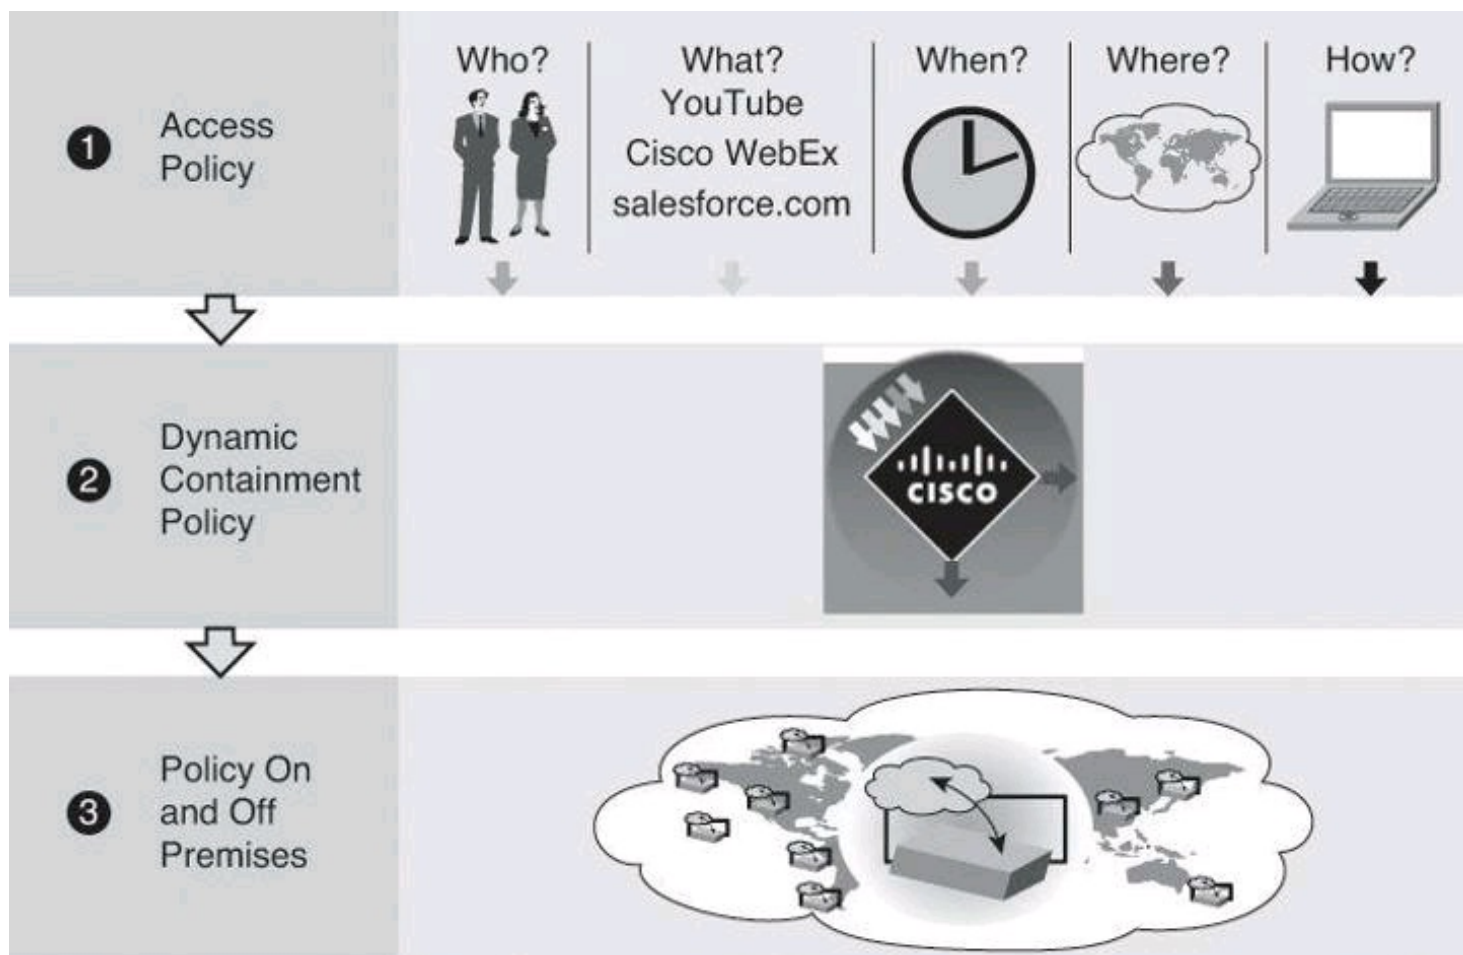


Figure 2-3. Intelligent Access Policies Provides Context Awareness for Adaptive Securing of Endpoints

Given the “borderless” aspect of Borderless Networks, security defenses must also extend beyond the enterprise network to mobile devices and cloud services. Using the Cisco Borderless Network Architecture, corporate LANs and WANs coordinate activities with service provider networks, distributed endpoints, and the Internet itself. No one vendor has all of the security products or resources to solve this challenge alone. Borderless network security requires industry ingenuity and collaboration. Cisco partners with leading security and networking vendors to extend the reach of the system beyond the confines of corporate networks.

Borderless Data Center

The borderless data center component of the architecture provides solutions that adjust to data center environments. [Figure 2-4](#) illustrates the progression of service deployment in a heavily virtualized environment, such as that of data center.

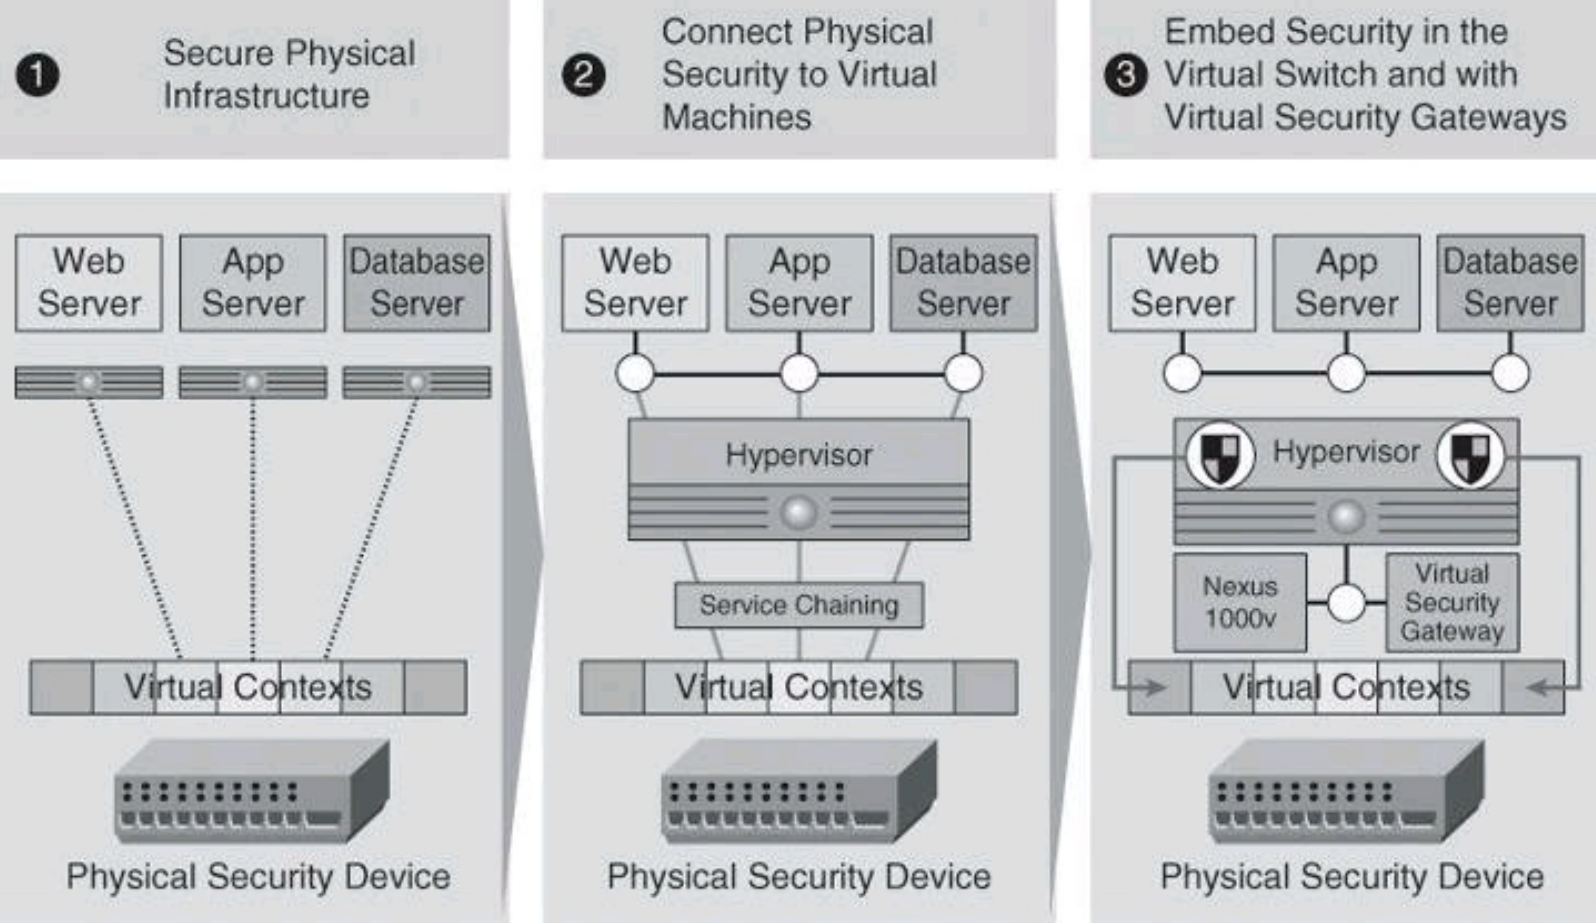


Figure 2-4. Evolution of Security in the Data Center

Virtual appliances provide trusted access to virtualized data centers in enterprise and cloud provider environments while meeting virtualization requirements of dynamic policy-driven operation, mobility-transparent enforcement, and scale-out deployment for dense multitenancy, where one service provider might be hosting for multiple different clients. But it does not stop there. Network connectivity is also virtualized and embedded into virtual machines. Virtual switches connect virtual machines within a single physical device, integrating network switch ports into the hypervisor.

An example of this architecture is the Cisco Virtual Security Gateway (VSG), a security virtual appliance that leverages the Cisco Nexus 1000V switch to provide zone-based trust services to virtual machines. Cloud providers can leverage the Cisco VSG to significantly enhance their multitenant scaling and also generate additional revenue by offering security and monitoring services to their customers. In effect, the Cisco VSG enables a broad set of multitenant workloads having varied security profiles to share a common compute infrastructure on premises (in the enterprise virtual data center and private cloud) or off premises (in the service provider cloud). By associating one or more virtual machines (VM) into distinct trust zones, the Cisco VSG ensures that access to trust zones is controlled and monitored through established security policies.

Policy Management Layer

The policy management layer is a higher-level function that can span multiple operational devices. Cisco's vision is to make the interface between policy and enforcement systems open and built on industry standards. Therefore, if a customer chooses to use a Cisco application entitlement system and an enterprise data loss prevention (DLP) policy manager from RSA, both should work smoothly with existing network infrastructure. This significantly reduces operational complexity for Cisco customers

because one set of policies can be enforced across a wide variety of infrastructure.

Borderless Network Services

Cisco Borderless Network Services covers many areas listed in [Table 2-1](#). However, our focus for this book is restricted to the security solutions. [Table 2-1](#) provides only a restricted list of SecureX core components. SecureX solutions and components will be presented in the next section of this chapter.

Table 2-1. Cisco Borderless Network Services

Category	Network Service	Solutions
Motion	Cisco Motion	Cisco CleanAir
		Cisco ClientLink
		Cisco VideoStream
Energy Management	Cisco EnergyWise	Cisco Services for Energy Management
Security	Cisco SecureX	Cisco Security Intelligence Operations
		Cisco TrustSec
		Cisco AnyConnect
Application Performance	Cisco Application Velocity	Cisco Wide Area Application Services
		Cisco UCS Express
Multimedia Optimization	Cisco Enterprise Medianet	Cisco Medianet Readiness Assessment Service
Management	Management solutions	Cisco Prime for Enterprise
		Cisco Prime LAN Management Solution
		Cisco Prime Network Analysis Module
		Cisco Prime Collaboration Manager
		Cisco Prime Network Control System

Borderless Security Products

The architectural approach to security found in the Cisco Borderless Network Architecture, shown in [Figure 2-5](#), results in distinct categories of Cisco products, technologies, and solutions. These product categories are built into the design of borderless network security and apply to multiple components of the architecture. For instance, a Cisco ASA can be found as an on-premises scanning engine in the borderless Internet component, but also as a firewall in the borderless data center.

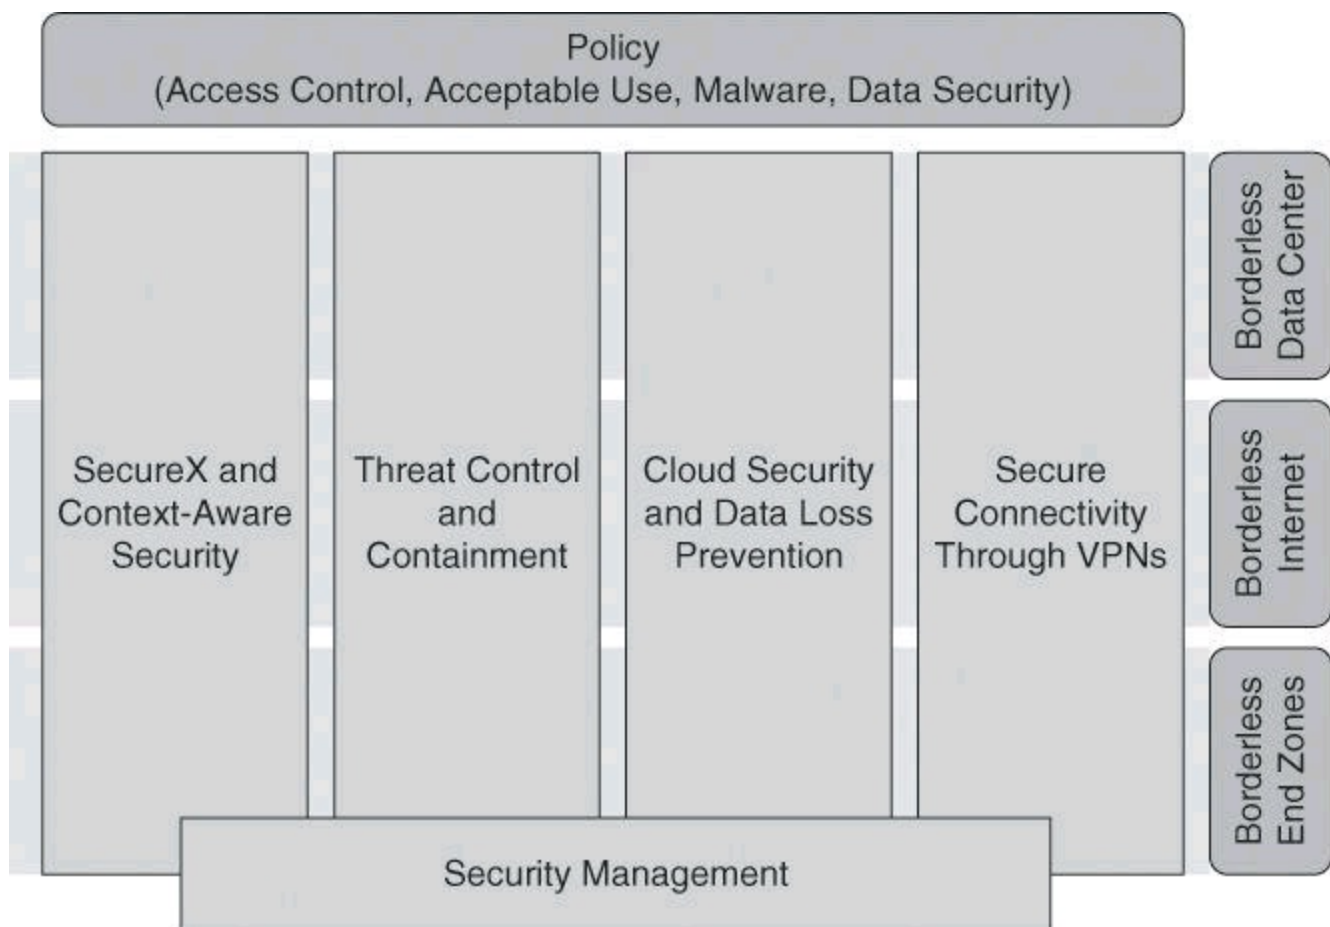


Figure 2-5. Policy Management Provided by Cisco Borderless Security Products

In the following sections, we look at the pillars of security found in the Cisco Borderless Network approach.

SecureX, a Context-Aware Security Approach

Cisco SecureX is an access control strategy that allows for more effective, higher-level policy creation and enforcement for mobile users, connecting from anywhere, at any time, using any device. Because Cisco SecureX uses a broad array of parameters for policy, it allows for much more effective security and enables situational awareness, as represented in [Figure 2-6](#). Instead of many complex firewall rules, security policy can now be based on context. For example, a salesperson accesses a global sales forecast from her laptop while in California; when she attempts to access that same sales forecast from the business center of the hotel where she is staying in China, the computer used is not detected as a corporate asset, and the request would be disallowed.

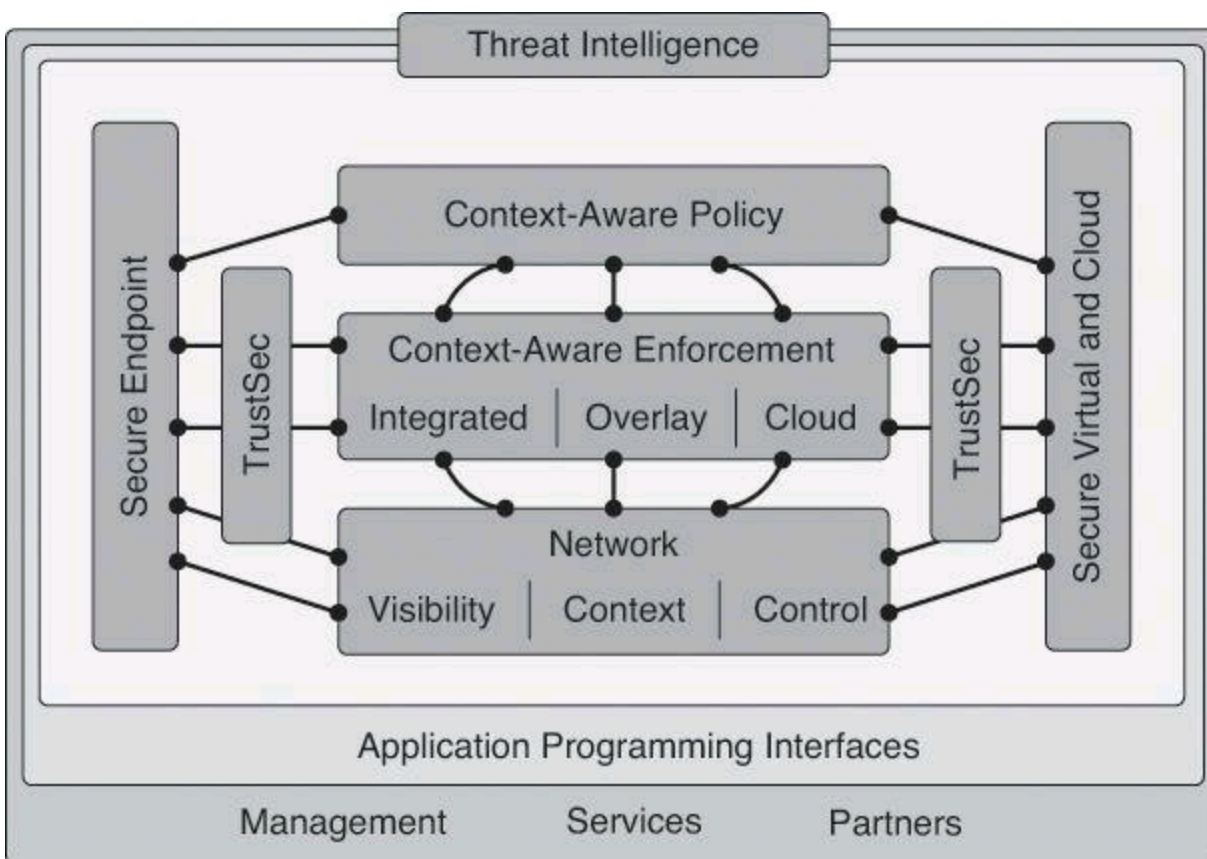


Figure 2-6. SecureX Context Awareness

This sort of intelligent policy enforcement uses next-generation scanning elements that are meshed into the new Cisco SecureX Architecture. Within this new architecture, the next-generation endpoint is able to automatically find the nearest scanning element somewhere in the virtual security fabric and to make a seamless connection. The behavior of a hacker halfway around the world is noted, that information is shared, and traffic from the servers of the hacker is blocked because your network now knows that it comes from someone that it cannot trust. Policy is centrally managed, but intelligence is gathered globally, with highly distributed enforcement.

[Table 2-2](#) provides a list of Cisco security solutions that are part of the Cisco SecureX Architecture as presented in the Cisco SecureX product brochure.

Table 2-2. Cisco SecureX Security Products

Category	Product
Secure Network	Cisco ASA 5500 Series Adaptive Security Appliance
	Cisco Intrusion Prevention System
	Cisco Integrated Services Router Generation 2
	Cisco Security Manager
Secure Email and Web	Cisco IronPort Email Security—Cloud, Hybrid, and On Premises
	Cisco Web Security—Cloud and On Premises
	IronPort Security Management Appliance
Secure Mobility	Cisco AnyConnect Secure Mobility Client
	Cisco Adaptive Wireless IPS Software
	Cisco Virtual Office
Secure Data Center	Cisco ASA 5585-X Adaptive Security Appliance
	Cisco Catalyst 6500 ASA Services Module
	Cisco Virtual Security Gateway (VSG)
Secure Access	Cisco Identity Services Engine
	Cisco Secure Access Control System

SecureX Core Components

To enforces security policies across the entire distributed network, not just at a single point in the data stream, SecureX relies on the following core components.

Context-Aware Policies

At the core of the Cisco SecureX Architecture is context awareness, which allows enforcement elements (such as firewalls, web proxies, IPS sensors, and even network infrastructure elements such as routers and switches) to define the access policy by using information such as the identity of the user, the security posture of the connecting device, the point of access to the network, and many other components.

Cisco Security Intelligence Operations (SIO)

Many of the Cisco security products listed in [Table 2-2](#) stay ahead of the latest threats by using real-time threat intelligence from Cisco Security Intelligence Operations (SIO). Cisco SIO is the world's largest cloud-based security ecosystem, using almost a million live data feeds from deployed Cisco email, web, firewall, and IPS solutions. Cisco SIO will be discussed further later in this chapter.

TrustSec

Cisco TrustSec relies on context awareness to enforce policy and at the same time provide access flexibility and situational access control.

Traditional access control parameters, such as the user identity, are considered. The user identity is the base component that the network gathers to initiate the determination of context. IEEE 802.1X is typically used to define the identity of those entering the network. Nonauthenticated devices, such as IP phones and printers, are also considered.

Other conditions are gathered from the Cisco AnyConnect client and the array of enforcement devices. Allowing an employee with high security clearance into the network on a weekday, using an uninfected corporate laptop from a campus office wired network, is not the same as allowing the same person (same identity) on a weekend evening to use an unknown PC, possibly infected with a virus, from a hotspot Internet café. AnyConnect and the device at the point of entry to the network determine this information.

The result is a different access policy and conditions that are based on the situational context of connectivity of individuals and devices, as illustrated by [Figure 2-7](#).

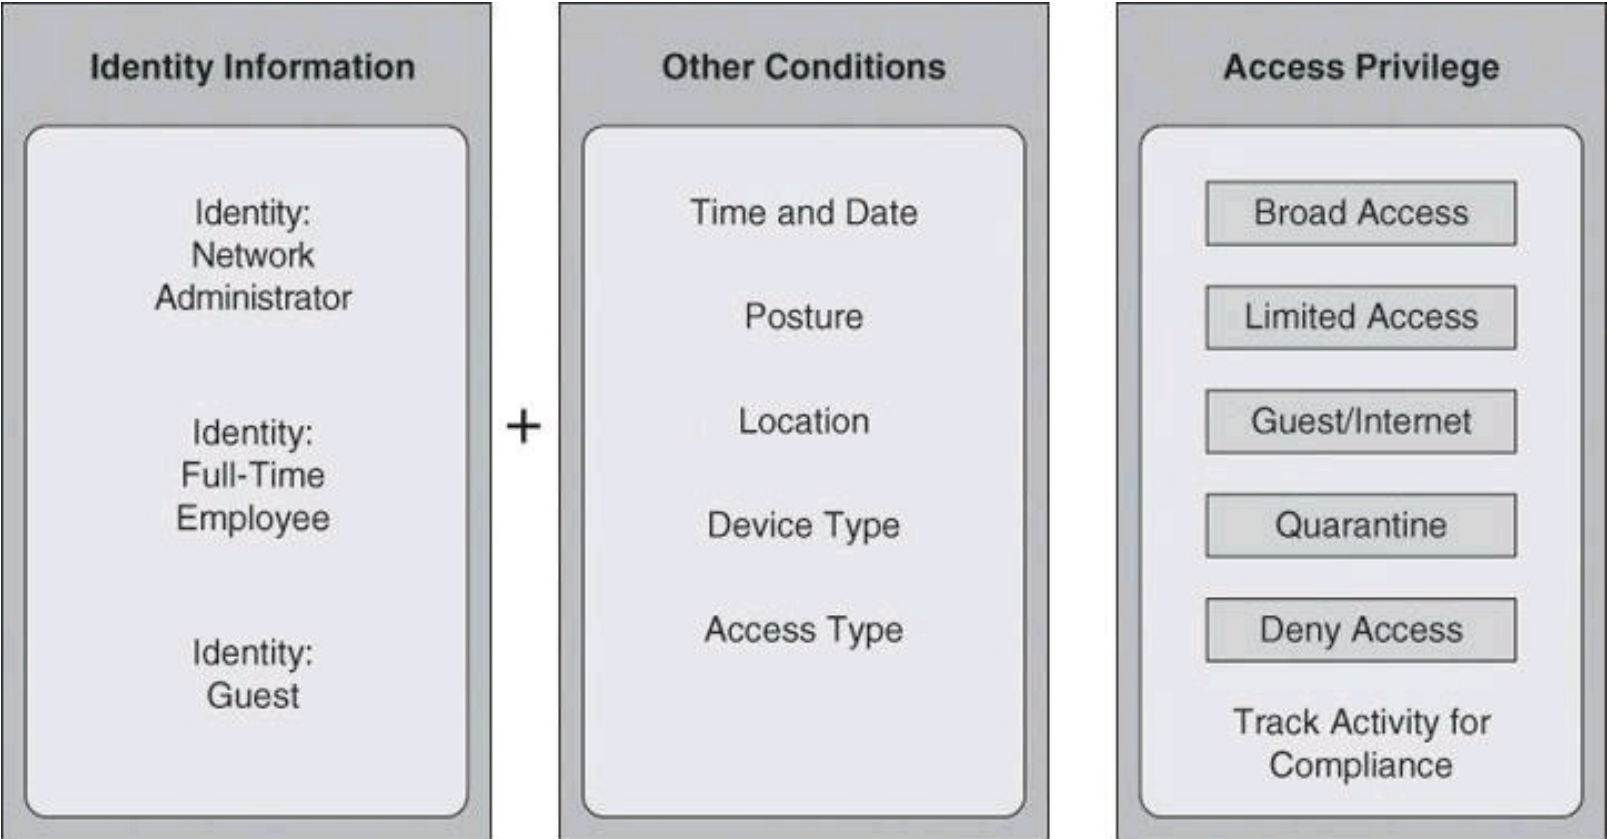


Figure 2-7. Context-Aware Policies Decide Access Privileges

The same individual might be allowed into corporate resources with broad access under certain conditions, while forced into a quarantine area to update the antivirus software or unlock a user password under other conditions. The network, as a system, adapts to each circumstance.

So, as explained above, Cisco TrustSec extends context awareness through policy-based access control for any user and any device seeking access to the distributed network. [Figure 2-8](#) illustrates the components of this function.

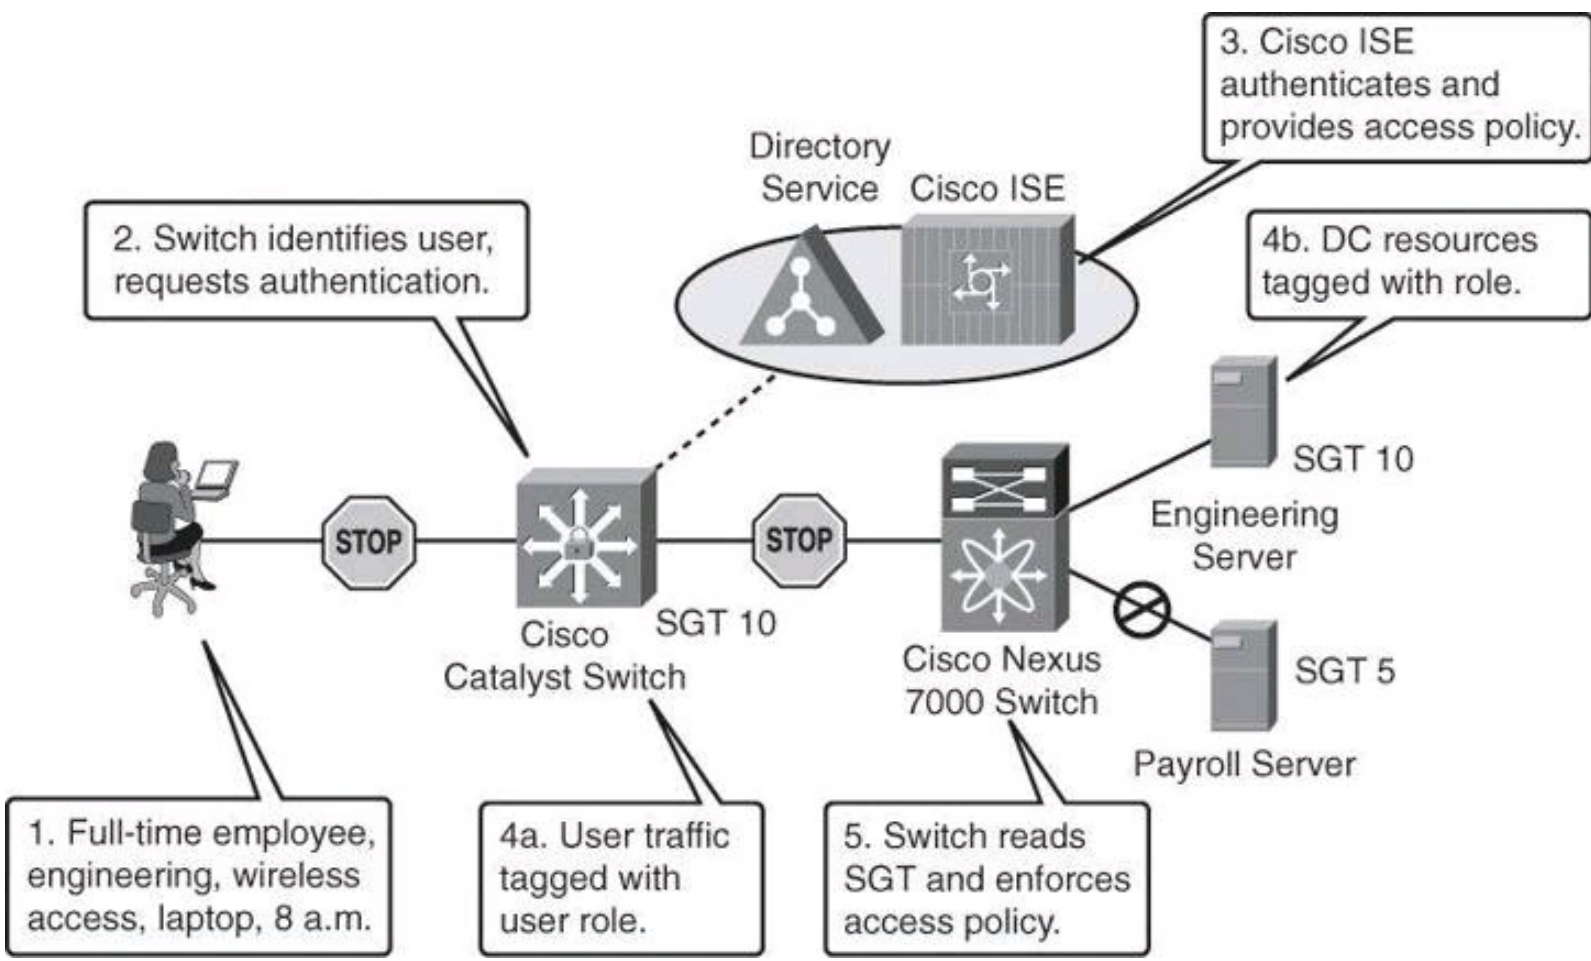


Figure 2-8. Secure Resources Using TrustSec

First, users are authenticated and authorized based on who, what, where, when, and how. Endpoint devices are also analyzed to determine if they meet corporate security policy before being granted access. Nonauthenticating devices, such as IP phones or video cameras, are also automatically identified and inventoried. Note that more and more endpoint devices, such as network printers, are 802.1x capable.

Next, traffic from any authenticated device is tagged with a unique, context-based access policy marker known as a security group tag (SGT). Network devices along the data path read this tag and enforce its associated policy by restricting access to predetermined network destinations and resources. The devices do so by using security group access control lists (SGACL).

Finally, Cisco TrustSec provides data confidentiality. This is known as MAC Security (MACsec). For example, a policy may require that any employee from the finance department accessing the payroll server must have their data secured. Cisco TrustSec understands this policy and can direct the network to dynamically encrypt the data of the user.

In summary, TrustSec restricts user access using Security Group Access (SGA). Access policy is inserted as SGTs into devices. The SGA reads and enforces policy tags on TrustSec-enabled Cisco switches along data paths using SGACLs.

AnyConnect

The Cisco AnyConnect client provides a secure connectivity experience across a broad set of PC- and smartphone-based mobile devices. The enforcement devices provide posture assessment, access control services, and policy enforcement. They also start the encryption process to provide

confidentiality via MACsec.

The Cisco AnyConnect Secure Mobility Solutions provide an innovative way to protect mobile employees on PC-based or smartphone platforms. These solutions use SSL or IP Security (IPsec) VPNs to deliver a more seamless, always-on, and always-protected experience to end users, while enabling IT administrators to enforce policies and block malware with cloud-based or hybrid web security. Cisco AnyConnect Secure Mobility is available to computers, tablets, and smartphones. For a complete list of Cisco AnyConnect Secure Mobility Client for Mobile Platforms, visit http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5743/ps5699/ps10884/datasheet_c78-678242.html.

The Cisco AnyConnect Secure Mobility Client is a lightweight, highly modular security client that provides easily customizable capabilities. The client provides broad operating system support from traditional operating systems such as Windows and Linux, to mobile devices such as smartphones and tablets (Apple iOS, Windows Mobile, Android, and others), as illustrated in [Figure 2-9](#).

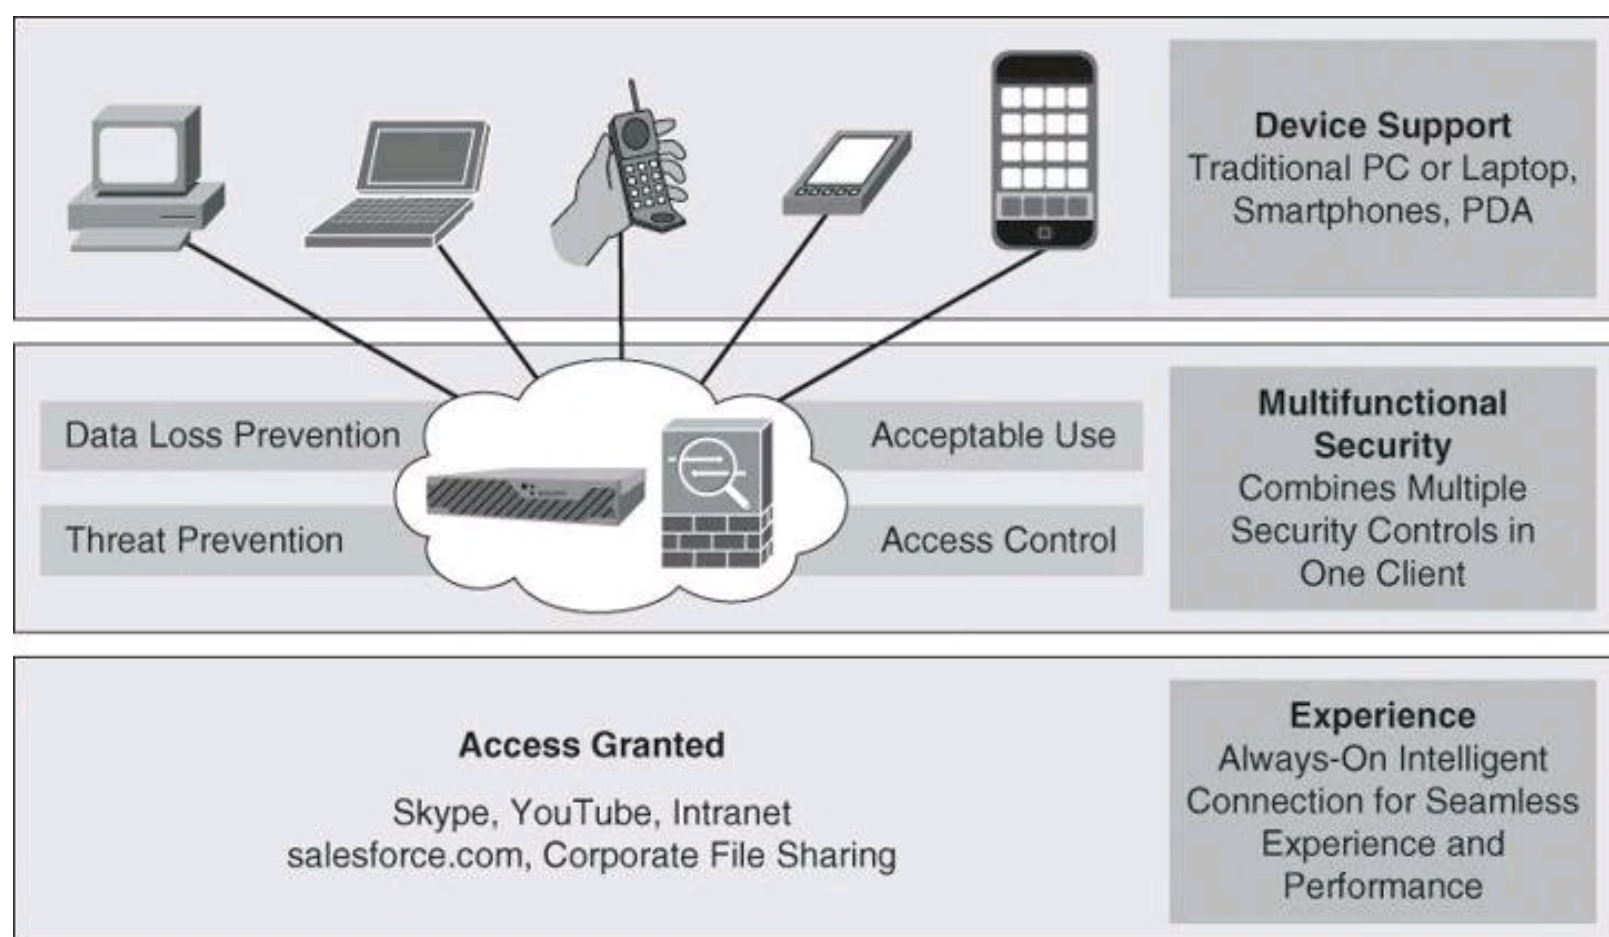


Figure 2-9. AnyConnect at Work

Cisco AnyConnect provides an always-on solution that finds the nearest enforcement element and enforces multiple types of policies. This includes the use of SSL and IPsec VPNs (IKEv2 only), 802.1 X authentications and authorization, Cisco Network Admission Control (NAC), and other posture and enforcement elements. Using this consolidated client, the network implements data loss prevention, threat prevention, acceptable use policies, and access control mechanisms, all with one lightweight client that eliminates the need for multiple software elements enforcing individual security controls.

Cisco Identity Services Engine

Cisco Identity Services Engine (ISE) is the centralized policy engine for business-relevant policy definition and enforcement. Cisco ISE complements global contextual information that is offered by Cisco SIO, with localized context awareness provided by Cisco AnyConnect and the scanning and enforcement elements for effective access policy enforcement.

Cisco ISE combines the functionality of other Cisco products, such as the Cisco Secure Access Control Server for authentication, authorization, and accounting (AAA) services and Cisco NAC, into this next-generation policy server.

Threat Control and Containment

The Cisco threat control and containment solution offers comprehensive protection for your network through network-wide visibility, simplified policy control, and proactive system protection. The Cisco threat control and containment solution regulates network access, isolates infected systems, prevents intrusions, and protects critical business assets. This solution counteracts malicious traffic such as worms, viruses, and malware before it affects your business, by using centralized policy, configuration, and threat event management.

Threat prevention products have been a flagship of Cisco security for a long time. Integrated security controls found in Cisco Adaptive Security Appliance (ASA), Cisco Integrated Services Routers (ISR), and IPS sensors provide a multidimensional approach to detecting, mitigating, and responding to threats at various levels.

- **Cisco ASA:** Cisco ASAs provide proven firewall services and context- and application-aware capabilities for comprehensive, real-time threat defense. ASAs implement a truly unified threat management service, with the integration of VPN and intrusion prevention technologies in various forms.
- **Cisco ISRs:** Through software- and hardware-integrated security functions, ISRs can easily become part of the army of security controls in networks of all kinds. These functions include zone-based policy firewall, Cisco IOS IPS, and the integration of SSL and IPsec VPN technologies.
- **Cisco IPS:** Intrusion prevention is accomplished in a distributed fashion, from IPS 4200 appliances to integrated hardware modules such as the Advanced Inspection and Prevention Security Services Module (AIP-SSM) for ASA or the Intrusion Detection Services Module (IDSM) for Cisco Catalyst 6500. These IPS sensors support a variety of IPS technologies, including signature-based, anomaly-based, policy-based, and reputation-based techniques.

These threat prevention mechanisms are integrated into the fabric of the network. Flexible deployment options include standalone appliances, virtualized devices, and service modules that are embedded into routers and switches.

Note

Cisco TechWiseTV has produced excellent videos on different security products. Check them out:

“Fundamentals of Intrusion Prevention”: <http://youtu.be/w-z2kS9dlcI>

“Fundamentals of ISE”: <http://youtu.be/sel1F7mKdtI>

“Fundamentals of TrustSec”: <http://youtu.be/78-GV7Pz18I>

Many other informative videos from TechWiseTV can be found at <http://www.youtube.com/user/techwisetv?ob=0>.

Cisco Security Intelligence Operation

Cisco SIO is the back-end security ecosystem that detects threat activity, researches and analyzes the threats, and provides real-time updates and best practices to keep organizations informed and protected. Cisco SIO, considered the largest threat analysis system providing blended threat protection, consists of three pillars:

- Threat intelligence, which is called Cisco SensorBase
- The automatic and human development process, called the IronPort Threat Operations Center
- The automated and best practices content that is pushed to network elements in the form of dynamic updates

Cisco SIO is a security intelligence center that baselines the current state of threats on a worldwide basis, and provides the network as a system with valuable information to detect, prevent, and react to threats. SIO acts as an early warning system by correlating threat information from the SensorBase, analyzed by the Threat Operations Center. SIO then feeds this information to enforcement elements, for live threat prevention based on malware outbreaks, current vulnerabilities, and zero-day attacks.

Cisco SIO weighs and processes the data, automatically categorizing threats and creating rules using more than 200 parameters. Security researchers also collect and supply information about security events that have the potential for widespread impact on networks, applications, and devices.

Rules are dynamically delivered to deploy Cisco security devices every three to five minutes. The Cisco SIO team also publishes security best practice recommendations and tactical guidance for thwarting threats.

Cisco's products participating and benefiting from Cisco SIO are IPS, ASA, IronPort Email Security Appliance (ESA), and IronPort Web Security Appliance (WSA). The AnyConnect telemetry module can even be used to send endpoint web malware traffic information back to the web filtering infrastructure of IronPort WSA.

Some interesting statistics on Cisco SIO at the time of printing:

- It is estimated that more than 700,000 sensors participate in Cisco SIO.
- Cisco SIO analyzes more than 5 billion web requests daily coming from Cisco equipment such as IronPort WSA.
- It is estimated that 35 percent of all email traffic worldwide is checked against Cisco SIO.

Cisco SIO is a powerful way of analyzing traffic based on reputation, blocking malware and spam, and providing a granular web categorization and application classification.

Cloud Security, Content Security, and Data Loss Prevention

Cloud computing represents one of the most significant shifts in IT that many of you are likely to see in your lifetime. Reaching the point where computing functions as a utility has great potential, promising innovations that cannot yet be imagined. However, this breakthrough business model also imposes new security risks.

Not in any order of importance or severity, following is a list of the top threats to cloud computing, according to the Cloud Security Alliance (<https://cloudsecurityalliance.org/>):

- Abuse and nefarious use of cloud computing
- Insecure interfaces and APIs
- Malicious insiders
- Shared technology issues
- Data loss or leakage
- Account or service hijacking
- Unknown risk profile

Vendors, providers, and customers alike are concerned about the consequences if cloud computing is not properly secured, and the loss of direct control over systems for which they are nonetheless accountable.

Cloud computing is a double-edged sword: it introduces more security concerns but it provides security opportunities. Regarding security concerns about cloud computing, consider, for example, an organization that has embraced Software as a Service (SaaS) and uses cloud-based services such as Google Apps for Business and Salesforce. In a traditional cloud service, if a user leaves the organization, the administrator would need to visit both the Salesforce website and the Google Apps for Business website and proceed to nullify that user's accounts at each of those sites. However, if the organization has implemented the Cisco SaaS Access Control solution offered on Cisco IronPort WSA (discussed later), the administrator would only need to cancel the user's account from the WSA.

Content Security

Before we jump in to cloud computing, let's review some of Cisco's offerings for content security.

Content security refers to the threat prevention capabilities that provide malware protection against spyware, viruses, spam, and inappropriate content. Examples of content security products from Cisco include the following:

- Cisco IronPort Web Security Appliance (WSA), which scans the content of a web page for malware prior to passing it to the inside user who initiated the request
- Cisco IronPort Email Security Appliance (ESA), which scans incoming email against spam
- Cisco ASA 5500 Series Content Security and Control Security Services Module (CSC-SSM) running Trend Micro's antivirus and anti-spyware technologies.

Cisco IronPort WSA and ESA are discussed further later in this section.

Data Loss Prevention

Also known as data leak prevention systems or extrusion prevention systems, data loss prevention (DLP) systems monitor outbound traffic to ensure that confidential information doesn't fall into the wrong hands.

Cisco IronPort WSA and ESA both offer solutions to scan outbound traffic against preconfigured and customizable rules, thus preventing corporate or personal information from leaving the organization's premises.

Cloud-Based Security

Traditional threat prevention is now augmented to extend the system to email and web security, two key elements of the new mobile, cloud-based business environment.

Cisco offers both on-premises and cloud-based security for web and email traffic. For web security, Cisco offers the IronPort WSA as an on-premises device and IronPort ScanSafe as a cloud-based solution. For email security, Cisco offers the IronPort ESA either as an on-premises device or as a cloud-based service.

Web Security

Cisco offers both on-premises and cloud-based web security solutions, ScanSafe and the Cisco IronPort WSA, respectively. The WSA is what we commonly refer to as a web-proxy server.

Web malware is a common vehicle used to materialize threats in cloud services. Instead of creating their own malicious websites, hackers exploit the vulnerabilities of an open and dynamic web to distribute their malware. Web malware infection from reputable websites that have been compromised is not only a reality, but is now hackers' preferred route to infect victims.

This change has made traditional methods of control, such as antivirus software, less effective, and it requires an alternative approach to security.

Web security architecture from Cisco presents a flexible approach and deployment options for implementing services such as URL filtering, web content filtering, real-time scanning, web traffic and application control, malware prevention, and others. ScanSafe is a cloud-based solution that enhances security while enabling cost savings by eliminating the need to purchase, deploy, and maintain hardware required for on-premises solutions. IronPort WSA is an on-premises solution that offers data loss prevention capabilities, URL filtering, and other web security controls.

ScanSafe Web Security analyzes every web request to determine if content is malicious, inappropriate, or acceptable based on the defined security policy. This offers effective protection against threats, including zero-day threats, that would otherwise be successful.

ScanSafe Web Security is powered by Outbreak Intelligence, which is composed of numerous correlated detection technologies, automated machine-learning heuristics, and multiple "scanlets." Outbreak Intelligence builds a detailed view of each web request and the associated security risk to ensure that ScanSafe customers use the web safely.

[Figure 2-10](#) illustrates the integration of ScanSafe with the Cisco AnyConnect client, to provide seamless deployment of context-aware policies, protected at the application layer against web malware.

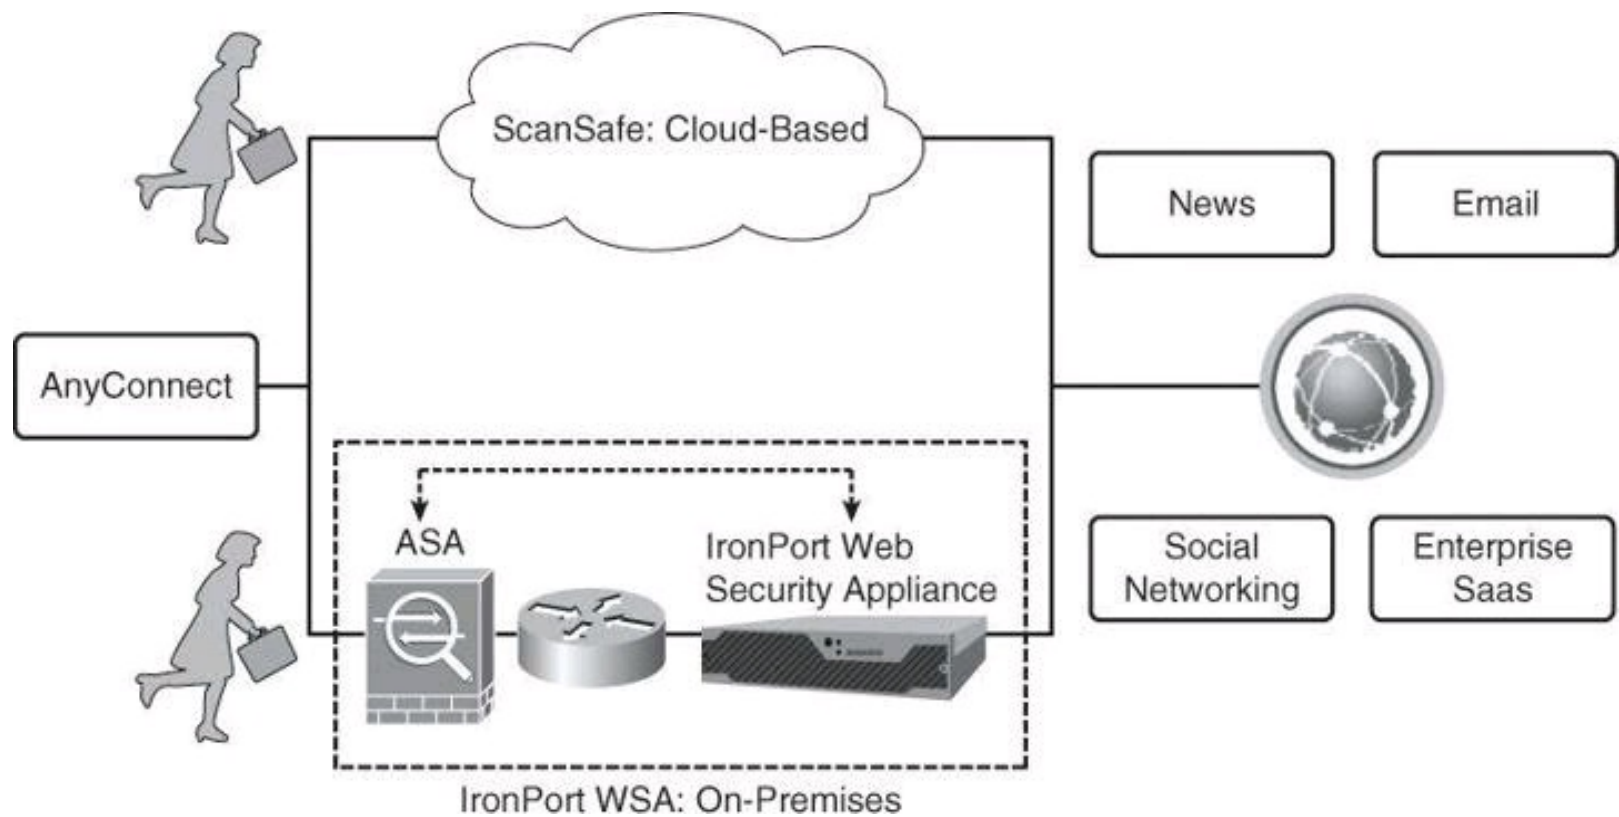


Figure 2-10. ScanSafe Integration with Cisco AnyConnect

This is the highlighted functionality:

- Integrated management and reporting that covers all aspects of the solution
- Consistent policy and security for all users, regardless of location, whether they connect from their cubicle or from the boardroom in an adjacent corporate building, as an example
- Numerous ways to integrate with existing network infrastructure and authentication services
- Bidirectional content-based policy enforcement
- Dynamic content classification
- Control over HTTP and HTTPS communications
- Accurate zero-day threat protection
- All security extended to remote and roaming users in addition to on-premises users

[Figure 2-11](#) also illustrates the use of IronPort Web Security Appliances on premises. In combination with ASAs, this solution implements acceptable use policies, malware prevention, and single sign-on to SaaS providers.




Figure 2-11. Cisco IronPort WSA on Premises

Cisco AnyConnect users use VPN technologies to connect to the corporate network via ASA, and their identity is then shared with WSAs for an integrated, context-aware web security solution.

Email Security

Cisco offers both on-premises and hosted email security solutions, using Cisco IronPort Email Security Appliance, as shown in [Figure 2-12](#).

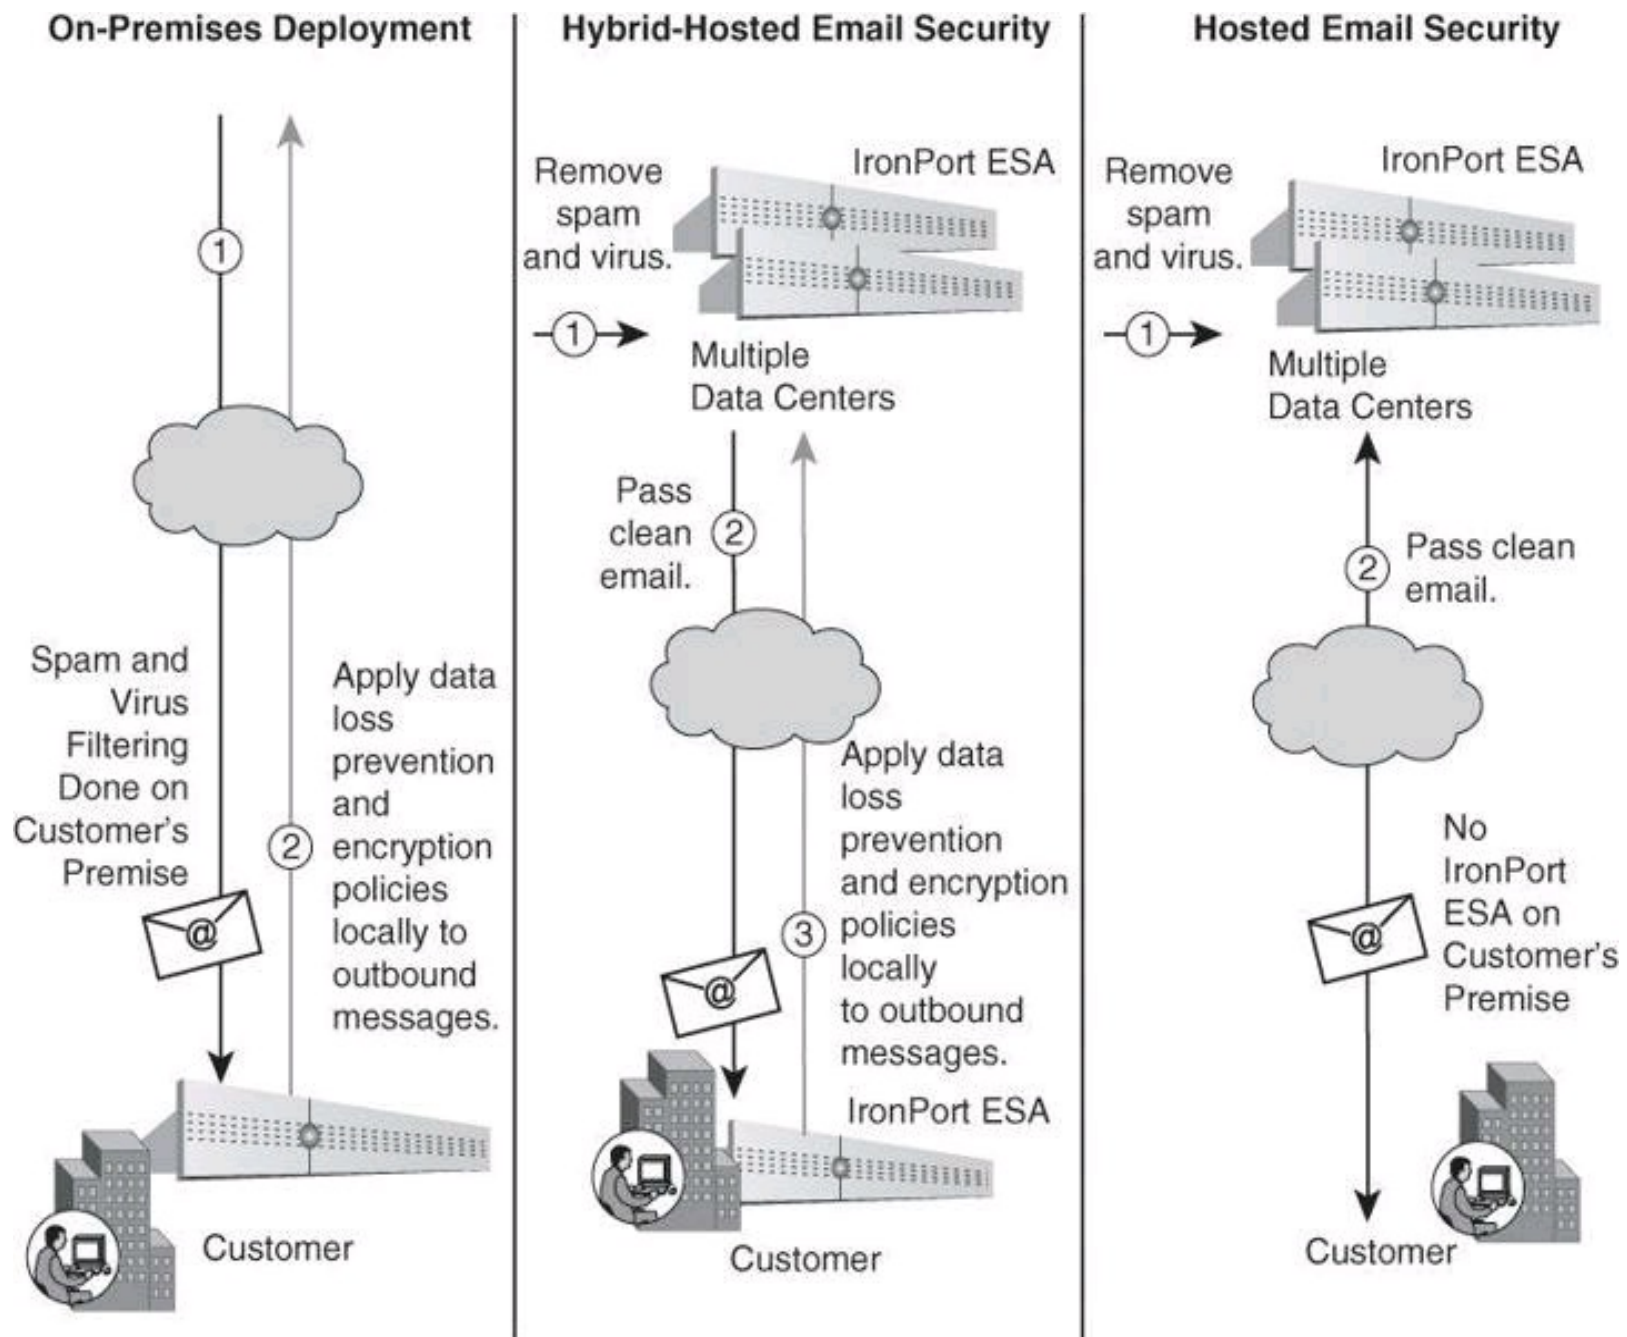


Figure 2-12. Cisco IronPort Email Security Solutions

Cisco IronPort Email Security solutions provide outstanding protection to organizations of all sizes. Sophisticated and scalable mechanisms help to do the following:

- Minimize the downtime that is associated with email-borne malware
- Simplify the administration of corporate mail systems
- Reduce the burden on technical staff, while offering insight into mail system operation

Best-in-class technologies work together to prevent and respond to multilevel threats. They include the following:

- **Spam protection:** The Cisco spam protection combines best-of-breed conventional techniques with IronPort's breakthrough context-sensitive detection technology to eliminate the broadest range of known and emerging email threats.
- **Data loss prevention:** Cisco has partnered with RSA, a leader in data loss prevention technology, to provide an integrated data loss protection solution, RSA Email DLP, on

Cisco IronPort ESAs. This solution ensures compliance with industry and government regulations worldwide and helps prevent confidential data from leaving customer networks.

- **Virus defense:** Cisco IronPort Virus Outbreak Filters detect new virus outbreaks in real time, then quarantine suspicious messages, offering protection before traditional antivirus solutions.
- **Email encryption tracking and reporting tools:** Cisco IronPort Preboot Execution Environment (PXE) encryption technology revolutionizes email encryption, meeting compliance requirements while delivering powerful, business-class email features.

IronPort solutions offer a choice of features and functionality available as appliance-based, cloud-based, hybrid, or managed solutions (where the solution is on premises but managed by Cisco's expert staff).

Secure Connectivity Through VPNs

Ensuring the privacy and integrity of all information is vital to your business. You can achieve privacy and integrity by using IPsec and SSL VPNs, as shown in [Figure 2-13](#). As your company uses the flexibility and cost effectiveness of the Internet to extend its network to branch offices, telecommuters, customers, and partners, security is paramount. You must create a manageable, cost-effective communications infrastructure that will do the following:

- Improve productivity
- Enable new business applications
- Help you comply with information privacy regulations
- Enhance business efficiency

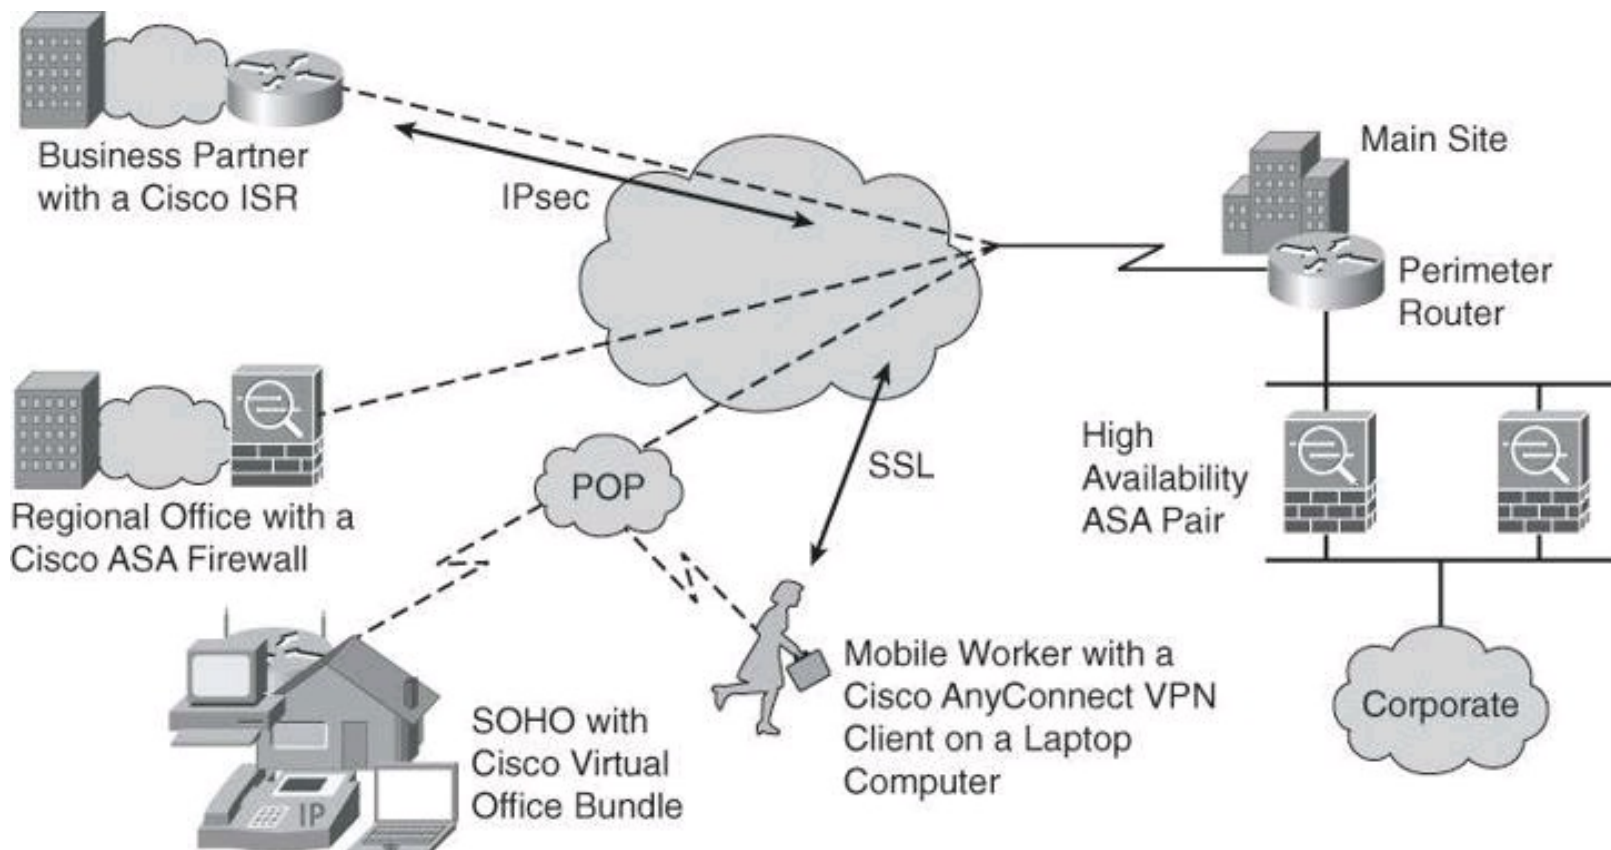


Figure 2-13. VPN Deployment Options

Cisco VPN solutions implement a set of products and security life-cycle services that are an essential element of the Borderless Network Architecture. By incorporating capabilities that secure the network, the endpoints, and the applications and messages, this systems-based approach delivers comprehensive security of your communications. The solution has two major elements:

- **Secure communications for remote access:** This element provides highly secure, customizable access to corporate networks and applications by establishing an encrypted tunnel across the Internet.
- **Secure communications for site-to-site connections:** This element provides an Internet-based WAN infrastructure for connecting branch offices, home offices, or the sites of business partners to all or portions of your network.

Security Management

Cisco network management systems help you automate, simplify, and integrate your network to reduce operational costs and improve productivity. The tools within the network management systems provide innovative ways to centrally and consistently manage your network to achieve critical functions such as availability, responsiveness, resilience, and security.

These network management systems also help reduce the troubleshooting and planning time that is associated with the introduction of new services such as voice, wireless, and security management. Solutions-focused tools simplify network management systems, including the management of devices, configurations, users, and services.

The tools available for security management are

- **Device managers:** Embedded into security devices in the form of HTML content. All you have to do is point your browser to an active IP address of the device, and the device manager immediately provides a comprehensive tool for one-to-one configuration and monitoring. ASDM, the Adaptive Security Device Manager, is an example of this approach for the ASA 5500 security appliances.
- **Cisco Configuration Professional:** A GUI-based device management tool for Cisco access routers. This tool simplifies routing, firewall, IPS, VPN, unified communications, WAN, and LAN configuration through GUI-based wizards. Cisco Configuration Professional is a valuable productivity enhancing tool for network administrators and channel partners for deploying routers with increased confidence and ease. This tool offers a one-click router lock-down and an innovative voice and security auditing capability to check and recommend changes to router configuration. This tool also monitors router status and troubleshoots WAN and VPN connectivity issues.
- **Cisco Security Manager:** Provides configuration and monitoring capabilities for firewalls, routers, switches, IPS sensors, and other security solutions. This tool uses multiple approaches to security management, including device-focused and policy-focused options.

Cisco Security Manager

Cisco Security Manager is a powerful but easy-to-use solution that enables you to centrally

provision all aspects of device configurations and security policies for the Cisco family of security products. The solution is effective for managing even small networks consisting of fewer than ten devices, but also scales to efficiently manage large-scale networks that are composed of thousands of devices. Scalability is achieved through intelligent policy-based management techniques that can simplify administration and promote policy uniformity among devices. As a simple example, the administration might come up with a new login banner that is pushed to Cisco devices, such as routers, firewalls, and switches, using CSM. If in the future the legal department changes the wording of the login banner, CSM could deploy the new banner in a matter of minutes to all the devices.

Some of the features of Cisco Security Manager include the following:

- Supports provisioning for Cisco router platforms running a Cisco IOS Security Software image, Cisco ASA 5500 Series Adaptive Security Appliances, Cisco PIX 500 Series Security Appliances, Cisco IPS 4200 Series Sensors, and Cisco Catalyst 6500 Series Advanced Inspection and Prevention Security Services Module (AIP-SSM)
- Responds faster to threats by allowing you to define and assign new security policies to thousands of devices in a few simple steps
- Provides a rich GUI for superior ease of use
- Offers multiple views that provide flexible methods to manage devices and policies, including the ability to manage the security network visually on a topology map
- Contains extensive animated help for the new user, which reduces the learning time
- Allows you to centrally specify which policies are shared and automatically inherited by new devices to ensure that corporate policies are implemented consistently, while providing optional flexibility
- Integrates with Cisco Secure Access Control Server (ACS) to provide granular role-based access control to devices and management functions, though this functionality is removed once you upgrade to CS ACS 5.2
- Provides the ability to assign specific tasks to each administrator during the deployment of a policy, with formal change control and tracking, and allows the security and network operations staff to work together as a single team with effective coordination

Summary

In this chapter, you learned about the Cisco Borderless Network Architecture. This chapter examined the Cisco Security portfolio of products and, more specifically, reviewed the following:

- Cisco SecureX Architecture (at a high level), highlighting its features and benefits and providing examples of Cisco products that fall within this category
- Cisco threat control and containment products and technologies, such as the Cisco ASA and Cisco IPS, and illustrating their high-level features and benefits
- Cisco content security and data loss prevention technologies, such as Cisco IronPort WSA and ESA, and illustrating their high-level features and benefits
- Cisco VPN solutions and technologies, and illustrating their high-level features and benefits
- The different security management products and technologies, focusing at a high level on

References

For additional information, refer to these Cisco.com resources:

“Borderless Networks,” <http://www.cisco.com/en/US/netsol/ns1015/architecture.html>

Cisco AnyConnect Secure Mobility Client Administrator Guide, Release 3.0, Chapter 7,

“Configuring AnyConnect Telemetry to the WSA,”
http://www.cisco.com/en/US/docs/security/vpn_client/anyconnect/anyconnect30/administra

“Cisco AnyConnect Secure Mobility Solution,”
http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5743/ps5699/ps10884/at_a_glanc
[578609.pdf](http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5743/ps5699/ps10884/at_a_glanc/578609.pdf)

“Cisco SecureX,” <http://www.cisco.com/en/US/netsol/ns1167/index.html>

“Cisco TrustSec,” <http://www.cisco.com/en/US/netsol/ns1051/index.html>

“Cisco Network Identity and Access Policy Solution,”
http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5712/ps2086/network_id_vds.htm

Cisco’s TechWiseTV, <http://www.youtube.com/ciscoCIN>

Securing the Borderless Network, <http://www.ciscopress.com/bookstore/product.asp?isbn=1587058863>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

1. Match each Cisco SecureX component with its definition.

a. Context-aware enforcement

b. Cisco TrustSec

c. Cisco SIO

d. Cisco AnyConnect

_____ 1. Extends the access control functionality end to end, using security group tags.

_____ 2. Used for real-time insight into the global threat environment.

_____ 3. Provides a consistent user interface and consolidates traditional function-specific client software products into one.

_____ 4. Allows enforcement elements to use identity and location information to define access policy.

2. The ASA, ISR, and IPS are all part of which product category?

a. Threat control and containment

b. Cisco TrustSec

c. Cisco SecureX

d. Cisco SIO

e. Content security

- 3.** Which two Cisco products provide data loss prevention in cloud services environments?
(Choose two.)
- a.** Cisco TrustSec
 - b.** Cisco ScanSafe
 - c.** Cisco SecureX
 - d.** Cisco IronPort ESA
 - e.** Cisco ASA
- 4.** Using Cisco TrustSec functionality, what needs to happen before security group tags are assigned?
- a.** MACsec confidentiality policy needs to be created
 - b.** Access policies are built in Cisco ISE
 - c.** Security group ACLs need to be defined
 - d.** User authentication mode needs to be selected
- 5.** The Cisco ISE policy engine combines the functionality of which two Cisco products?
(Choose two.)
- a.** Cisco AnyConnect
 - b.** Cisco Secure ACS
 - c.** Cisco Security Manager
 - d.** CiscoWorks
 - e.** Cisco NAC

Part II: Protecting the Network Infrastructure

Chapter 3. Network Foundation Protection and Cisco Configuration Professional

In this chapter, you learn about Cisco IOS Network Foundation Protection (NFP) as a framework for infrastructure protection, all its components, and commonly used countermeasures as found in Cisco IOS devices. You also learn how to use Cisco Configuration Professional (CCP) to implement security controls on Cisco IOS routers.

More precisely, you learn how to do the following:

- Categorize common threats against the network infrastructure
- Describe Network Foundation Protection as a framework to develop and implement security controls to protect the network infrastructure
- List and compare security controls that protect the control plane
- List and compare security controls that protect the data plane
- List and compare security controls that protect the management plane
- Articulate the features and benefits of CCP, describing its requirements and installation options
- Demonstrate the CCP GUI, showcasing the most relevant options and features
- Describe the unique components of CCP that are used for effective security policy deployment and configuration
- Describe and implement the One-Step Lockdown and audit features found on CCP

Threats Against the Network Infrastructure

A key element in the overall security posture of an organization is the security of the network infrastructure. The network infrastructure is the foundation that is built with routers, switches, and other equipment that provides the fundamental network services to keep a network running. The infrastructure is often the target of denial of service (DoS) and other attacks that can directly or indirectly disrupt the network operation. In order to ensure the availability of the network, it is critical to implement the security tools and best practices that help protect each network element and the infrastructure as a whole.

Cisco Network Foundation Protection (NFP) provides an umbrella strategy for infrastructure protection by encompassing Cisco IOS security features.

The current network environment is complex, as networking devices offer a feature-rich set of services to cater to a variety of business needs. Business continuity requires security features and services for both network devices and infrastructure, thus ensuring the availability of the network devices under all circumstances.

Because connecting to the Internet is imperative, network devices and infrastructure are exposed to many risks and threats. Deploying security best practices helps secure the network foundation by protecting network elements and their interactions. In Internet and cloud computing environments, this means distrusting every packet and implementing policies that are aimed at categorizing traffic.

Different traffic classes are then processed under different policies to ensure not only proper forwarding, but also dropping suspicious traffic, thus protecting the network devices themselves.

[Table 3-1](#) lists some of the common threats to the network infrastructure. These threats are used not only to disrupt service via DoS attacks and routing protocol exploits, but also to create confidentiality exposure via trust exploitation attacks.

Table 3-1. Common Issues for Network Infrastructure

Vulnerabilities	Threats	Impact
Design errors	Trust exploitation attacks	Exposed management credentials
Protocol weaknesses	Login, authentication, and password attacks	High route processor CPU utilization (near 100 percent)
Software vulnerabilities	Routing protocol exploits	Loss of protocol updates keepalives and routing
Misconfiguration	Spoofing	Route flaps and major network transitions
Multiple categories of vulnerabilities	Denial of service attacks	Slow or unresponsive management sessions
Multiple categories of vulnerabilities	Confidentiality and integrity attacks	Indiscriminate packet drops

These threats exploit known vulnerabilities such as protocol weaknesses. For instance, the absence of confidentiality mechanisms in Telnet and certain versions of Simple Network Management Protocol (SNMP) can expose passwords and community names to malicious attackers. Unknown OS/software vulnerabilities can also be exploited using zero-day attacks. This exploitation can occur with software bugs and weaknesses, sometimes found in the operating systems of network devices.

The impact of these exploits and incidents is at the core of business operations. High resource utilization slows down network devices, and can even bring them down. Routing protocol operations can be disrupted, causing connectivity problems for all kinds of traffic. The problem is augmented when management sessions are slowed down in such a way that fixing the problem becomes a challenge.

The following are some of the symptoms and impact of network infrastructure security incidents:

- High route processor CPU utilization (near 100 percent)
- Loss of line protocol keepalives and routing protocol updates, leading to route flaps and major network transitions
- Slow or completely unresponsive interactive sessions via the CLI, due to high CPU utilization
- Route processor resource exhaustion, making resources such as memory and buffers unavailable for legitimate IP data packets

- Packet queue backup, leading to indiscriminate drops (or drops due to lack of buffer resources) of other incoming packets

It is important to apply a structured approach to network infrastructure protection. Different threats exploit different vulnerabilities in different functional areas of the network. As an example, spoofing attacks are commonly used to power DoS attacks from the external network, to exhaust edge device resources and link bandwidth. On more internal security domains, such as server farms and the campus access layer, Layer 2 attacks are typically aimed at trust exploitation and information theft using man-in-the-middle mechanisms.

Deploying the right policy and feature to the appropriate device is critical to the effective mitigation of these threats.

Cisco NFP Framework

To understand Cisco NFP architecture, it is important to understand the basics of the architecture of devices with routing capabilities. A router can be logically divided into three functional components or planes:

- Data plane
- Management plane
- Control plane

Most traffic travels through the router via the data plane. However, the route processor must manage certain packets, such as routing updates, keepalives, and network management. This traffic is often referred to as control and management plane traffic. The management plane, however, is considered separate from a functional perspective, in order to define different protection mechanisms to management traffic. In other words, the control and data planes use the central resources (CPU, memory, and so on) in Cisco network devices more intensely than does the management plane. This is depicted in [Figure 3-1](#) by grouping them inside the same segmented lines. The data plane is typically processed in a fast-switching cache using distributed processing or in high-priority, process-independent code in the kernel.

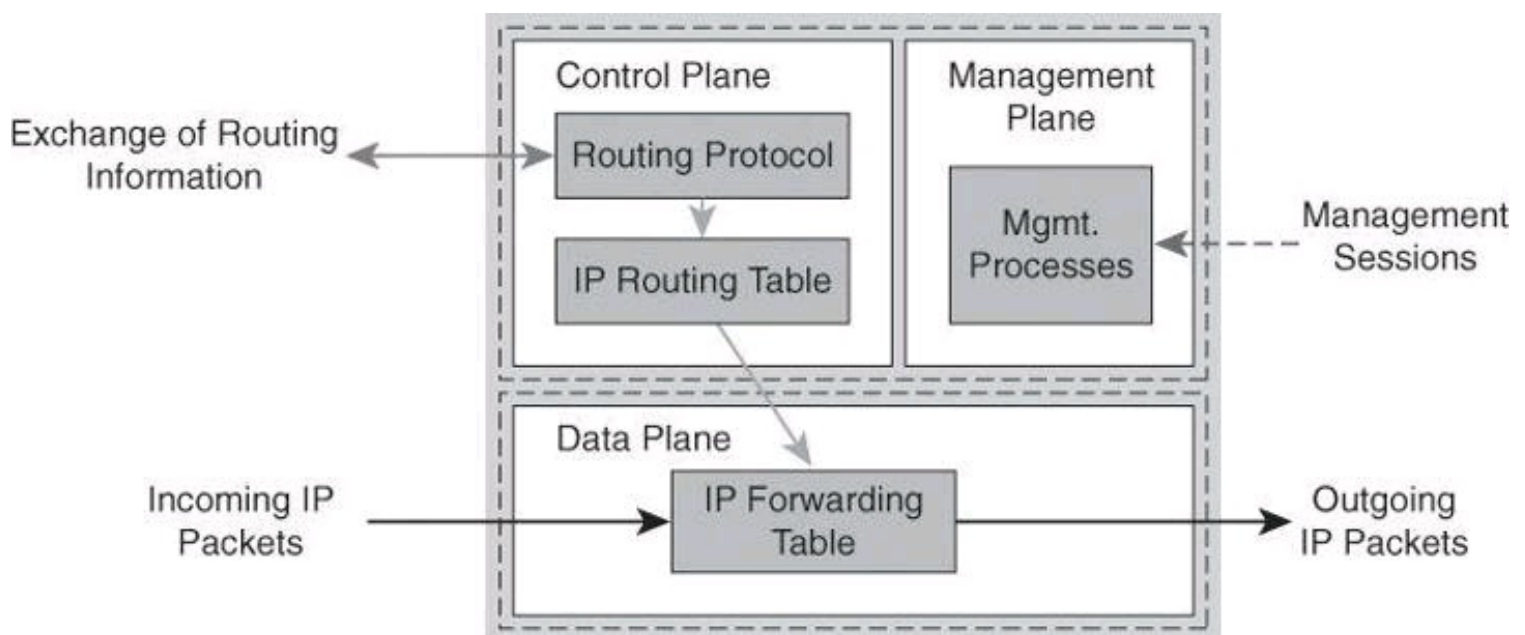


Figure 3-1. Device Planes

From the perspective of network traffic, a similar classification applies:

- **Data plane packets:** End-station, user-generated packets that are always forwarded by network devices to other end-station devices. From the perspective of the network device, data plane packets always have a transit destination IP address and can be managed by normal, destination IP address–based forwarding processes. Data plane packets are typically processed in a fast-switching cache.
- **Control plane packets:** Network device–generated/received packets that are used for the creation and operation of the network itself. From the perspective of the network device, control plane packets always have a receive destination IP address and are managed by the CPU in the network device route processor. Examples include protocols such as Address Resolution Protocol (ARP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and other protocols that hold the network together.
- **Management plane packets:** Network device–generated/received packets, or management station–generated/received packets, that are used to manage the network. From the perspective of the network device, management plane packets always have a receive destination IP address and are managed by the CPU in the network device route processor. Examples include protocols such as Telnet, Secure Shell (SSH), TFTP, SNMP, FTP, Network Time Protocol (NTP), and other protocols that are used to manage the device or network.

Differentiating Between Control Plane, Data Plane, and Management Plane

The control path code, such as routing protocols, runs on the **control plane**.

Data packets, such as user traffic, are forwarded by the **data plane**.

Management sessions to the device, such as Secure Shell, use the **management plane**.

To deploy a protection policy to protect the different operation planes of a router, a structured approach is needed. The three distinct functions will benefit from a combined effort to protect the device, but each one affects the router operations differently. The control plane affects the ability to route packets, because it implements routing protocol functions. The management plane affects the ability to configure and monitor the device, and eventually implement the desired security controls. The data plane affects the ability to forward traffic, a critical function that must be performed efficiently.

Packet overloads on the control plane of a router can slow down routing processes and, as a result, degrade network service levels and user productivity. Packets that traverse the control plane are those destined for the CPU of that router, as opposed to network endpoints. All packets entering the control plane are redirected by the forwarding plane.

Cisco NFP is an umbrella strategy encompassing Cisco IOS security features that provides the tools, technologies, and services that enable organizations to secure their network foundations. Cisco NFP helps to establish a methodical approach to protecting router planes, forming the foundation for continuous service delivery. NFP is not a single feature or technology, nor is it a single management tool or set of commands. It is a series of Cisco IOS features designed specifically to protect the device control plane by “locking down” services and routing protocols, protect the device data plane

from malicious traffic, and protect the device management plane.

Cisco NFP is not a single feature of technology but rather is the deployment of a protection strategy that is based on several mechanisms that affect all three planes. [Table 3-2](#) summarizes some of the protection techniques available under the Cisco NFP framework. This list is only a sample of the available tools, some of which will be described and demonstrated in this chapter. More information on Cisco NFP can be found at <http://www.cisco.com/go/nfp>. The NFP features listed in [Table 3-2](#) will be explained in the upcoming chapters of this book.

Table 3-2. Some Components of Cisco NFP

Plane	Feature	Benefit
Control plane	Control Plane Policing (CoPP)	Filter or rate limit control plane traffic with no regard to physical interface
	Control Plane Protection (CPPr)	Extend CoPP with granular traffic classification
	Routing protocol authentication	Integrity of routing and forwarding
	Cisco AutoSecure	Automate device hardening
Management plane	NTP, syslog, SNMP, SSH	Secure management and reporting
	CLI views	Obtain the benefits of role-based access control (RBAC) for command line
	Authentication, authorization, and accounting (AAA)	A comprehensive framework for RBAC
Data plane	Access control lists	Traffic filtering consistent across security device platforms
	Layer 2 controls (private VLANs, STP guards, others)	Protect the switching infrastructure
	Zone-based firewall, Cisco IOS IPS	Deployment flexibility in Cisco IOS form factor

Your network benefits from applying a structured approach to infrastructure protection using Cisco NFP. [Figure 3-2](#) shows some of the benefits, which are explained next.

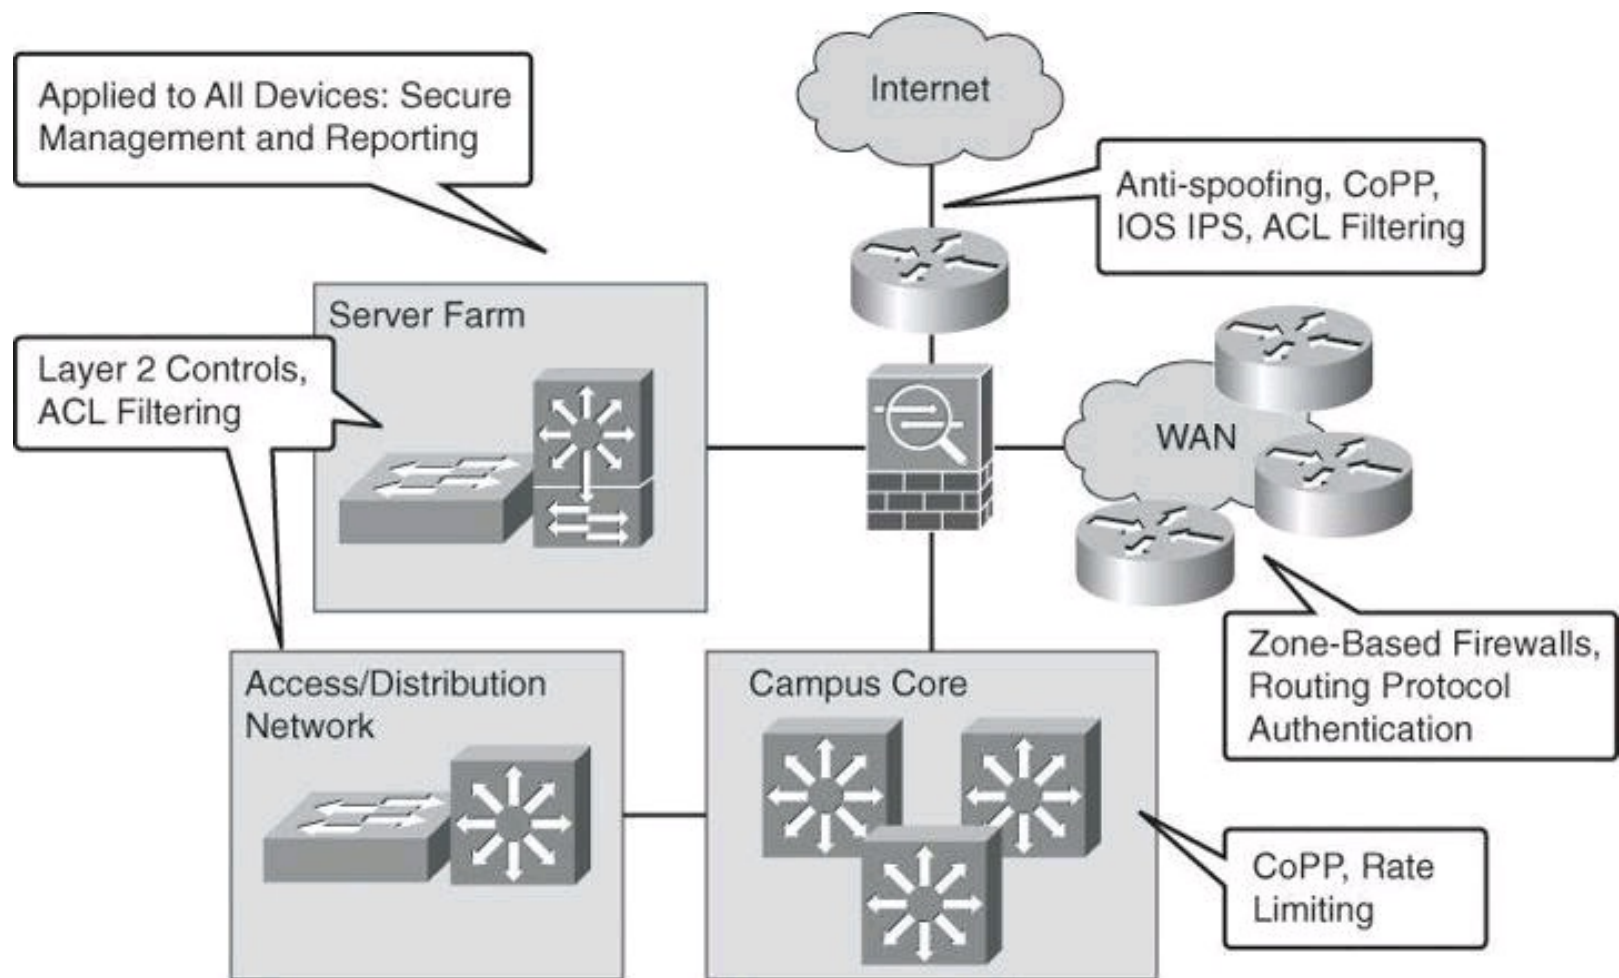


Figure 3-2. Some of Cisco NFP in a Network

Overall, network elements benefit from a strategy that is based on role-based access control (RBAC) at the management level. Role-based access control is a strategy that associates the authority to execute certain tasks with a specific privilege level, such as certain tasks that could be executed by senior network administrators but not by junior network administrators. A simple example would be that senior administrators are allowed to perform the **show running-config** command, but junior administrators are not allowed. Network devices are organizational assets like any other, and their exploitation can cause network outages, which could be prevented using proper RBAC deployment.

Secure management and reporting mechanisms, such as SSH and SNMPv3, could be planned across the board for all network devices for device hardening and management plane protection.

On the other hand, traditional perimeters such as the Internet edge benefit from antispoofing mechanisms such as infrastructure access control list (ACL) filtering, as well as external threat protection using Cisco IOS Intrusion Prevention System (IPS). These NFP countermeasures protect not only the device itself but also the assets that those devices connect to in the organization.

Internal perimeters, such as the WAN edge, benefit from a distributed approach to security countermeasures, in the form of zone-based firewalls. These firewalls benefit from the Cisco IOS form factor to leverage existing network elements, implementing access control and data plane security in multiple edge points of the WAN in a distributed fashion.

More internal perimeters, such as server farms and the access and distribution layers, benefit from Layer 2 controls that protect the switching infrastructure.

Cisco NFP does not prevent poor network design. In that sense, building some level of redundancy into the network design is critical to the deployment of some of the NFP features. There is no substitute to designing a layered approach to security, with redundant countermeasures at different levels. Physical security and high availability are also design building blocks that need to be in place for an NFP strategy to be effective.

It is critical to apply the proper selection of Cisco NFP features to the proper device, according to a risk management plan. In that sense, Cisco NFP also assumes that a security policy is in place.

In terms of operations, the network design should follow recommended practices in change management and disaster recovery procedures.

Control Plane Security

Control Plane Policing (CoPP) is a Cisco IOS feature designed to allow users to manage the flow of traffic that is managed by the route processor of their network devices. Ideally, traffic is fast-switched from one interface to the other, as represented by the arrow in [Figure 3-3](#). However, sometimes traffic needs to be process-switched. The CPU is shared among three main functions: routing protocols, slow data plane, and management process. Excessive traffic to any of the three main functions can overwhelm the CPU and affect the other two functions. CoPP is designed to prevent unnecessary traffic from overwhelming the route processor and the slow data path.

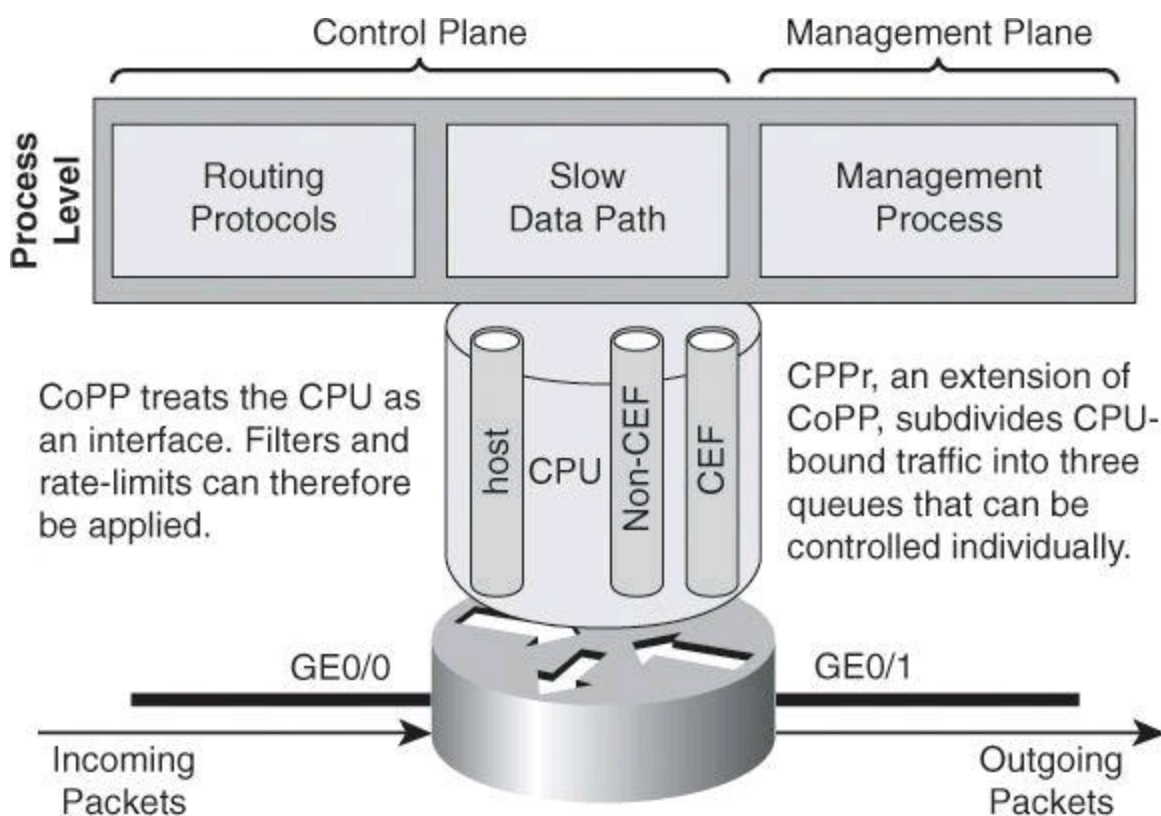


Figure 3-3. Goal of CoPP: Treat the CPU as an Interface

To protect the control plane on a router from DoS attacks, and to provide packet quality of service (QoS), the CoPP feature treats the control plane as an interface, as shown in [Figure 3-3](#). Because the CoPP feature treats the control plane as a separate entity, a set of rules can be established and associated with the ingress and egress functions of the control plane, rules such as ACLs or rate limits.

These rules are applied only after the packet has been determined to have the control plane as its

destination or when a packet exits the control plane. Thereafter, you can configure a service policy to prevent unwanted packets from progressing after a specified rate limit has been reached. For example, a system administrator can limit all TCP synchronization (SYN) packets that are destined for the control plane to a maximum rate of 1 Mbps.

Control plane security consists of implementing the following features:

- **CoPP:** CoPP lets users configure a QoS filter that manages the traffic flow of control plane packets to protect the control plane of Cisco IOS routers and switches against reconnaissance and DoS attacks.
- **CPPr:** CPPr is an extension of the policing functionality that the existing CoPP feature provides. CoPP allows QoS policing of aggregate control plane traffic that is destined to the route processor. The CPPr feature extends this policing functionality by allowing finer policing granularity by considering the CPU as three queues (subinterfaces), as shown in [Figure 3-3](#): the host subinterface used by the routing and management processes, the Cisco Express Forwarding (CEF) subinterface, and the non-CEF subinterface.
- **Control Plane Logging:** CPPr lets you filter and rate-limit the packets that are going to the control plane of the router and discard malicious packets, or malformed packets or packets containing errors. The Control Plane Logging feature enables logging of the packets that these features drop or permit. The Control Plane Logging feature provides the logging mechanism that you need to deploy, monitor, and troubleshoot CoPP features efficiently.

CoPP

The CoPP feature provides the following benefits:

- Protection against DoS attacks that are targeted toward the network infrastructure by traffic flows and protocols that must be permitted, but where rate limiting offers substantial protection.
- Easier deployment. CoPP uses the existing Modular QoS CLI (MQC) infrastructure, which allows customers to preserve the existing interface configurations and add global control plane–specific commands to address security goals.
- A consistent implementation strategy across all Cisco hardware.
- Increased reliability, security, and availability of the network.

CPPr

Cisco CPPr extends the CoPP feature by enabling classification of the control plane traffic based on packet destination and information that is provided by the forwarding plane, allowing appropriate throttling for each category of packet. CPPr provides a finer policing. This is where you might add granularity to your policy, to accomplish tasks such as differentiating between Telnet, SSH, and HTTPS management traffic.

The functionality that is added with CPPr includes a traffic classifier, which intercepts traffic and classifies it into three control plane categories (queues, as shown in [Figure 3-3](#)). New port-filtering and queue-thresholding features have also been added. The port-filtering feature provides for policing of packets going to closed or nonlisted TCP and UDP ports, while queue thresholding limits the number of packets for a specified protocol that will be allowed in the control plane IP input

queue.

As with CoPP, CPPr is implemented using Cisco IOS MQC, a highly flexible framework that allows users to create and attach traffic policies to interfaces. The Cisco MQC mechanisms are used by CPPr to define the classification and policing descriptions for its policies. In this way, in addition to the limited permit and deny actions that are associated with simple ACLs, specific packets may be permitted but rate-limited when using the MQC structure. For example, you may wish to permit certain Internet Control Message Protocol (ICMP) packet types, but rate limit them so that the route processor is not adversely impacted. This action adds tremendously to the capabilities and flexibility of developing and deploying a useable CPPr policy.

Key Concepts

What does Control Plane Protection provide you?

- Extends CoPP protection by providing mechanism for finer policing granularity for control plane traffic
 - Allows rate limits for each traffic class individually
 - Provides ease of configuration for control plane policies using modular policy CLI infrastructure
 - Provides better platform reliability, security, and availability
 - Provides CPU protection so it can be used for important jobs, such as routing
-

Traffic Classes

Before developing an actual CPPr policy, required traffic must be identified and separated into different classes. Multiple classification schemes can be used, but one recommended methodology involves classifying traffic into distinct groups that are based on relative importance. This approach can be illustrated with a simple classification example composed of four basic classes: critical, normal, undesirable, and default.

[Table 3-3](#) illustrates a more complex classification scheme, along with the action to apply to each traffic class. This example uses seven different classes, which provides greater granularity and is more suitable for real-world environments. The actual number and type of classes that are needed for a given network may differ and should be selected based on local requirements, security policies, and a thorough analysis of baseline traffic.

Table 3-3. Example of a Complex CPPr Classification Scheme

Traffic Class Rate	Rate (pps)	Conform Action	Exceed Action
Routing Protocols (OSPF, and so on)	N/A	Transmit	Transmit
Management (HTTPS, SSH, and so on)	250	Transmit	Drop
Reporting (syslog and so on)	50	Transmit	Drop
Monitoring (SNMP and so on)	75	Transmit	Drop
Spanning Tree	25	Transmit	Drop
Undesirable	10	Drop	Drop
Default (anything else)	10	Transmit	Drop

The goal in this book is to introduce the complex topics of CoPP and CPPr. A deeper discussion of these topics can be found in the Cisco Press book *CCNP Security SECURE 642-637 Official Cert Guide*.

Routing Protocol Integrity

Other control plane security mechanisms are specific to routing protocols, such as routing protocol neighbor authentication. When configured, neighbor authentication occurs whenever neighbor routers exchange routing updates. This authentication ensures that a router receives reliable routing information from a trusted source.

Without neighbor authentication, unauthorized or deliberately malicious routing updates could compromise the security of your network traffic. A security compromise could occur if an unfriendly party diverts (rogue default gateways) or analyzes that traffic (man-in-the-middle attacks). For example, an unauthorized router could send a fictitious routing update to convince your router to send traffic to an incorrect destination.

The unfriendly party could analyze the diverted traffic to learn confidential information about your organization or merely use it to disrupt the ability of your organization to communicate effectively using the network. Neighbor authentication prevents your router from receiving any such fraudulent routing updates.

These routing protocols support neighbor authentication for IPv4: BGP, OSPF, Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (ISIS), and Routing Information Protocol version 2 (RIPv2). For IPv6, the list is complete with OSPF version 3 (OSPFv3) and RIP next generation (RIPng).

The mechanism behind routing protocol integrity is hashing, which will be explained in detail in [Chapter 11, “Intrusion Prevention Systems.”](#)

Cisco AutoSecure

Another example of Control Plane Protection is Cisco AutoSecure. This feature also implements protection mechanisms for the data and management planes. Cisco AutoSecure provides vital security requirements to enterprise and service provider networks by incorporating a straightforward one-touch device-lockdown process. Cisco AutoSecure enables rapid implementation of security policies and procedures to simplify the security process, without having to understand all of the Cisco IOS

Software features and execute each of the many CLI commands manually. This feature uses a single command in CLI, or a single option in CCP, a one-touch device-lockdown process (discussed in the next section) that instantly configures the security posture of routers and disables nonessential system processes and services, thereby eliminating potential security threats.

Cisco AutoSecure allows two modes of operation:

- **Interactive mode:** Prompts users to select their own configuration of router services and other security-related features
- **Noninteractive mode:** Configures security-related features of the router based on a set of Cisco defaults

Interactive mode provides for greater control over the router security-related features than noninteractive mode. However, when a user needs to quickly secure a router without much human intervention, noninteractive mode is appropriate. Use AutoSecure with caution and thoroughly research its functionality prior to using it, because securing one aspect of the router might break the security of another aspect. An example of a potentially undesirable effect of using AutoSecure is that it suggests disabling Proxy ARP, which is required for the NAT process.

[Table 3-4](#) lists some of the protection mechanisms that are activated when using Cisco AutoSecure. In general terms, Cisco AutoSecure protects the router functional planes by doing the following:

- Disabling often unnecessary and potentially insecure global services
- Enabling certain services that help further secure often necessary global services
- Disabling often unnecessary and potentially insecure interface services, which can be configured on a per-interface level
- Securing administrative access to the router
- Enabling appropriate security-related logging

Table 3-4. Cisco AutoSecure Protection for All Three Planes

Plane	Action
Control plane	Disables often unnecessary and potentially insecure global services (finger, HTTP, Cisco Discovery Protocol, and so on)
Management plane	Secures administrative access to the router (password existence and minimum length, AAA, SSH, and others)
Data plane	Disables often unnecessary and potentially insecure interface services, which can be configured on a per-interface level (IP redirects, IP proxy ARP, and others)

Management Plane Security

Securing the network infrastructure requires securing the management access to the infrastructure devices. If infrastructure device access is compromised, the security and management of the entire network can be compromised. Consequently, it is critical to establish the appropriate controls to prevent unauthorized access to infrastructure devices.

Network infrastructure devices often provide a range of different access mechanisms, including console and asynchronous connections, as well as remote access based on protocols such as Telnet, rlogin, HTTP, and SSH. Some mechanisms are typically enabled by default with minimal security associated with them. For example, Cisco IOS Software–based platforms are shipped with console and modem access enabled by default. For this reason, each infrastructure device should be carefully reviewed and configured to ensure that only supported access mechanisms are enabled and that they are properly secured.

The recommended practices to secure both interactive and management access to an infrastructure device are as follows:

- **Restrict device accessibility:** Limit the accessible ports and restrict the permitted communicators and the permitted methods of access by enforcing a password policy.
- **Present legal notification:** Display legal notice developed with company legal counsel for interactive sessions.
- **Authenticate access:** Ensure that access is only granted to authenticated users, groups, and services. RBAC and authentication, authorization, and accounting (AAA) services provide mechanisms to effectively authenticate access.
- **Authorize actions:** Restrict the actions and views that are permitted by any particular user, group, or service. Actions can be restricted using ACLs and views can be limited using CLI views (discussed later in this book).
- **Ensure the confidentiality of data:** Protect sensitive data stored on management servers, such as syslog or SNMP servers, from viewing and copying. Consider the vulnerability of data in transit over a communication channel to sniffing, session hijacking, and man-in-the-middle attacks. Using management protocols with strong authentication, such as SNMPv3, mitigates confidentiality attacks that are aimed at exposing passwords, device configurations, and so on.
- **Log and account for all access:** Record who accessed the device, what occurred, and when for auditing purposes.
- **Integrity of logs and digital certificate:** Ensure the integrity of logs and digital certificate validation by synchronizing network time across all devices and network elements by using mechanisms such as Network Time Protocol (NTP). NTP helps the Security Information and Event Management (SIEM) server to perform more accurate correlation between multiple alarms that would have been received around the same time. NTP and its configuration will be covered in greater detail in [Chapter 4, “Securing the Management Plane on Cisco IOS Devices and AAA.”](#)

Secure Management and Reporting

[Figure 3-4](#) shows a management module with two network segments that are separated by a Cisco IOS router that acts as a firewall and a VPN termination device. The segment outside of the firewall (the production network) connects to all of the devices that require management. The segment inside of the firewall contains the management hosts themselves and the Cisco IOS routers that act as terminal servers. Remote access, such as SSH, to network devices connected on the production network is either disabled or highly restricted.

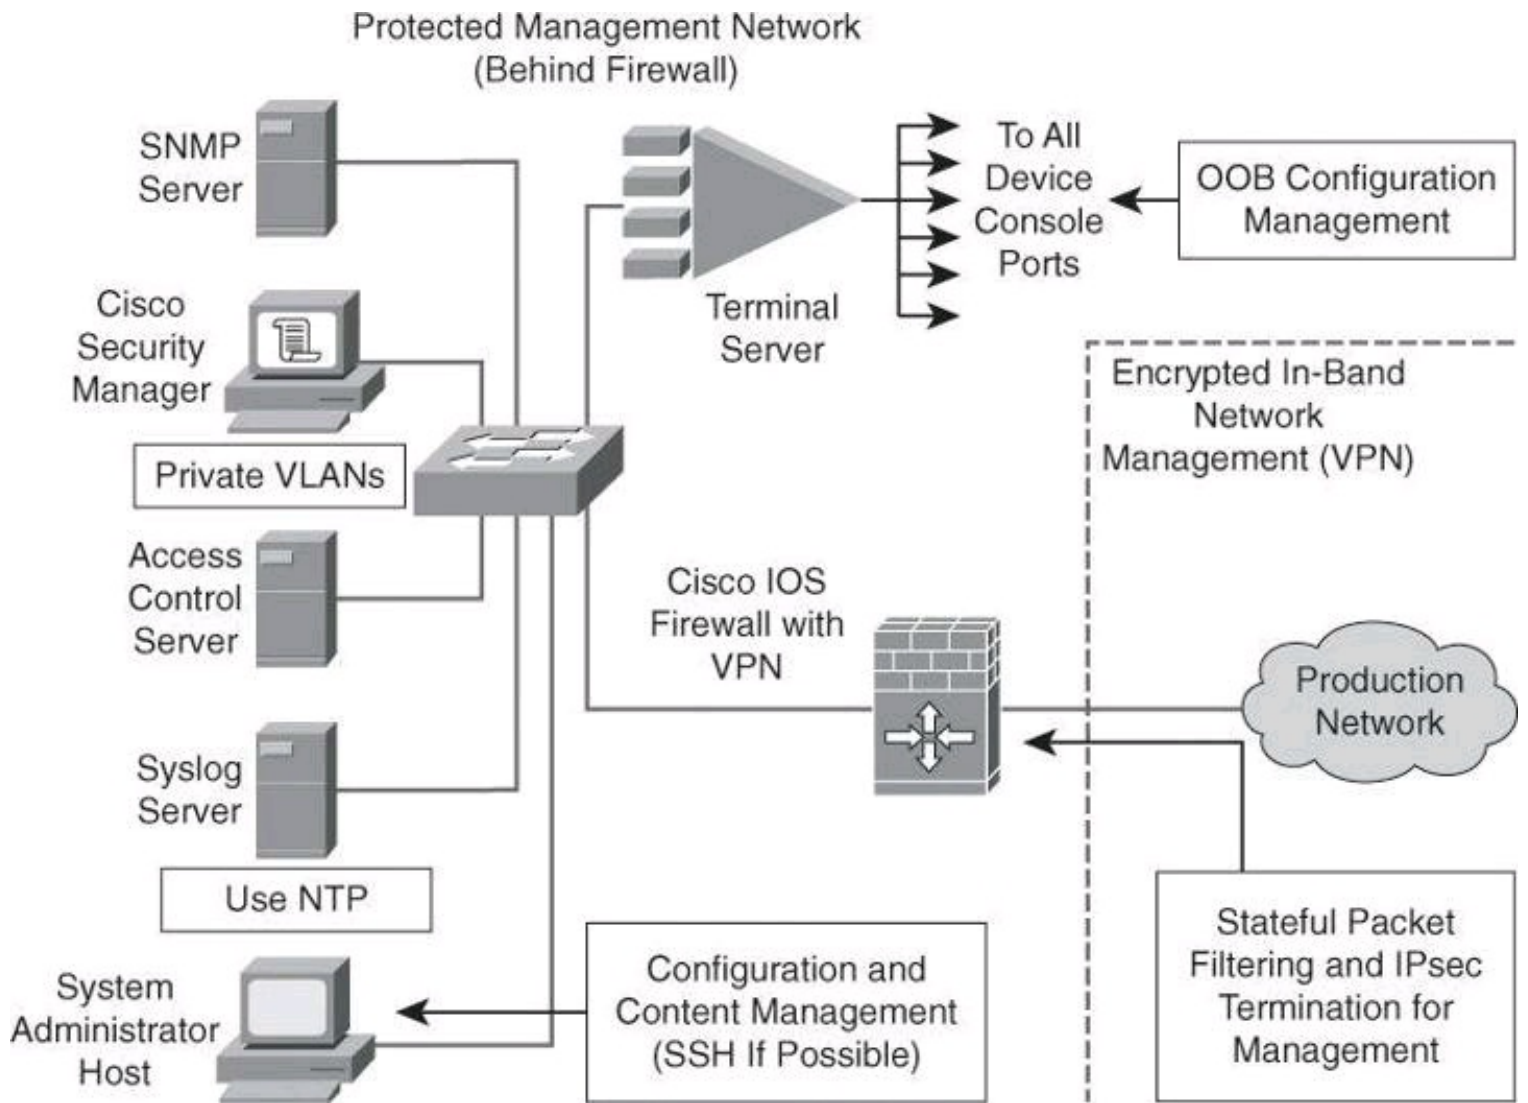


Figure 3-4. Example of Secure Management and Reporting

The information flow between management hosts and the managed devices can take two paths:

- **Out-of-band (OOB):** Information flows within a network on which no production traffic resides.
- **In-band:** Information flows across the enterprise production network, the Internet, or both.

The connection to the production network is only provided for selective Internet access, limited in-band management traffic, and IP Security (IPsec)-protected management traffic from predetermined hosts. In-band management occurs only when a management application does not function OOB, or when the Cisco device being managed does not physically have enough interfaces to support the normal management connection. This latter case employs IPsec tunnels. The Cisco IOS Firewall is configured to allow syslog information into the management segment and, in addition, Telnet, SSH, and SNMP, if these services are first initiated by the inside network. Stateful packet filtering and IPsec will be covered in detail later in this book.

Because the management network has administrative access to nearly every area of the network, it can be a very attractive target to hackers. The management module has been built with several technologies designed to mitigate such risks. The first primary threat is a hacker attempting to gain access to the management network itself. You can mitigate this threat only through the effective

deployment of security features in the remaining modules in the enterprise. All of the remaining threats assume that the primary line of defense has been breached. To mitigate the threat of a compromised device, strong access control is implemented at the firewall, and at every other possible device, to prevent exploitation of the management channel. A compromised management device cannot even communicate with other hosts on the same management subnet because private VLANs (PVLAN) on the management segment switches force all traffic from the management devices directly to the Cisco IOS Firewall, where filtering takes place.

Network administrators need to securely manage all devices and hosts in the network. Management includes logging and reporting information flow, including content, configurations, and new software, from the devices to the management hosts.

From an architectural perspective, providing OOB management of network systems is the best first step in any management and reporting strategy. Devices should have a direct local connection to such a network where possible, and where this is not possible (because of geographic or system-related issues), the device should connect via a private encrypted tunnel over the production network. Such a tunnel should be preconfigured to permit only the traffic that is required for management and reporting. The tunnel should also be locked down so that only appropriate hosts can initiate and terminate tunnels.

It is important to note, however, that OOB management is not always desirable. Often, the decision depends on the type of management applications that you are running and the protocols that are required. For example, consider a management tool with the goal of determining the reachability of all the devices on the production network. If a critical link were to fail between two core switches, you would want this management console to alert an administrator.

If this management application is configured to use an OOB network, it may never determine that the link has failed, because the OOB network makes all devices appear to be attached to a single OOB management network. With management applications used to test reachability, it is therefore essential to run the management application in-band, as user traffic would do. In-band management needs to be configured in a secure manner.

SNMP management has its own set of security needs. Use SNMP version 3 (SNMPv3) where possible, because SMNPv3 supports authentication and encryption. Keeping SNMP traffic on the management segment allows the traffic to traverse an isolated segment when it pulls management information from devices. To reduce security risks, SNMP management should only pull information from devices and should not be allowed to push changes to the devices. To ensure that management information is pulled, each device is configured with a read-only SNMP community string. You can configure an SNMP read-write community string when using an OOB network. However, be aware of the increased security risk of a plaintext string that allows modification of device configurations if an older SNMP version is used. A rudimentary example of a read string would be “GET TEMPERATURE”, where the network management server inquires about the internal temperature of a router. An example of a more dangerous write string would be “SET RELOAD”, ordering a router to perform a reboot.

As mentioned earlier, it is important that all network devices be on the same time for proper log analysis. Time synchronization is best achieved by using NTP, which automates the process. The time distributed to the devices by NTP can be authenticated, through a process similar to routing

authentication discussed earlier in this chapter. Authenticated time protects against spoofed NTP updates, which can corrupt logs and render digital certificates unusable.

Role-Based Access Control

RBAC is an access control approach that restricts user access based on the role of the user, along with their individual identity. Roles are created for job or task functions and assigned access permissions to specific assets. Users are then assigned to roles, through which assignment they acquire the permissions that are defined for the role. This means users are not assigned permissions directly, but only acquire them through role assignment, as shown in [Figure 3-5](#). Managing the rights of an individual user becomes as simple as assigning appropriate roles to the user.

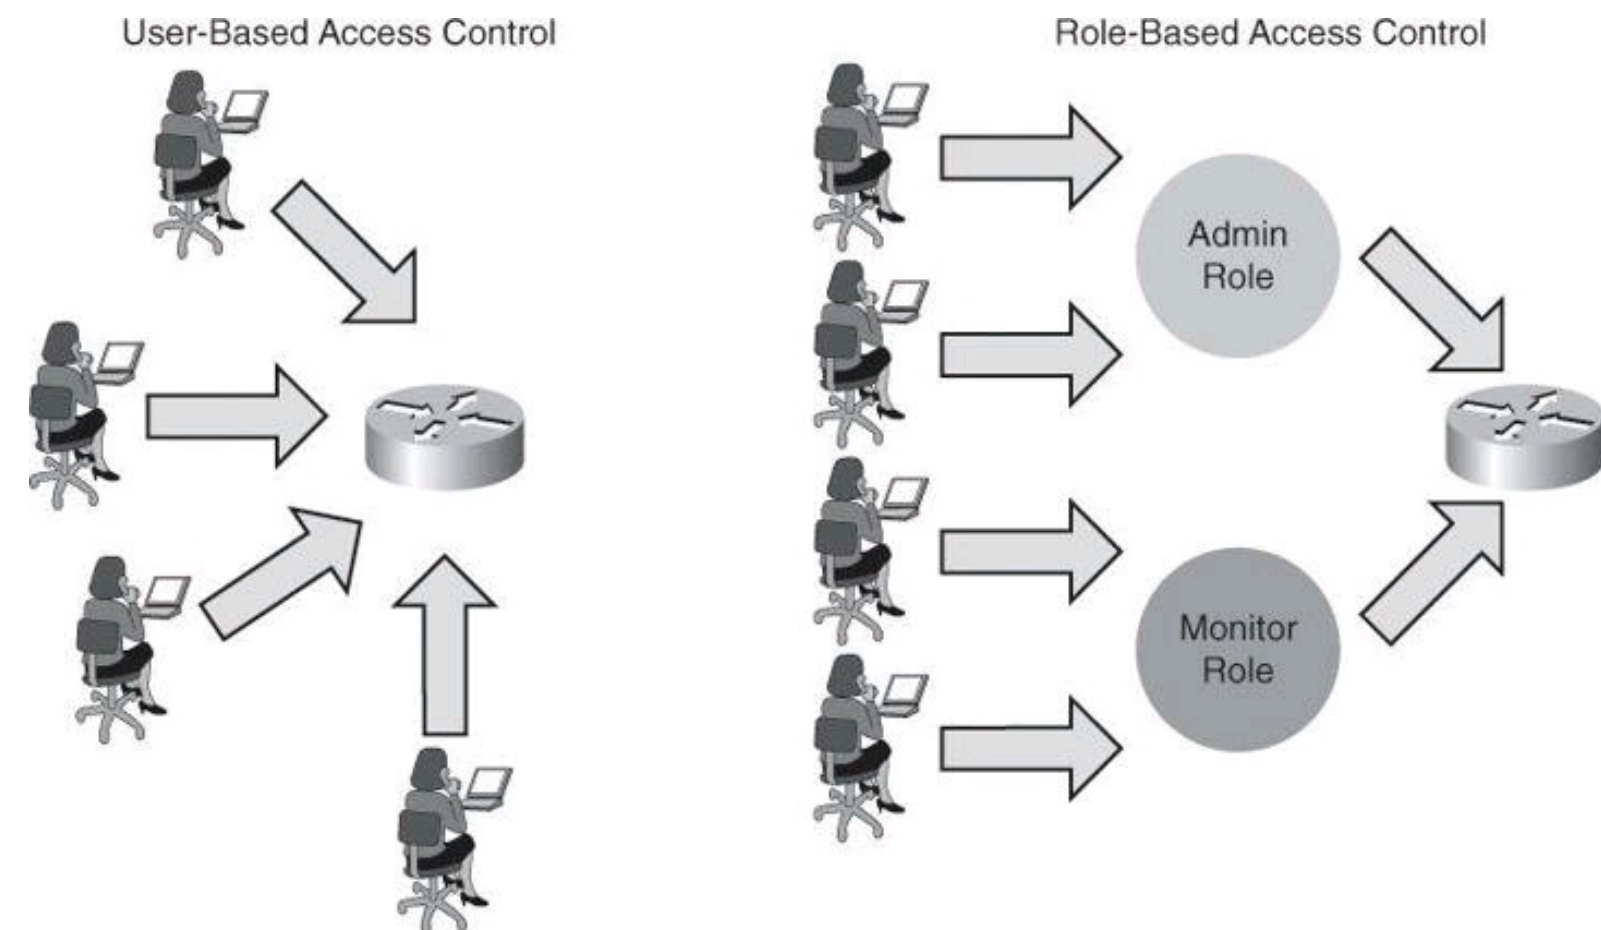


Figure 3-5. Advantage of Role-Based Access Control

With Cisco IOS, the role-based CLI access feature implements RBAC for router management access. This feature allows you to create different “views” of router configurations for different users. Views define which commands are accepted from different users and which configuration information is visible to them. With role-based CLI access, you can exercise better control over Cisco networking devices.

For scalability, users, permissions, and roles are usually created and maintained in a central repository, to make the access control policy available to multiple devices using it.

Deploying AAA

AAA servers are typically used as a central repository of authentication credentials (the users, answering the question “who is trying to access the device?”), authorization rules (the “what” users can accomplish), and accounting logs (the “what users did” part of the equation).

[Figure 3-6](#) shows an example of the authentication and authorization process using an external AAA server such as Cisco Secure Access Control System (ACS) to provide AAA services to a network for management purposes.




Figure 3-6. Authentication and Authorization Using an External AAA Server

In this chapter, we are concerned with limiting management access to our network equipment, and as such we will be discussing how to use AAA to limit the traffic to routers and switches for the purpose of owning the devices to manage them. In a later chapter, we will discuss the concept of using AAA to control traffic going through the devices.

In step 1, the client establishes a connection with the router.

In step 2, the router prompts the user for a username and password. The user sends her credentials to the router.

In step 3, the router passes the user's credentials to the AAA server and advises the router that the user is authenticated.

In step 4, the user is authorized to access the router (administrative access) based on information found in the AAA server authorization database.

Data Plane Security

Among the laundry list of ways to protect the data plane, some that we will see in this book include

- Access control lists
- Private VLAN
- Firewalling
- Intrusion Prevention System (IPS)

Before discussing the details for data plane security, let's review the most common means of protecting it, access control list (ACL) filtering, and then discuss briefly Layer 2 protection. Both of these topics are discussed in greater detail in subsequent chapters.

Access Control List Filtering

ACLs perform packet filtering to control which packets move through the network and where. Such control provides security by helping to limit network traffic, restrict the access of users and devices to the network, and prevent traffic from leaving a network. IP access lists can reduce the chance of spoofing and DoS attacks and allow dynamic, temporary user access through a firewall.

The following are the most common reasons to use ACLs:

- **Block unwanted traffic or users:** ACLs can filter incoming or outgoing packets on an interface, thus controlling access based on source addresses, destination addresses, or user authentication. You also can use ACLs to determine which types of traffic are forwarded or blocked at the router interfaces. For example, you can permit email traffic to be routed but at the same time block all Telnet traffic.
- **Reduce the chance of DoS attacks for internal devices:** There are a number of ways to reduce the chance of DoS attacks. For example, by specifying IP source addresses, you can control whether traffic from hosts, networks, or users accesses your network. You can filter on specific Time to Live (TTL) values in packets to control how many hops a packet can take before reaching a router in your network. By configuring the TCP Intercept feature, you can prevent servers from being flooded with requests for a connection.
- **Mitigate spoofing attacks:** ACLs allow security practitioners to implement recommended practices to mitigate spoofing attacks. The guidelines that are found in several RFCs provide basic filtering, and can be easily deployed using ACLs.
- **Provide bandwidth control:** An ACL on a slow link can prevent excess traffic
- **Classify traffic to protect other planes:** You can place an ACL on inbound vty (Telnet) line access from certain nodes or networks. For the control plane, ACLs can control routing updates being sent, received, or redistributed.

Antispoofing

One efficient use of ACLs is as an antispoofing mechanism. Spoofing protection involves discarding traffic that has an invalid source address. Network security baseline includes source IP spoofing protection that is based on BCP38/RFC 2827 ingress traffic filtering.

Packets with spoofed source IP addresses represent a security risk because they are often used to conduct an attack, to evade traceability and bypass access controls. These packets may also be used to direct an attack at a spoofed source, something that is known as a reflection attack, as shown in [Figure 3-7](#). In this figure, the hacker sends packets with a source IP address of a third party. Hosts on network 171.1.0.0 reply to what they think to be the initiator of the echo-request and send the echo-replies to 200.1.1.1, which is not the originator of the ping.

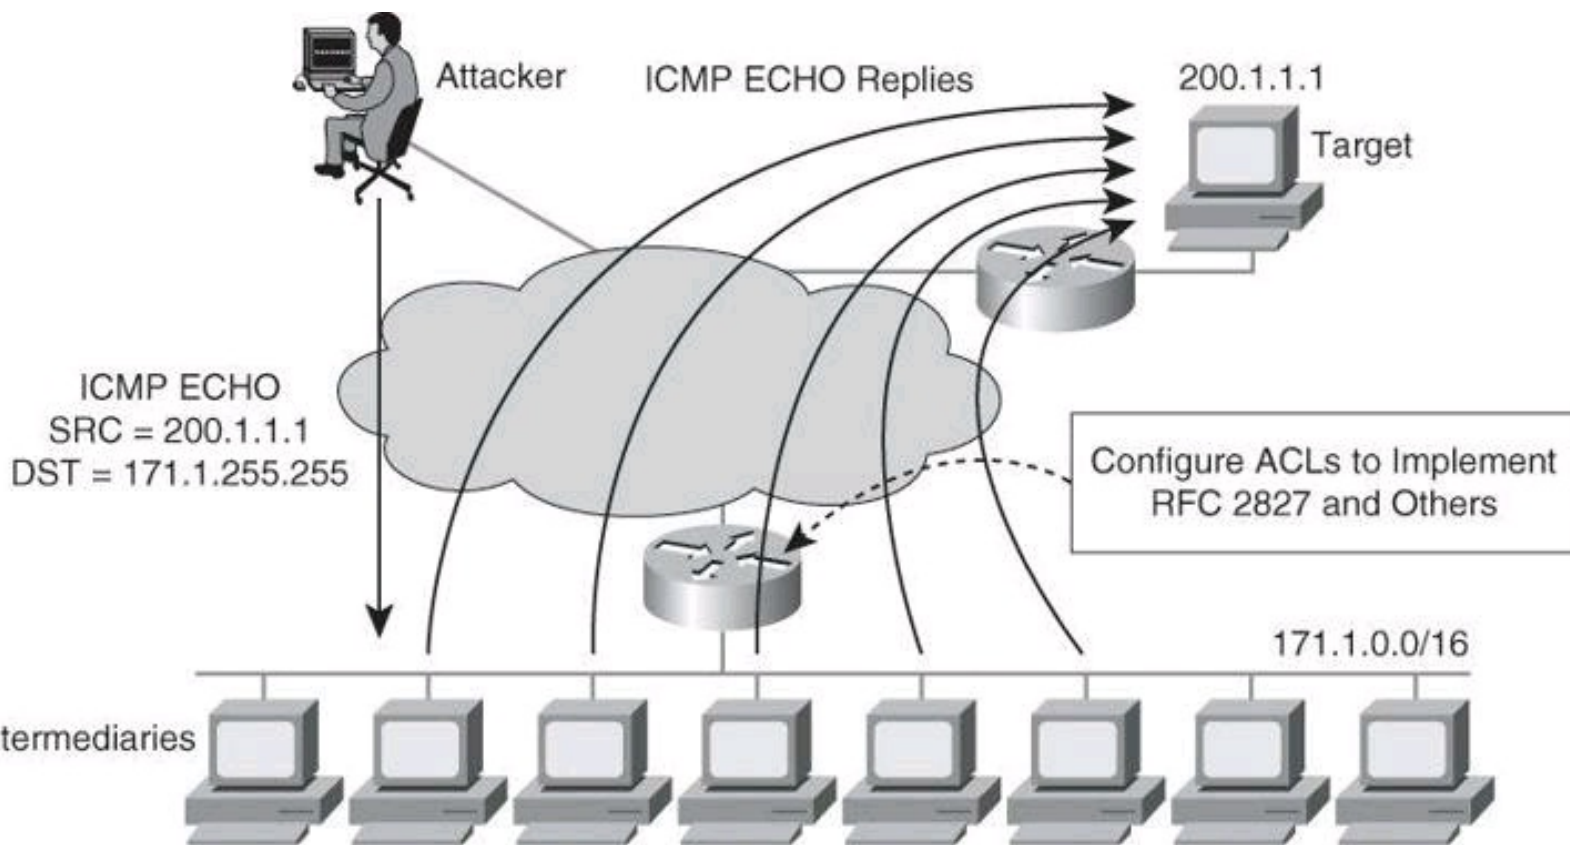


Figure 3-7. Attack Involving Spoofed Address

A variant of the attack shown in [Figure 3-7](#) would be for the attacker to turn the 171.1.0.0 subnet workstations into members of a botnet, as explained in [Chapter 1](#). The hacker could then order those zombies to launch TCP SYN packets to the victim, but spoofing as source IP address a third party. The victim would route its reply to the third-party address. This particular attack would create three casualties: the zombie site, which would have a high volume of outbound traffic; the destination victim; and the third site, whose address was invoked as the source IP address of the original request.

Spoofed traffic with an invalid source IP address may include traffic from either of the following:

- RFC 1918, Documented Special Use Addresses (DSUA), or unallocated IP address range
- Valid IP network address range, but not originating from the associated legitimate network

Implementing BCP38/RFC 2827 ingress traffic filtering to address source IP address spoofing renders the use of invalid source IP addresses ineffective, forcing attacks to be initiated from valid, reachable IP addresses. This implementation is beneficial because it enables greater success in tracing the originator of an attack.

ACLs are the traditional technique for filtering forged source IP addresses. However, ACLs are not dynamic in nature, requiring manual configuration changes, and may have an impact on the performance of a device. Other features, such as Unicast Reverse Path Forwarding (uRPF), can be used to complement the antispoofing strategy.

Note

Download document ID 13608, “Cisco Guide to Harden Cisco IOS Devices,” from Cisco.com for an informal discussion of some Cisco IOS system device configuration

settings that network administrators could consider to harden their routers, especially perimeter routers.

Layer 2 Data Plane Protection

Data plane protection mechanisms depend on feature availability for specific devices. In a switching infrastructure, these Cisco Catalyst integrated security capabilities provide data plane security on the Cisco Catalyst switches using integrated tools:

- Port security prevents MAC flooding attacks.
- DHCP snooping prevents client attacks on the DHCP server and switch.
- Dynamic ARP Inspection (DAI) adds security to ARP by using the DHCP snooping table to minimize the impact of ARP poisoning and spoofing attacks.
- IP Source Guard prevents IP spoofing addresses by using the DHCP snooping table.

Some of the above controls will be discussed later in this book.

Now that we have seen in general terms the need for protecting the network infrastructure, let's have a look at a GUI tool to help us harden our Cisco IOS routers.

Cisco Configuration Professional

CCP, shown in [Figure 3-8](#), is a GUI-based device management tool for Cisco access routers, specifically Integrated Services Routers (ISR) and Integrated Services Routers Generation 2 (ISR G2). This tool simplifies routing, firewall, IPS, VPN, unified communications, WAN, and LAN configuration through GUI-based wizards.

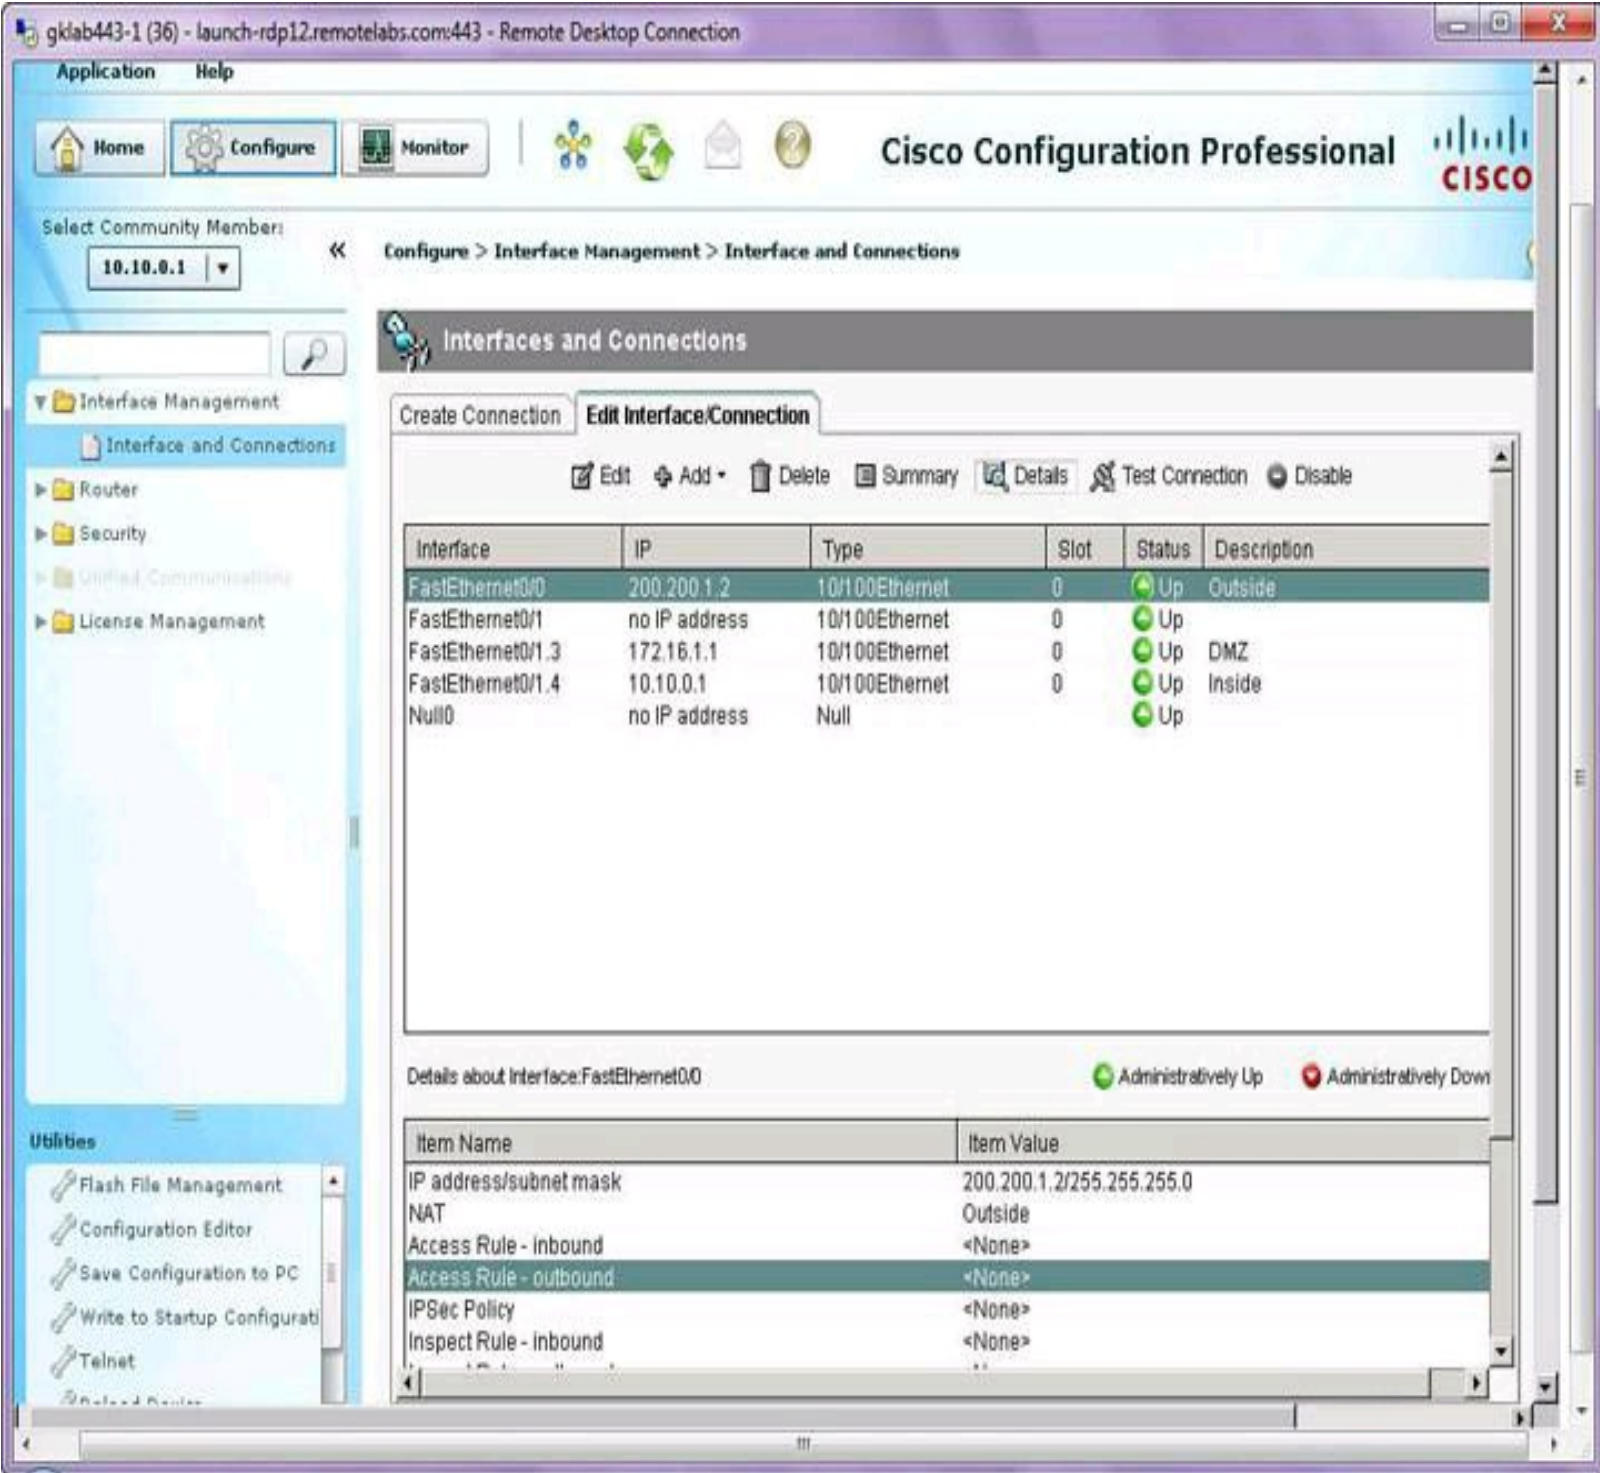


Figure 3-8. CCP GUI

CCP is a valuable productivity-enhancing tool for network administrators, used for deploying routers with increased confidence and ease. It offers a one-click router lockdown and an innovative voice and security auditing capability to check and recommend changes to router configuration. CCP also monitors router status and troubleshoots WAN and VPN connectivity issues. The GUI includes options to configure Cisco Services-Ready Engine (SRE) modules, facilitating the integration of hardware-based functionality and centralizing their configuration in the same place where the hosting router is configured.

CCP is free, and you can download it from <http://www.cisco.com/go/ciscocp>. CCP can be preinstalled in the flash memory of the router, typically in its express version.

The following list summarizes the benefits of using CCP:

- Supports configuration of advanced technology deployment on ISRs and ISR - G2 through a single integrated tool
- Simplifies routing, firewall, IPS, VPN, unified communications, WAN, and LAN configuration through GUI-based easy-to-use wizards
- Monitors router status and troubleshoots WAN and VPN connectivity issues
- Provides smart wizards for routing and security configuration
- Offers license management for ISR-G2

The following are features of CCP Express:

- A lightweight version of CCP
- Router flash
- LAN and WAN interfaces

By relying on best-practice configurations that are approved by the Cisco Technical Assistance Center (TAC), network administrators can take advantage of CCP, which does the following:

- Lowers the total cost of ownership (TCO) of Cisco routers
- Reduces human errors
- Simplifies initial setup in voice deployments
- Helps ensure proper linkage between users, dialing plans, and voicemail settings

CCP offers smart wizards and advanced configuration support for the following:

- LAN and WAN interfaces
- Network Address Translation (NAT)
- Stateful and application firewall policy
- IPS, IPsec, and SSL VPNs
- Quality of service (QoS)
- Cisco Network Admission Control (NAC) policy features

CCP enables IT managers to easily organize and manage multiple routers at a single site.

CCP Initial Configuration

Devices that are shipped with CCP have a default configuration that allows you to connect a PC to an Ethernet port on the device and start configuration immediately. This initial configuration is accomplished using CCP Express. CCP Express and a factory default configuration file are installed in flash memory on routers that are shipped with CCP. You can connect a PC directly to the device, and then use CCP Express to configure LAN and WAN connections, a firewall, and NAT and make security settings before you place the device on the network in which it will operate. After you have configured the device and connected it to the network, you will be able to use CCP to connect to the device over the network and make advanced configurations.

[Figure 3-9](#) provides a visual aid to the following configuration procedure:

Default Configuration

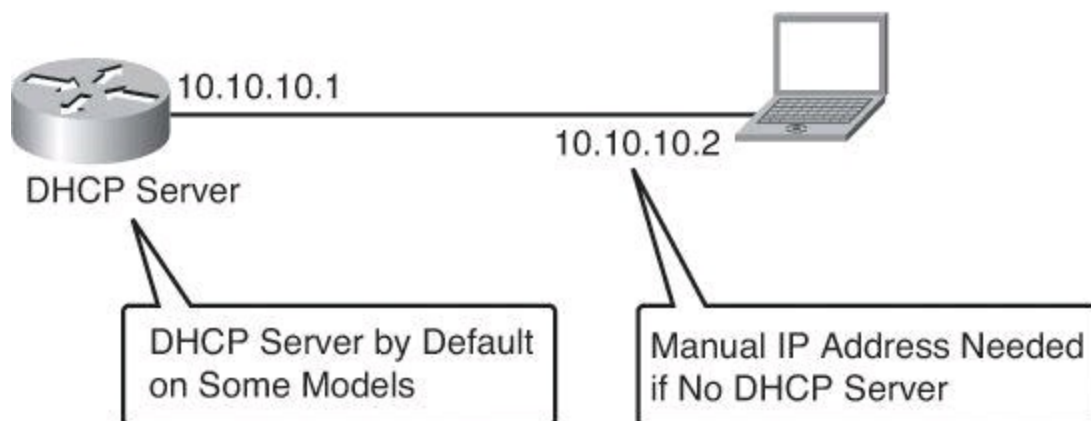


Figure 3-9. Connecting to the Default CCP Settings of an ISR

Step 1. The device default configuration file configures an IP address for one Ethernet interface, and that may configure the device as a DHCP server. See [Table 3-5](#) to determine whether the device is configured as a DHCP server, and which Ethernet port to connect the PC to.

Table 3-5. Cisco Configuration Professional—Defaults According to Models

Device Model	DHCP Server	Connect PC to the Applicable Ethernet Port
Cisco 815, Cisco 86x, Cisco 88x, Cisco 180x, Cisco 1805, Cisco 1811, Cisco 1812	Yes	Depending on the model, connect to any ACT Lnk port or LAN port or SWITCH port or PWR Lnk port
Cisco 1841, Cisco 1861, Cisco 2801, Cisco 2811	No	Fast Ethernet 0/0
Cisco 28xx, Cisco 38xx	No	Gigabit Ethernet 0/0
Cisco 19xx, Cisco 29xx, Cisco 39xx	No	Gigabit Ethernet 0/0

Refer to CCP and ISR documentation on Cisco.com for an up-to-date list of defaults and detailed information on applicable Ethernet ports depending on model.

Step 2. Connect the PC to the appropriate port listed in [Table 3-5](#).

Step 3. Configure the PC IP address by doing one of the following:

- If the device is configured as a DHCP server, ensure that the PC is configured to accept an IP address from a DHCP server.
- If the device is not configured as a DHCP server, configure the static IP address 10.10.10.2 on the PC, and use the subnet mask 255.255.255.248.

Step 4. Open an Internet Explorer browser window, and enter the IP address **10.10.10.1** to connect to the device and start CCP Express.

Step 5. Complete the CCP Express wizard to configure the device.

When you have completed initial setup and given the device an IP address on your LAN, you can use CCP to connect to the device and perform additional configuration.

To use CCP with an already deployed device, enter the commands shown in [Table 3-6](#).

Table 3-6. Command to Provision a Deployed Device with CCP Support

Feature	Requirement	Configuration
Secure Access	SSH and HTTPS	Router(config)# ip http secure-server
		Router(config)# ip http authentication local
		Router(config)# line vty 0 4
		Router(config-line)# login local
		Router(config-line)# transport input ssh
		Router(config-line)# transport output ssh
Nonsecure Access (clear text)	Telnet and HTTP	Router(config)# ip http server
		Router(config)# ip http authentication local
		Router(config)# line vty 0 4
		Router(config-line)# login local
		Router(config-line)# transport input telnet
		Router(config-line)# transport output telnet
User Privilege	Level 15	Router(config)# username cisco privilege 15 secret 0 cisco

Note

In the **username** command of [Table 3-6](#), the **0** indicates that the following password is clear text, which should be avoided. If the value would have been set to **5**, it would have indicated that the following password would appear encrypted in configuration files.

Also, connecting CCP to a new device for the first time can be accomplished by choosing **Application > Setup New Device** in CCP, which you have manually installed on your PC. Connect your local PC to the console port of the router. This approach, shown in [Figure 3-10](#), sounds counterintuitive—using a GUI through a serial port—but works, where CCP uses the console connection of your local PC to attempt to reach the router to configure it. For this to work, ensure that no other terminal monitor application is currently accessing the console port of the router.

1 - Introduction

2 - Configuring Device

3 - Configuration Summary

Step 1 - Introduction

Setting up a new device


Use this wizard to setup a device to enable Cisco Configuration Professional to manage it. This wizard can be used on a device if its new, has been reset to factory default or has existing configuration. The device will be configured with authentication credentials, vty lines and an https server.

These steps are described in the CCP Quick Start Guide.

Prerequisites :

1. Ensure the computer running CCP is connected to the powered up device over the console port.
2. Ensure the device baud is set to its default value
3. Ensure a community is selected

See below for a typical connection between a PC/laptop and a router. If a serial port is not available a USB to RS-232 converter may have to be used.



< Back

Next >

Cancel

Figure 3-10. CCP Used to Discover a Device Using the Serial Connection of a PC

Cisco Configuration Professional User Interface and Features

Although every aspect of CCP is important and useful, due to space constraints, we will focus in the following pages only on the most relevant options and features of CCP.

Note

If you would like to follow along with these explanations of CCP, you can download the Cisco Configuration Professional Demo version from Cisco.com, using your Cisco.com user ID. This simulator is prefilled with routers and configurations so you can practice your CCP navigational and configuration skills without interfering with equipment currently in production. However, some features are not available or configurable in Demo or Offline modes.

CCP eliminates the need for multiple device managers by providing a single tool to configure and manage devices. The user interface makes it easy to manage networking features. The main parts that define the user interface are shown in [Figure 3-11](#) and include the following:

- **Menu bar:** The row of menus across the top of the window. The menu bar offers application services, a list of open windows, and online help.
- **Toolbar:** The row of icons directly below the menu bar. These icons represent the most often used application services and most often configured networking features.
- **Left navigation pane:** This navigation pane is the scalable panel on the left side of the content pane in which you select the features to configure and monitor.
- **Content pane:** This pane is the right side of the workspace in which windows appear. You view reports here and enter information that configures networking features.
- **Status bar:** The status bar is the bar at the bottom of the window where CCP displays the status of the application. A closed lock icon refers to a secure connection and an opened lock icon refers to an unsecure connection.

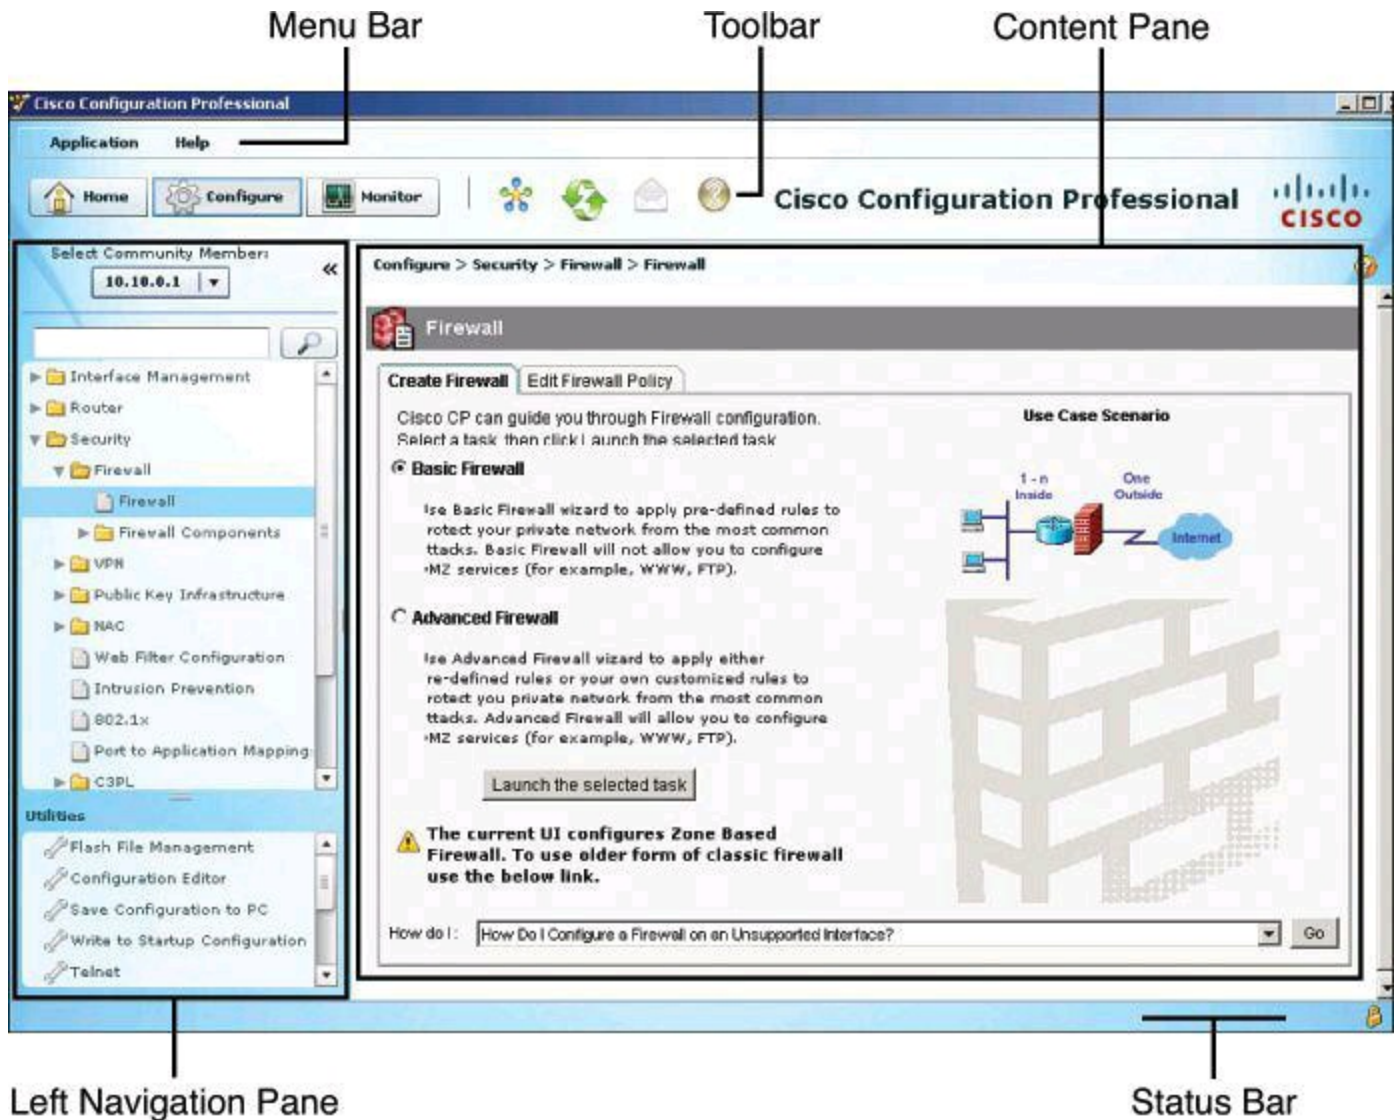


Figure 3-11. Navigating the GUI of Cisco Configuration Professional

Menu Bar

The row of menus across the top of the window offer application services and help. [Figure 3-12](#) identifies the options in the Application menu, which are described here:

- **Manage Community:** Enables you to create a new community or choose an existing community. A community member would be a router that you manage with CCP.

- **Setup New Device:** Enables you to set up a new device.
- **Create User Profile:** Enables you to restrict users from using all of the features that are available in the left navigation pane.
- **Import User Profile:** Enables you to import a user profile.
- **Options:** Enables you to set user preferences such as log level, show community at startup, and show CLI preview parameters.
- **Template:** Enables you to create, edit, or apply a template.
- **Work Offline:** Enables you to work with CCP in offline mode.
- **Exit:** Exits the CCP application.



Figure 3-12. CCP Toolbar

Toolbar

The following CCP features are available from the toolbar at the top of the window, as shown in [Figure 3-12](#):

- **Home:** Click **Home** to display the Community View page, which summarizes the community information and allows you to add, edit, and discover devices and view the discovery and router status of each device.
- **Configure:** Click **Configure** to display the features that you can configure on a chosen device. The features are displayed in the left navigation pane, as shown in [Figure 3-12](#). Note that if a feature (router, security, or voice) is not supported on a device, that feature is not displayed in the left navigation pane. However, if the Cisco IOS version that is installed on the device does not support a specific feature, but an upgrade does support it, then that feature is disabled (grayed out) in the left navigation pane.
- **Monitor:** Click **Monitor** to display the router and security features that you can monitor for a chosen device. The features are displayed in the left navigation pane.
- **Manage Community:** Click this icon to open the Manage Community dialog box, where

you can add a new community or edit an existing community.

• **Refresh:** Click this icon to do the following:

- Rediscover the selected device in the Select Community Member drop-down menu
- Rediscover and reload the current feature

Note that the Refresh function is not available for offline mode, and is available in online mode only after successful discovery of one or more devices. Clicking Refresh refreshes the device that is selected in the Select Community Member drop-down menu. Choosing Home > Dashboard, selecting a device, and clicking Refresh does not refresh that device.

Navigation Pane

The following options are found in the navigation pane on the left of the screen, as shown in [Figure 3-11](#).

Interface Management

The CCP connection wizards guide you through LAN and WAN configurations, and check the information that you enter against the existing configuration, warning you of any problems. The supported interfaces are

- LAN
- WAN
- Wireless LAN
- Cellular WAN (3G)
- Analog trunks
- Digital trunks

In addition, CCP can be used to configure EnergyWise features for power management.

The Interface Management folder on the navigation pane provides configuration support for interfaces and service modules. CCP can configure the modules management, install applications, and upgrade images. The following service modules are supported:

- Network modules
- Enhanced network modules
- Cisco Wide Area Application Services (WAAS) modules
- Advanced Integration Modules (AIM)
- Voice Network Modules (VNM)
- Services-Ready Engine (SRE) modules

Support for Services-Ready Engine Virtualization (SRE-V) allows you to do the following:

- Install SRE-V software on SM-SRE 700 and 900 modules
- Edit console manager and hypervisor IP configurations
- Launch the vSphere client user interface from CCP

- Configure users, groups, roles and permissions, and syslog
 - View license details
-

Note

Cisco SRE-V is a branch-office infrastructure platform that combines computing, networking, storage, virtualization, and unified management into a cohesive system. It enables the VMware vSphere Hypervisor to be provisioned on a Cisco Services-Ready Engine (SRE) Service Module and host one or more virtual machines running Microsoft Windows Server operating system. The entire system is integrated with Generation 2 of the Cisco Integrated Services Router (ISR-G2). Further information can be found at Cisco.com.

Router

The Router folder, in the left navigation pane, displays the routing window with the configured static routes and Routing Internet Protocol (RIP), Open Shortest Path First (OSPF), and Extended Interior Gateway Routing Protocol (EIGRP) configured routes. From this window, you can review routes, add new routes, edit existing routes, and delete routes.

Under the Router folder, the administration can also set the hostname, username, and password. Other supported routing functions include NAT, DHCP, Domain Name System (DNS), QoS, and Cisco Performance Routing (PfR).

The Router folder is also used for overall management options, related to AAA, management access, logging, and others.

Security

CCP can be used to manage functions that comprise a large portion of the topics discussed in this book. They include

- Device hardening via the One-Step Lockdown and Security Audit wizards
- VPN options for IPsec and SSL site-to-site and remote access deployments
- Zone-based firewall wizards, as well as classic Cisco IOS firewall configuration options for backwards compatibility
- Cisco IOS IPS wizards

Each is discussed in turn next.

Device Hardening Security Audit is a feature that examines your existing router configurations and then updates your router in order to make your router and network more secure. Security Audit is based on the Cisco AutoSecure feature. It performs checks on and assists in configuration of almost all of the AutoSecure functions.

Security Audit operates in one of two modes. The first mode is the Security Audit wizard, which lets you choose which potential security-related configuration changes to implement on your router. The second mode is One-Step Lockdown, which automatically makes all recommended security-related configuration changes.

VPN Comprehensive VPN configuration can be accomplished using CCP options. A wizard-like

VPN Design Guide is available, helping you to determine which kind of VPN to configure, among site-to-site, remote access (SSL and IPsec using Easy VPN and Virtual Tunnel Interface [VTI]), Dynamic Multipoint VPN (DMVPN), and Group Encrypted Transport (GET) VPN. The input parameters include information about what type of user you are, the type of equipment that the router establishes VPN connections with, the type of traffic that the VPN will carry, and other features that you need to configure. After you provide this information, the VPN Design Guide recommends a VPN type and allows you to launch the wizard that will enable you to configure that type of VPN.

For ease of configuration, CCP offers a VPN option to create a text file that captures the VPN configuration of the local router. This file can be used to configure a remote router that enables it to establish a VPN connection to the local router.

Firewall The firewall option simplifies access control configuration using smart wizards. Two firewall wizards are available:

- **Basic Firewall:** Creates a firewall using default rules. The use case scenario shows a typical network configuration in which this kind of firewall is used.
- **Advanced Firewall:** Leads you through the steps of configuring a firewall. You can create a demilitarized zone (DMZ) network and specify an inspection rule. The use case scenario that is shown when you choose this option shows you a typical configuration for an Internet firewall.

CCP provides preconfigured application security policies that you can use to protect the network. Using a simple slider bar, you can select the security level that you want and view a description of the security it provides. The wizard summary screen displays the policy name and the configuration statements in the policy.

Comprehensive editing tools are available to maintain the firewall ruleset, with quick edit options and a graphical view of the ruleset that allows you to understand the scope of the rule and the zones and interfaces involved.

You can also view the details of the policy by clicking the Application Security tab and choosing the name of the policy. Application inspection rules for deep packet inspection are available.

IPS The IPS option simplifies intrusion management using a wizard-based approach to configuration. CCP allows for a simplified initialization, update, and management of IPS signatures, including intuitive options for signature tuning in order to improve the accuracy and sensitivity of the system. This allows for improved chances of managing false positives and false negatives.

The IPS option also includes a Security Dashboard, which displays threat information in the form of top threats table, and allows the administrator to visualize the events and incidents generated by the IOS IPS functionality.

Management, Monitoring, and Troubleshooting Comprehensive monitoring and troubleshooting options are available by clicking the Monitor button in the toolbar shown in [Figure 3-12](#). Convenient tools allow detailed troubleshooting, including connectivity testing tools that generate real traffic and provide visual cues as to potential configuration or connectivity issues.

Content Pane

The content pane, located on the right side of the workspace, is where the configuration takes

place. The content pane changes depending on the selection you make on the navigation pane, located on the left of the window. In addition to configuring networking features from the content pane, you can view and monitor the equipment.

Status Bar

The status bar appears at the bottom of the window and displays the current communication status with the device. A locked padlock icon indicates that CCP has a secure connection with the chosen community member. The unlocked padlock icon indicates that CCP has a nonsecure connection with the chosen community member.

Cisco Configuration Professional Building Blocks

When using Cisco Configuration Professional, you can take advantage of the following tools to facilitate and direct the process. These building blocks simplify operations, allow RBAC, and provide bulk configuration and monitoring options for multiple devices.

- **Communities:** Groups of devices that share common components. Communities can be used to group devices based on geographic location, device function, and so on.
- **Templates:** Parameterized configuration files that can be created once with a common configuration and applied multiple times by using configuration variables that change value for each device.
- **User profiles:** GUI views that allow RBAC over CCP menus and options.

RBAC is discussed further in [Chapter 4](#).

Communities

A community is basically a group of devices (known as community members). A single community can contain a maximum of ten devices. You can create a community and then add devices to it based on some common parameters. For example, you can create communities that are based on the location of the devices. You can create a San Jose community and add devices to it, and then you can create a Bangalore community and add devices to it, and so on.

Before you begin using CCP, you must first create a community and then add devices to that community. When you start CCP for the first time, CCP automatically creates a community for you, to which you can add devices. Devices are discovered on the network, and their configuration is uploaded to CCP.

When you add a device to a community, you must specify its IP address or hostname, credential information (username and password), and whether the communication will be secured. CCP uses this information to discover the device. After you discover the device, you can configure and monitor it. Again, up to ten devices can be added to each community.

You can create and manage communities from the Manage Community dialog box.

Creating Communities

The following list shows how to create and manage communities in CCP:

Step 1. Use the Manage Community dialog box to create communities. The Manage Community dialog box automatically displays when you start Cisco Configuration

Professional, with a community called New Community created by default.

You can also open the Manage Community dialog box in the following ways, shown in [Figure 3-13](#):

- From the toolbar, click the **Manage Community** icon.
- From the menu bar, choose **Application > Manage Community**.

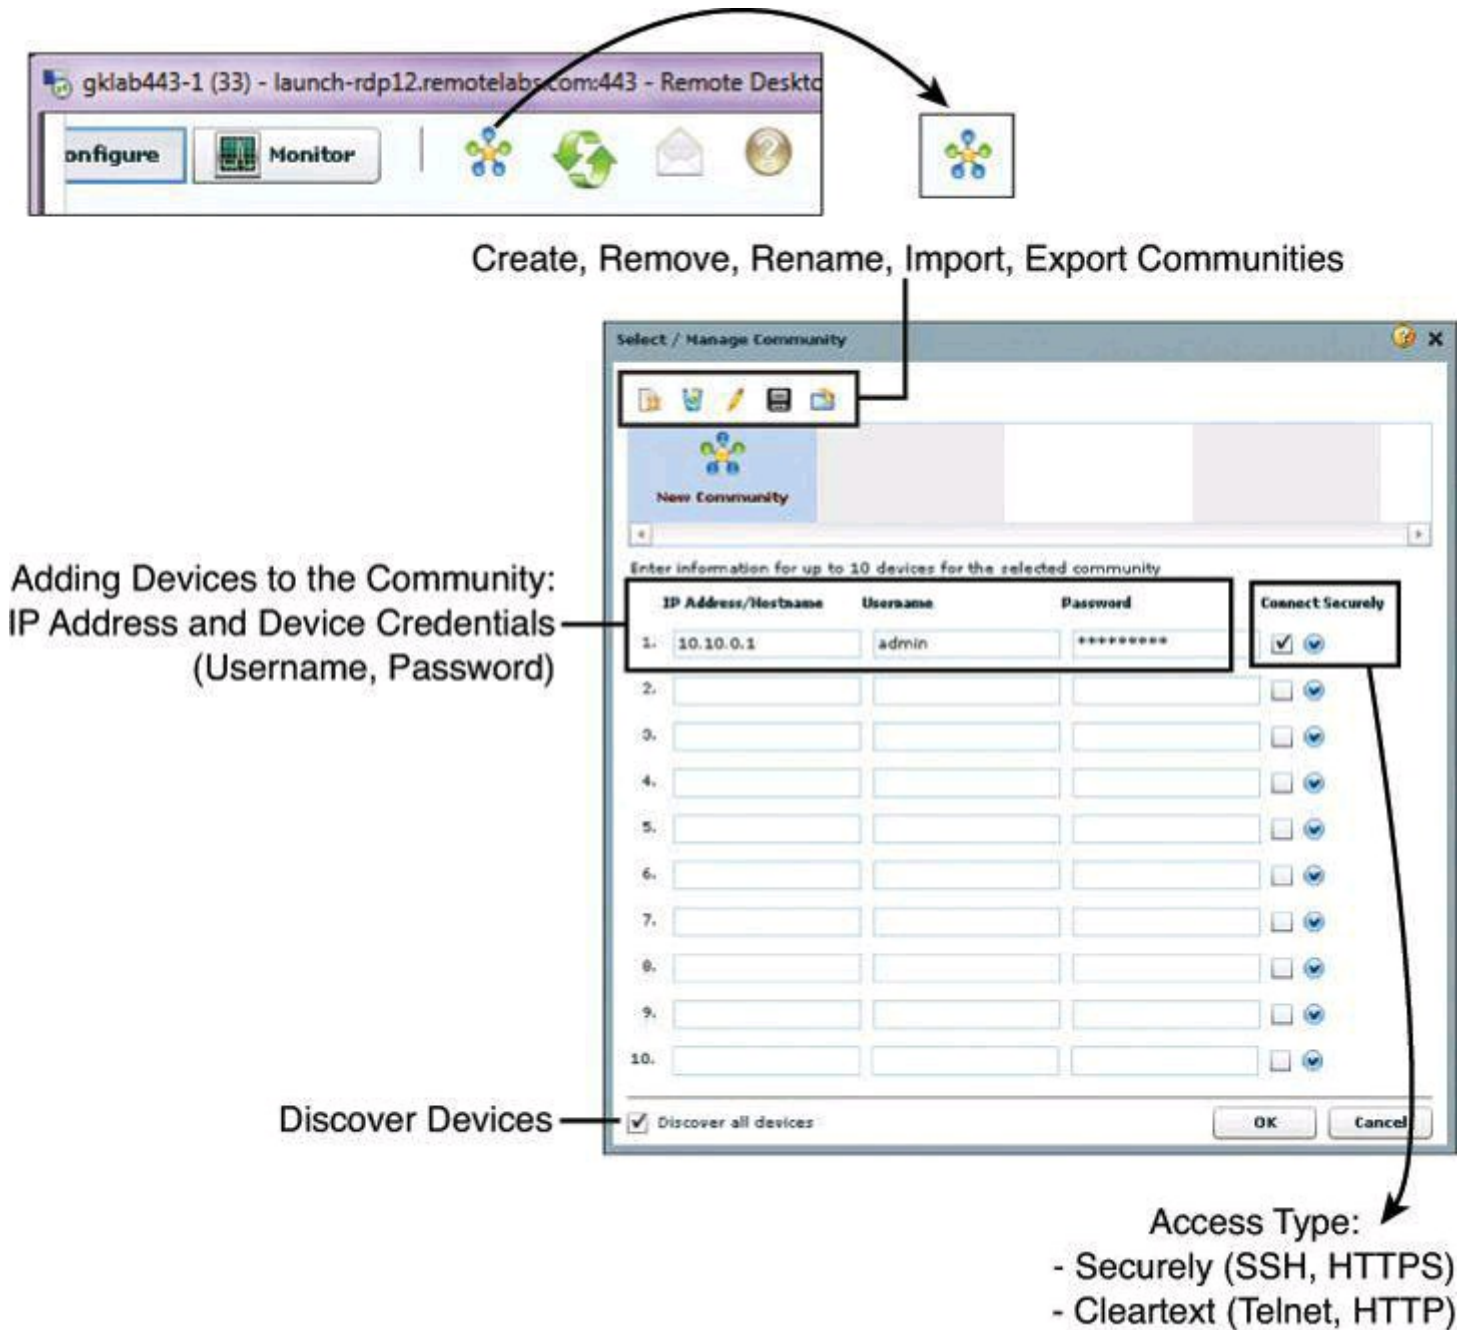


Figure 3-13. Creating Communities from the Application Menu or with the Icon on the Toolbar

Step 2. In the Manage Community dialog box, enter the IP address or hostname and the username and password information for the devices that you want to configure. If you enter the default username cisco and default password cisco, the Change Default Credentials dialog box opens. For security reasons, you must change the default credentials to new credentials.

Step 3. If you want CCP to connect securely with the device, check **Connect Securely**. To view the port information, click the down arrow next to the Connect Securely check box.

Step 4. If you want to change the default port information, click it, and then enter a new port value.

Make sure that CCP can access the device at the specified secure or nonsecure ports.

Step 5. If you want CCP to discover all the devices in a community, check **Discover All Devices**. If you want, you can choose to discover the devices later, from the Community View page.

Step 6. Click **OK**. The Community View page opens. It displays the information about the devices in the community.

Managing Communities

After you create a community and add devices to it, you can view the information for that community by clicking the only available selection in the navigation pane, a page called Community View. Clicking Community View refreshes the information appearing in the Community Information pane. From the Community Information pane, you can manage the devices (community members) in a community, such as add devices to a selected community, edit device information, delete devices, discover the devices, view information about the discovery process, and view hardware and software information about a selected device.

[Figure 3-14](#) illustrates the Community View option in CCP. The list of devices appears in the content pane, and the selected device can be quickly chosen using the Select Community Member drop-down menu. The buttons at the bottom of the window include the Discover button, which enables you to discover or rediscover devices, the Discovery Details button, which displays a log of events and errors that may arise during the discovery process, the Manage Devices button, which enables you to manage device IP addresses and credentials, and the Router Status button, which displays the status of the devices.

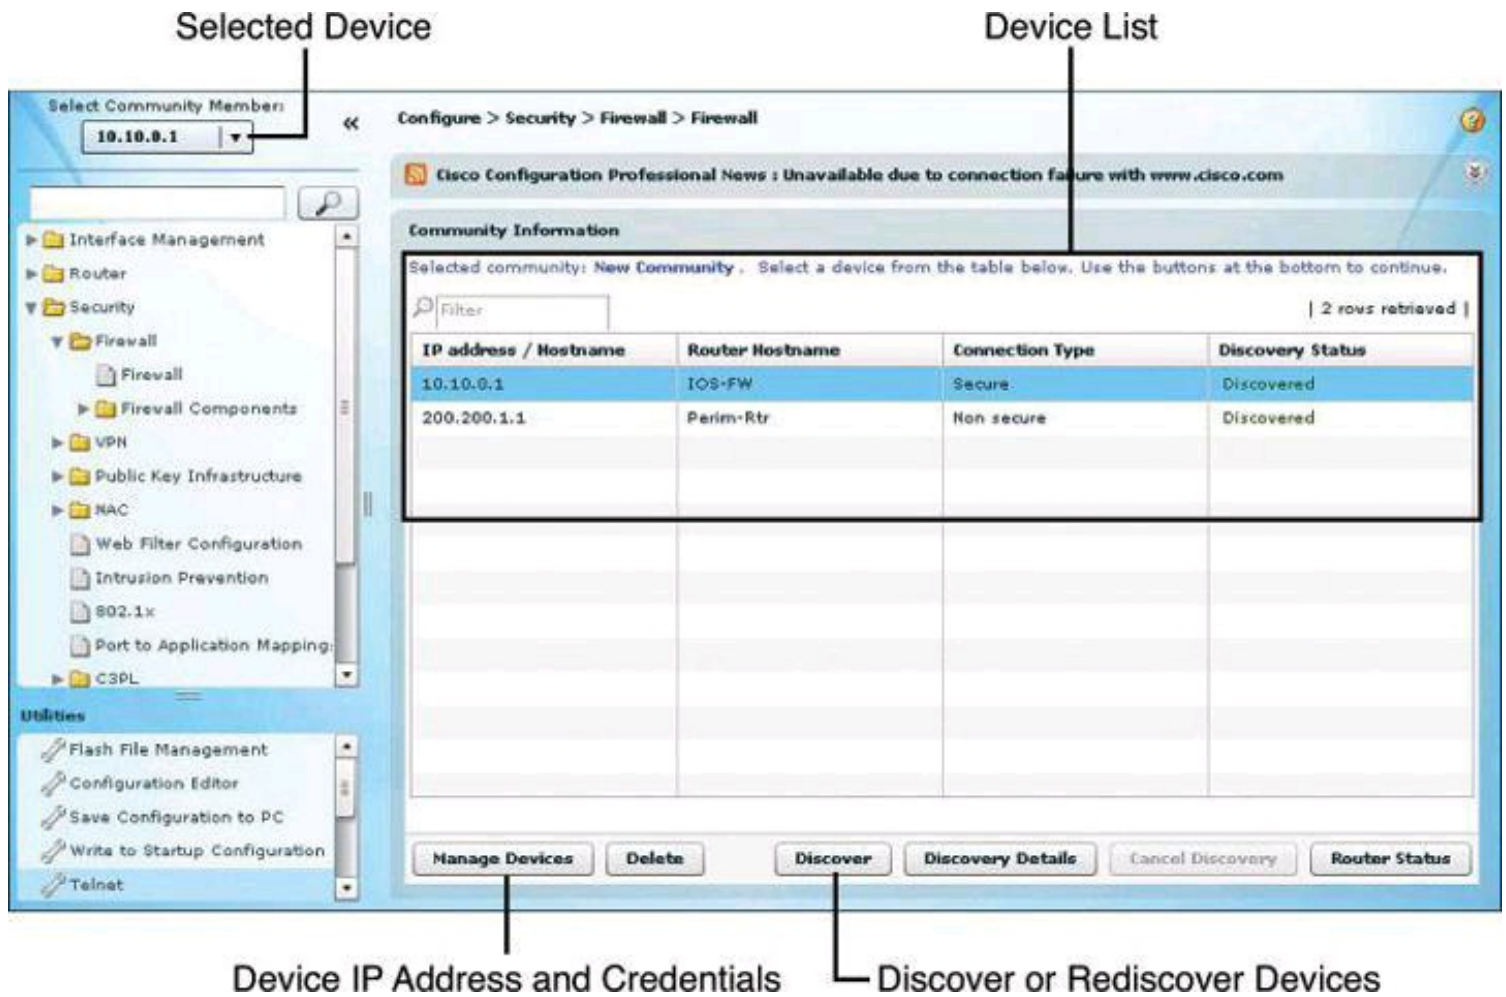


Figure 3-14. Managing Devices Within the Community

Templates

Templates facilitate management and consistency by attributing parameters to common components of configuration files (known as parameterizing). Using the template options, you can convert specific components of the configuration files into variables, which will acquire a different value for different devices that use the template.

After the template is created, it can be applied to multiple devices in different communities. The process includes options to provide values to the variables on the template, making them specific to each device. The template remains untouched. When applying the template, you can override the existing configuration, or merge the resulting configuration into the existing configuration.

[Figure 3-15](#) illustrates the template management options. The wizard-based process guides you through the following steps to create a template:

Step 1. Choose **Application > Template > Create**.

Step 2. Select the device from which you will obtain the template or select a configuration file from the local PC.

Step 3. Check the **Include This Configuration in the Template** check box.

Step 4. Find and select the attribute to be converted to variables, and click the **Parameterize**.

You can click Unparameterize to undo the action if the attribute is not to

become a variable. The selected attribute is wrapped within a double set of curly brackets `{{}}`. For instance, the word `demoone` becomes `{{demoone}}` after being parameterized.

Step 5. Click **Finish** and select the template filename and location.

Create variables that will acquire different values for different devices.

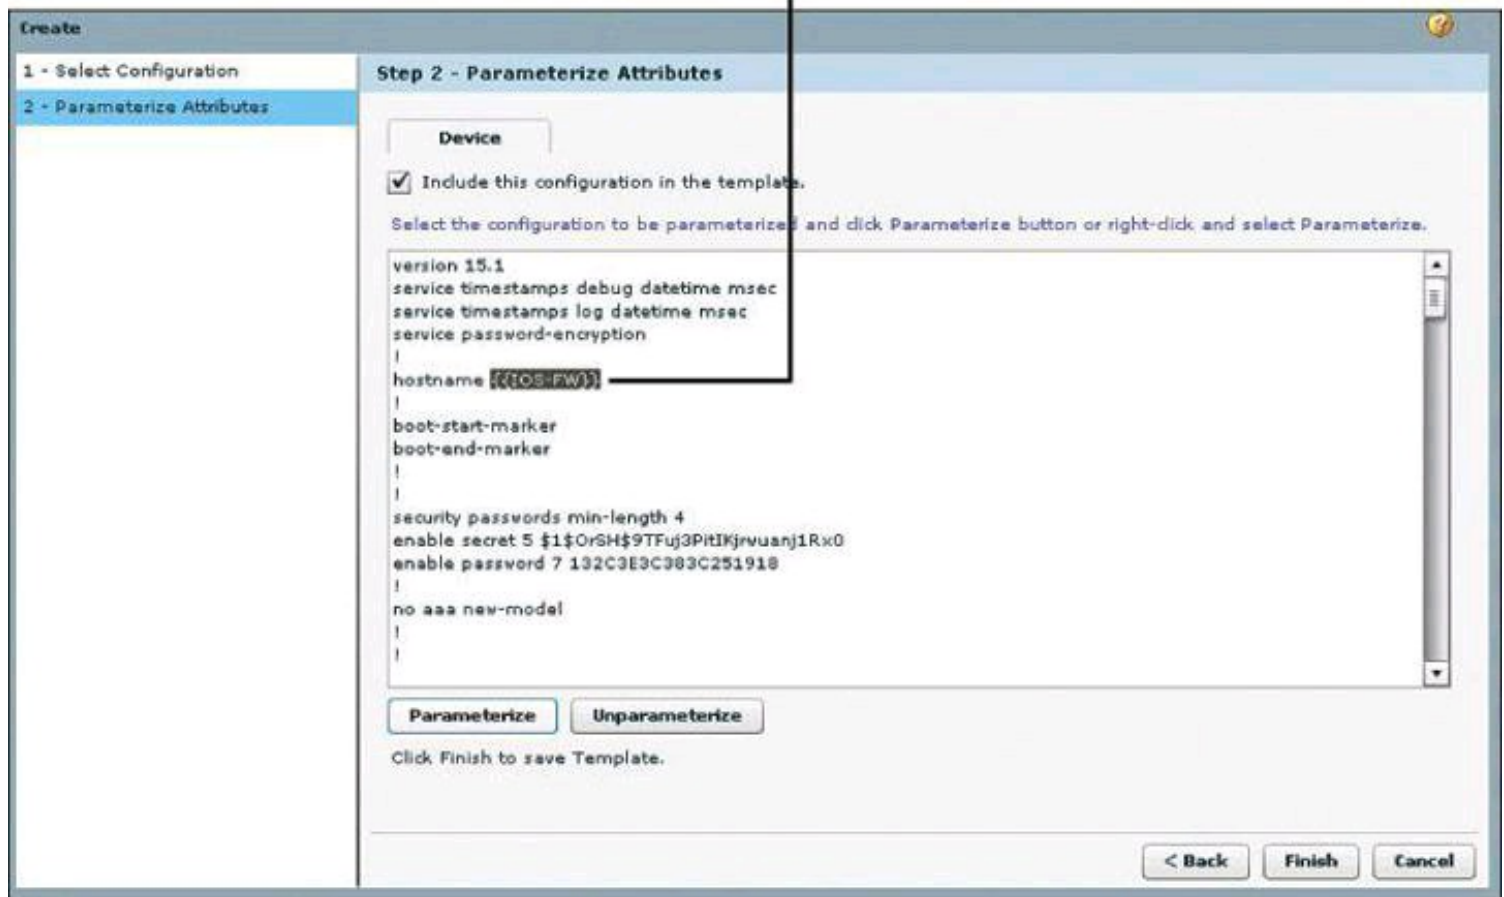


Figure 3-15. Creating Variables in CCP Templates

After the template is created, you can apply it to multiple devices by following these steps:

Step 1. Choose **Application > Template > Apply**.

Step 2. Click the browse button and navigate to select the template file.

Step 3. Click **Find Parameterized Attribute** to find and select specific text, and then change the value of the text to the value specific for the device.

Step 4. In the Apply to Device step, shown in [Figure 3-16](#), choose the router to apply the template from the Select Device drop-down list, and then click the respective radio button to select whether to merge with or override the running configuration. Check the **Enable Rollback** check box if rollback is desired.

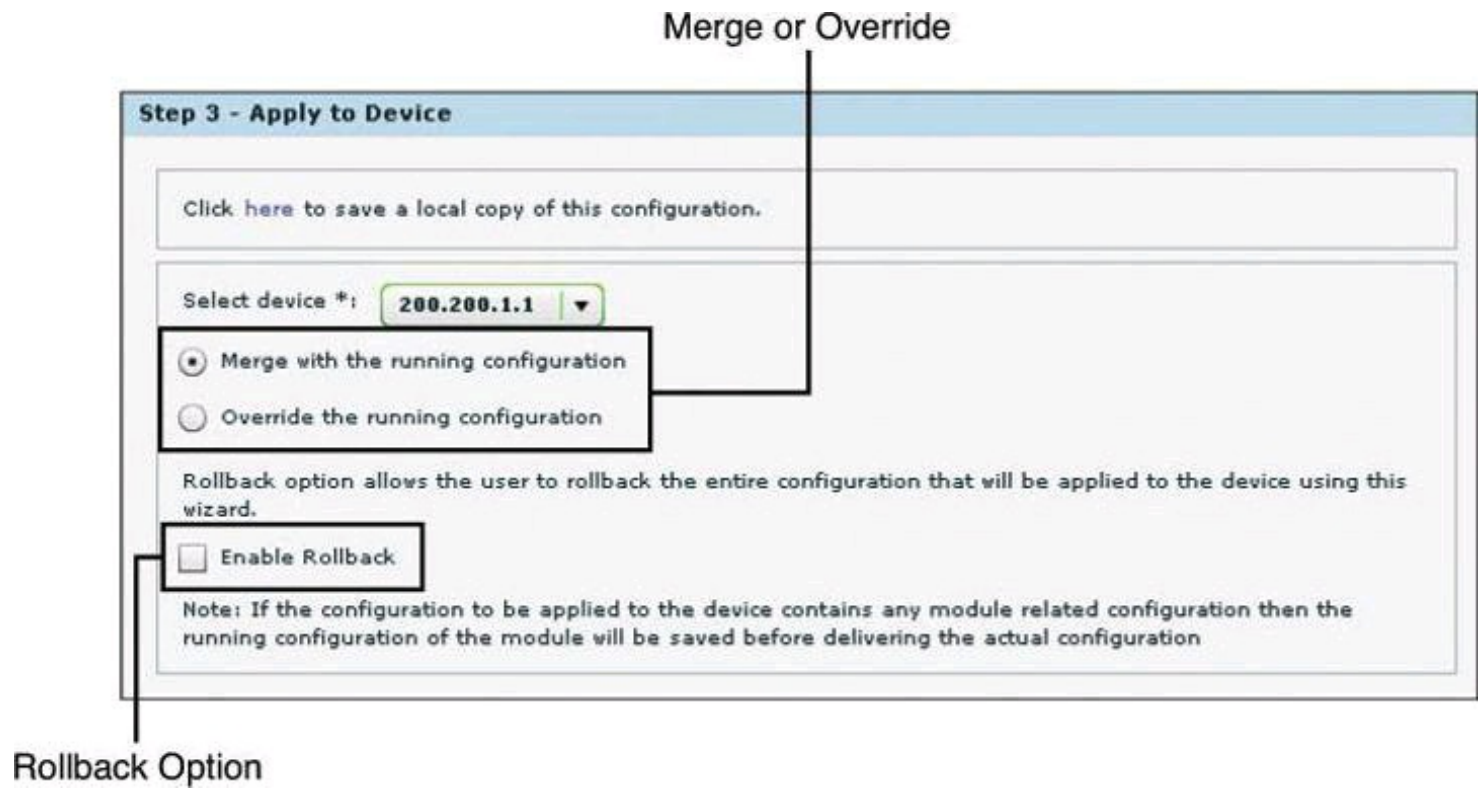


Figure 3-16. Applying a Template on a Discovered Device

Step 5. Click the Reload Device button to apply overridden configurations, or click Rollback to return to the previous configuration and discard the changes.

User Profiles

User profiles are a tool to define configuration views on CCP. Using profiles, administrators can define the view and configure which menus and screens are visible to the users who are acquiring the profile. This process effectively creates a different CCP option tree, where screens are available to profile users in full permission mode or in read-only mode.

[Figure 3-17](#) illustrates the profile management options. The wizard-based process guides you through the following steps to create a profile:

Step 1. Choose **Application > Create User Profile**.

Step 2. Click to select the device or devices that will be part of the user profile.

Step 3. The resulting folders represent the CCP option tree. Expand the folders to find the menus and screens.

Step 4. Select screens to set permissions.

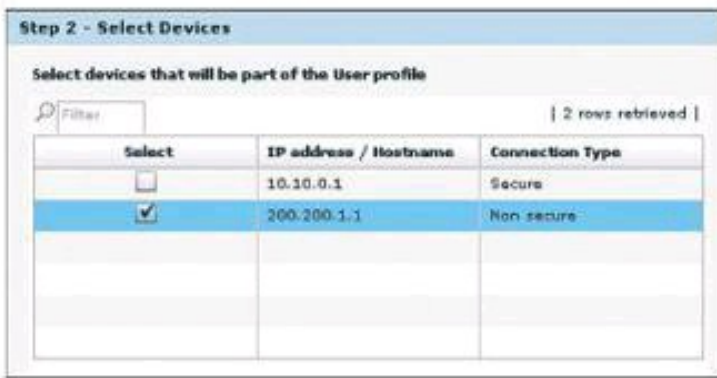
Step 5. Check the check boxes to define permissions. The following permissions are available:

- **Show this screen:** Makes the screen visible to users acquiring the profile, and defines full permission over the components of the screen.
- **Restrict to read operations only:** Changes the screen permissions to view-only or read-only.
- **Disabled:** This option does not appear explicitly, but can be accomplished by clearing the two check boxes for the previously listed permissions.

Step 6. Click **Save** to save the user profile as a file.

1. Select the device.

2. Define permissions on Cisco Configuration Professional menu options.



3. Verify the user view.

Full or Read-Only Permissions

Color Coding

This option is now disabled for this user.

Figure 3-17. Using User Profiles

Following this configuration, users will see options that are grayed out when they try to access screens that have been disabled for access, and they will see buttons and check boxes that are grayed out when they try to use options on screens that have been designated as read-only.

Using CCP to Harden Cisco IOS Devices

There are two options to harden your Cisco IOS devices, as shown in [Figure 3-18](#): the Security Audit or the One-Step Lockdown. The Security Audit approach is more granular than the One-Step Lockdown options. With the One-Step Lockdown approach, CCP evaluates and hardens your device without the option for you to review the changes about to be made.

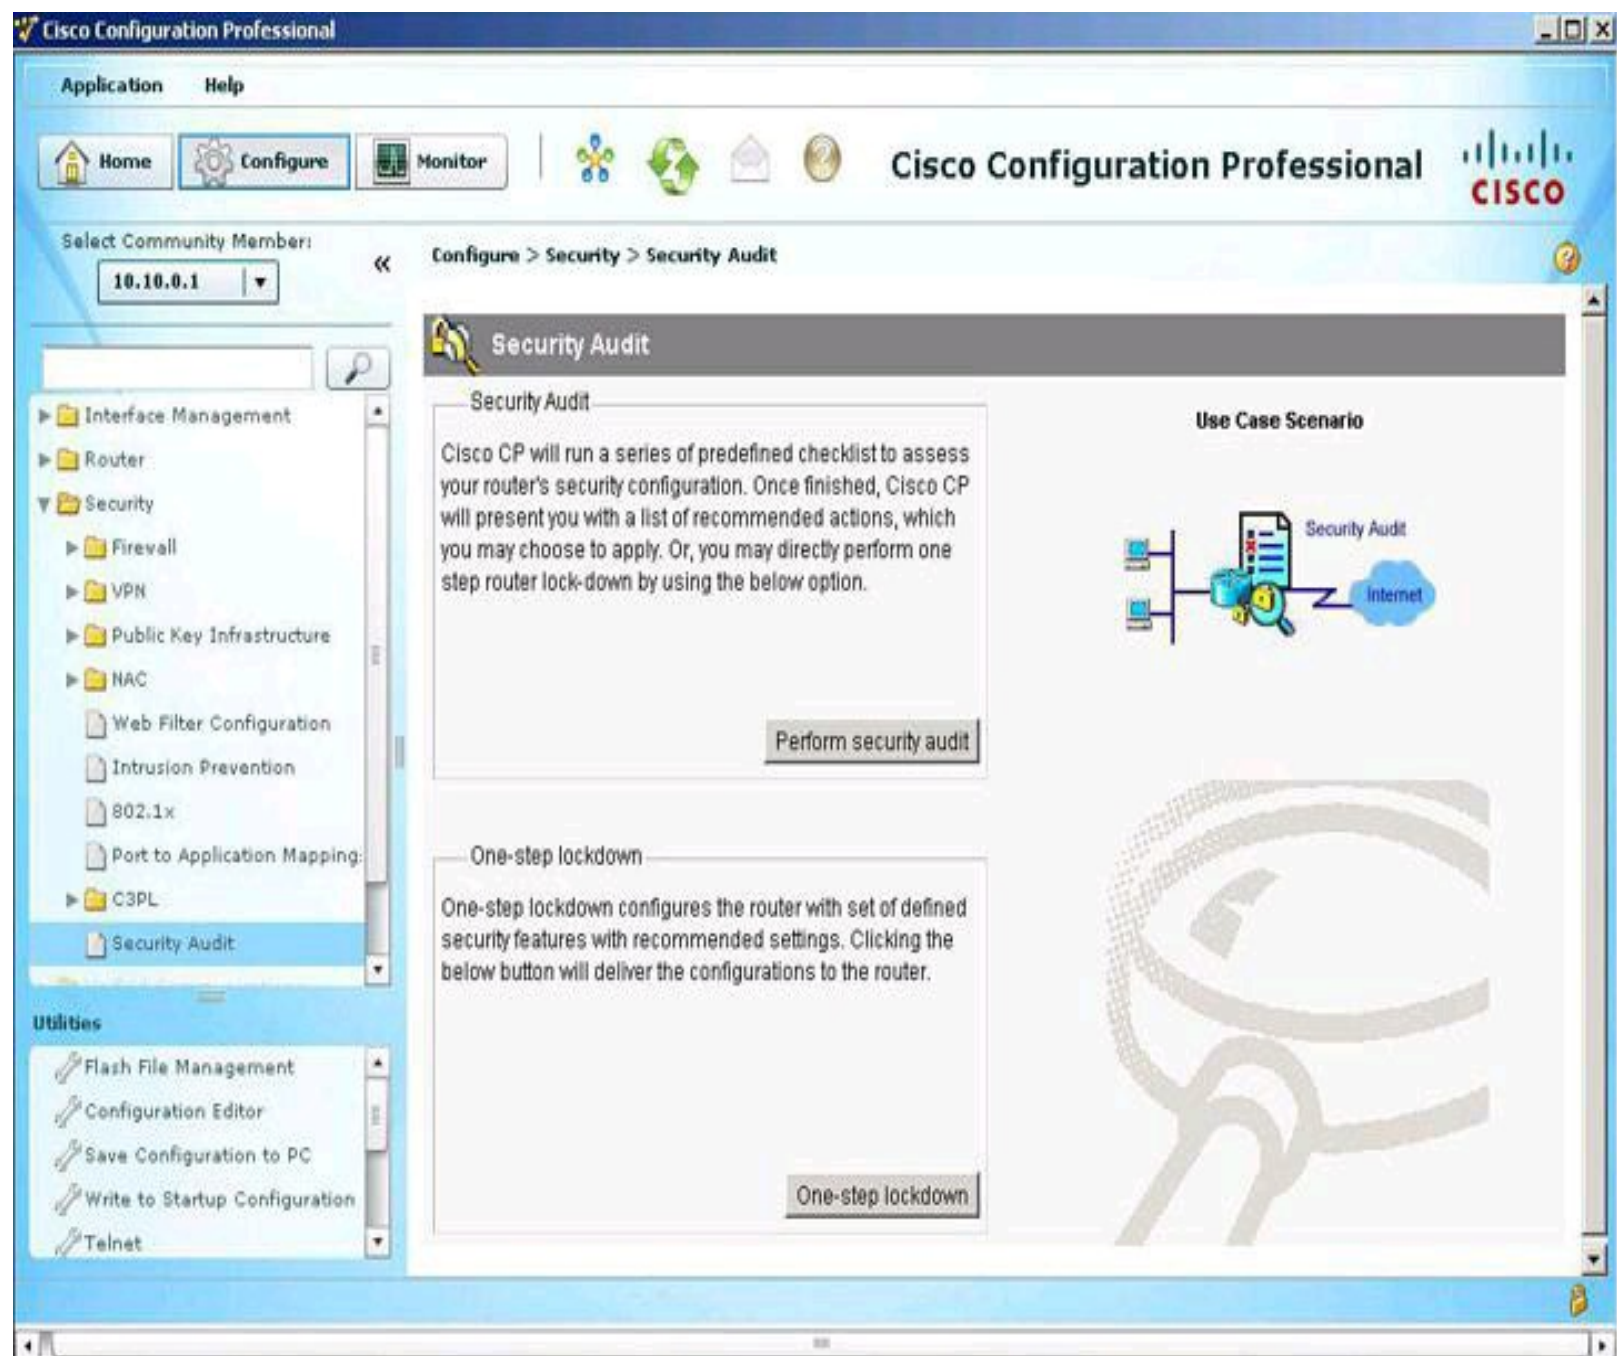


Figure 3-18. Security Audit Tools

Security Audit

The Security Audit Wizard implements device hardening by testing your router configuration to determine if any potential security problems exist in the configuration. The wizard then presents you with a screen that lets you determine which security problems you want to fix. Once determined, the Security Audit Wizard makes the necessary changes to the router configuration to fix those problems using One-Step Lockdown.

Following are the steps to perform the security audit:

Step 1. From the toolbar, choose **Configure > Security > Security Audit**.

Step 2. Click **Perform Security Audit**. The Welcome page of the Security Audit Wizard opens.

Step 3. Click **Next**. The Security Audit Interface Configuration page opens.

Step 4. Input information to tell the Security Audit Wizard which of your router interfaces connect to your inside network and which connect outside of your network. For each interface listed, check either Inside or Outside to indicate where the interface connects, as shown in [Figure 3-19](#).

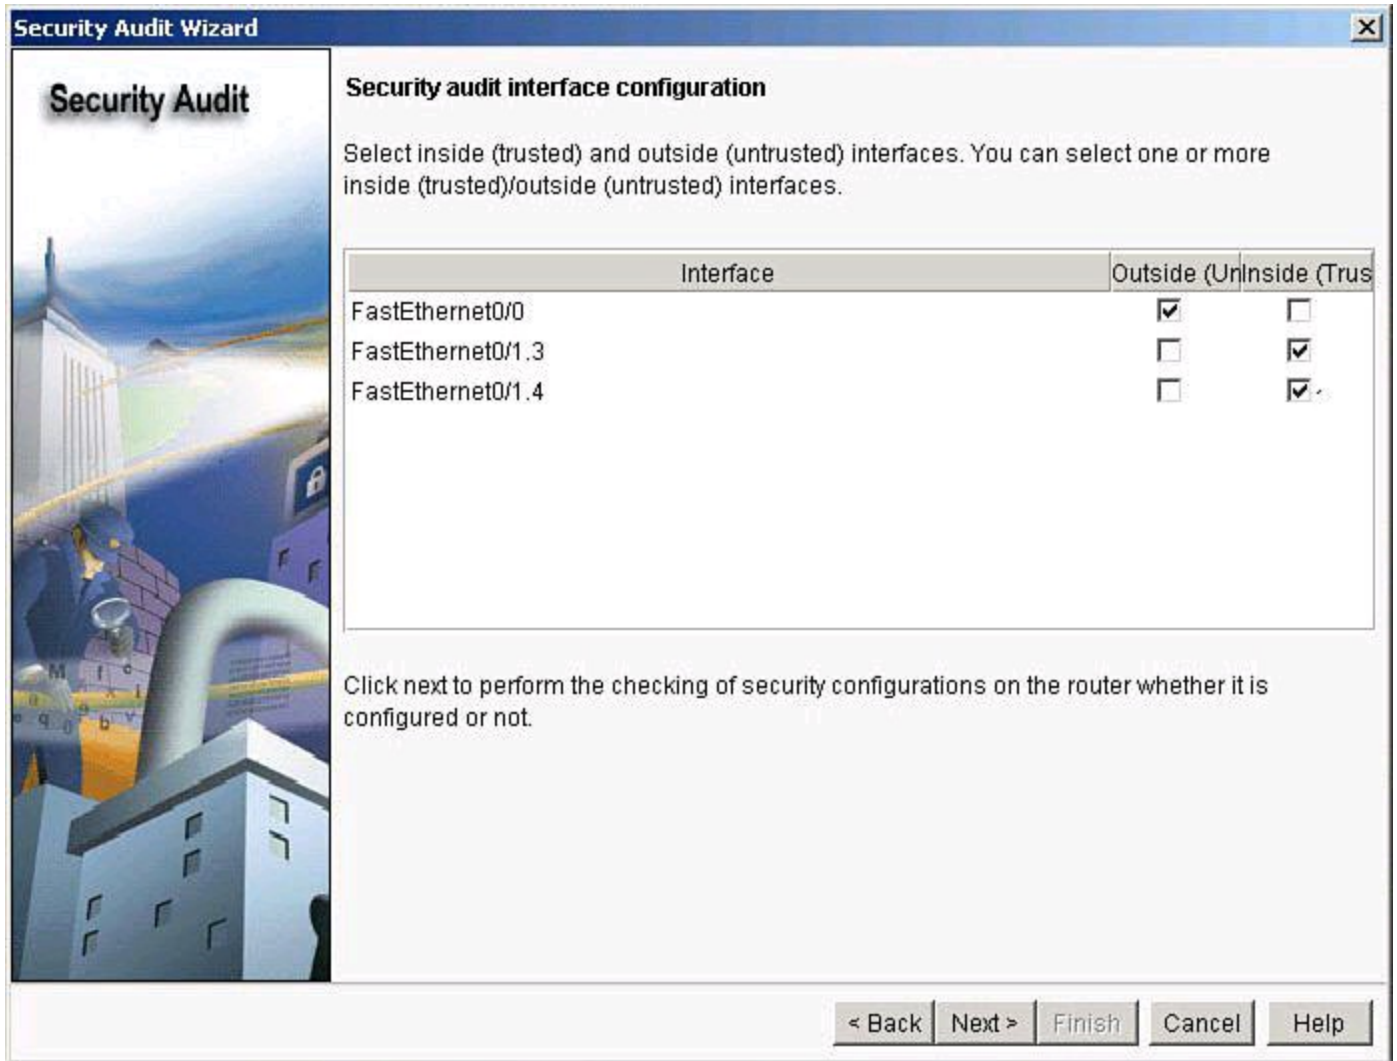


Figure 3-19. Selecting Interfaces for a Successful Security Audit

Step 5. Click **Next**. The Security Audit Wizard tests your router configuration to determine if possible security problems exist. The potential problems are determined by Cisco recommended practices. A screen showing the progress of this action opens, listing all of the configuration options being tested for, and whether the current router configuration passes those tests, as shown in [Figure 3-20](#). If you want to save this report to a file, click **Save Report**.

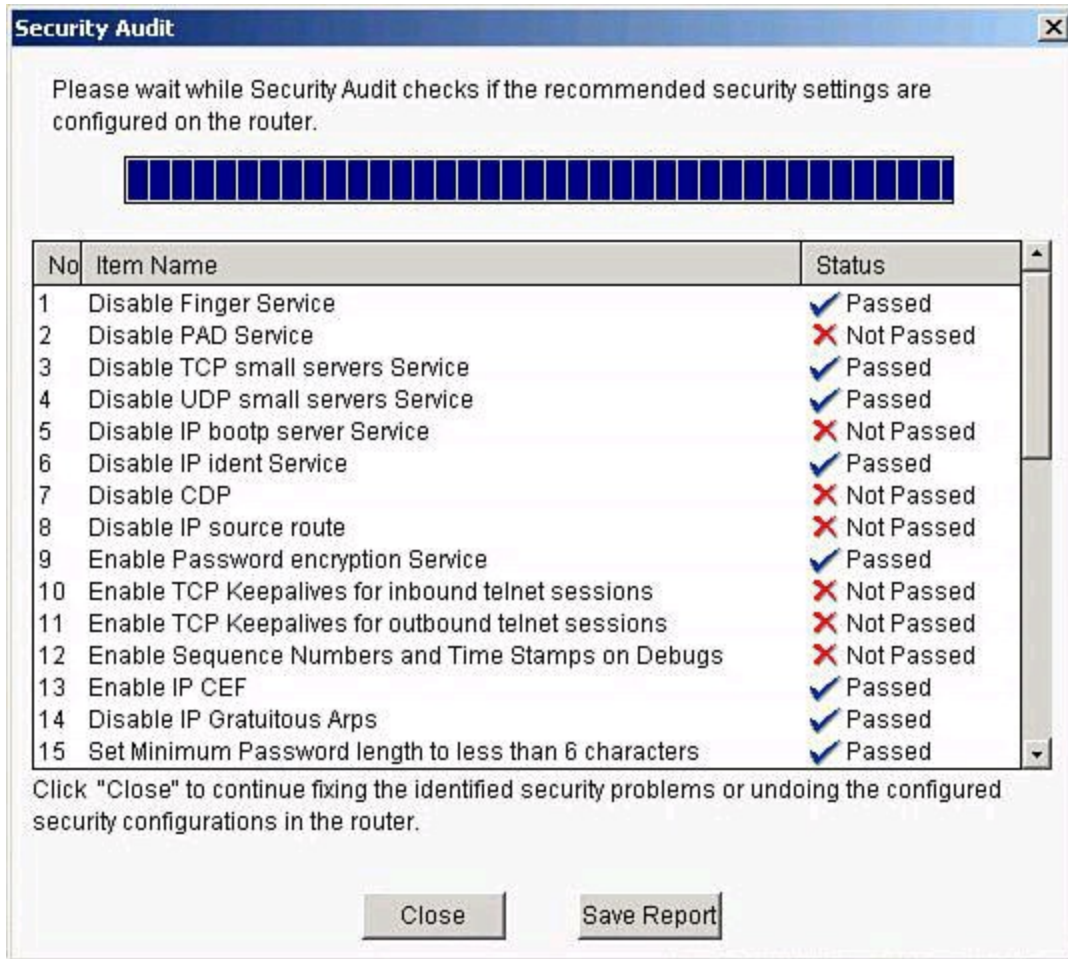


Figure 3-20. Security Audit Results

Step 6. Click **Close**. The Security Audit Wizard displays a “report card” showing a list of possible security problems, as shown in [Figure 3-21](#).

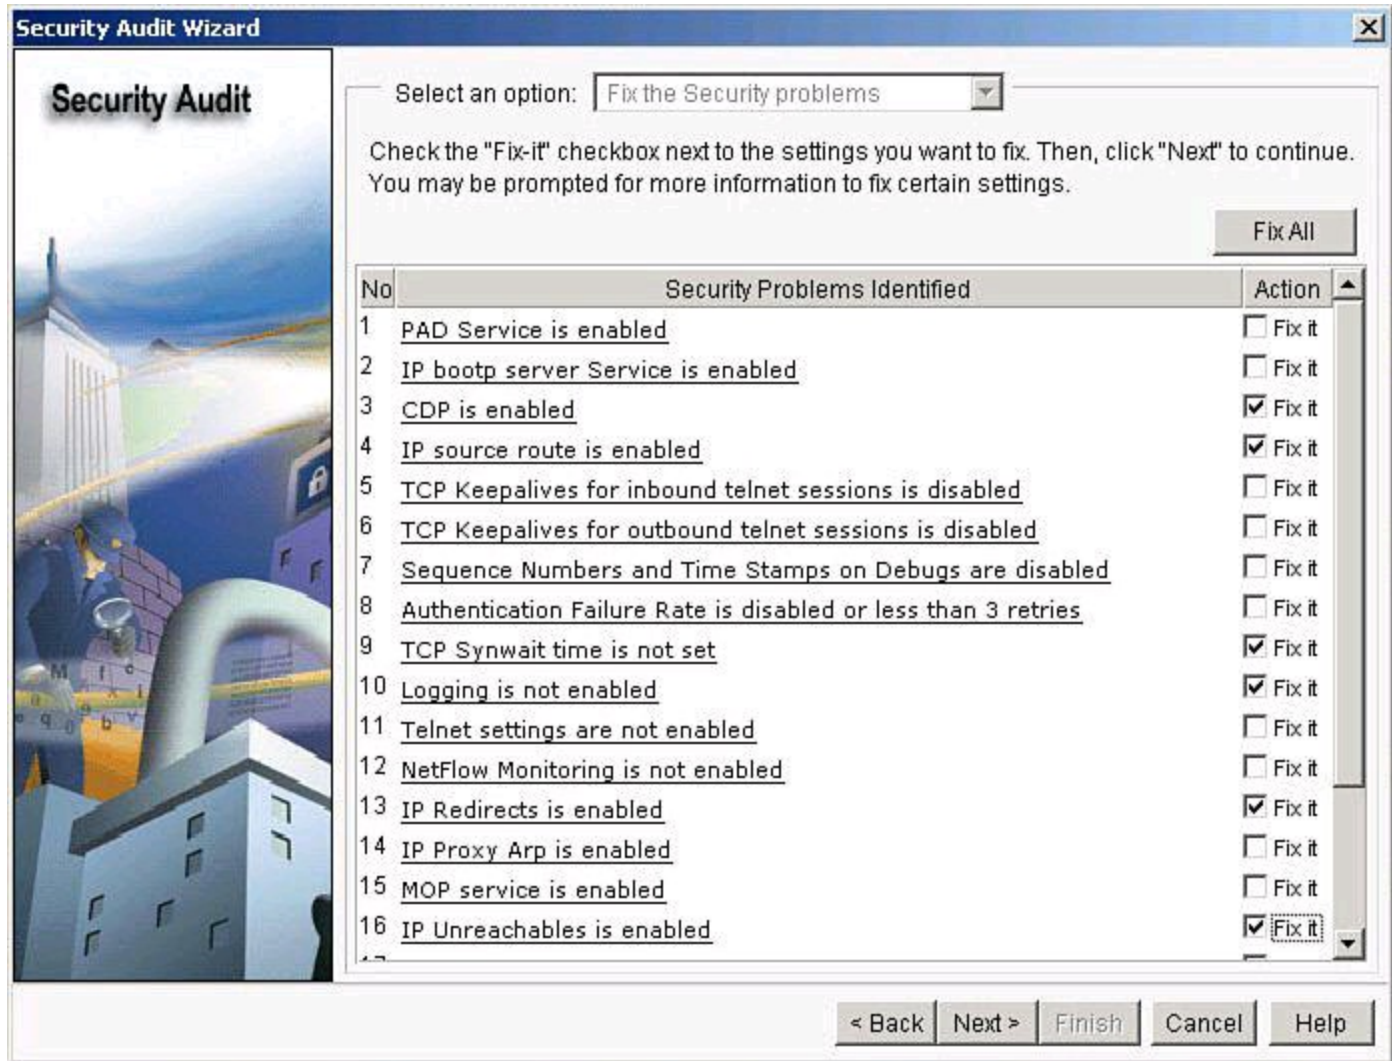


Figure 3-21. Addressing Issues Reported by the CCP Security Audit

Step 7. Check the **Fix It** check box next to any problems that you want CCP to fix. For a description of the problem and a list of the Cisco IOS commands that will be added to your configuration, click the problem description to display a help page about that problem.

Step 8. Click **Next**. The Security Audit Wizard may display one or more screens requiring you to enter information to fix certain problems. Enter the information as required and click **Next** for each of those screens.

Step 9. The Summary page of the wizard shows a list of all the configuration changes that Security Audit will make. Click **Finish** to implement those changes to your router.

CCP can undo this security fix. If you want CCP to remove this security configuration, run the Security Audit Wizard. In the report card window shown in [Figure 3-21](#), choose the option **Undo Security Configurations** from the drop-down list at the top of the window, and check the check box next to any configurations that you want to undo. Click **Next**.

One-Step Lockdown

The One-Step Lockdown option, available from the Security Audit main menu shown in [Figure 3-18](#), tests your router configuration for potential security problems and automatically makes necessary configuration changes to correct problems that are found. The configuration is evaluated against Cisco recommended practices, similar to the Security Audit Wizard. This time, however, changes are made

automatically and without prompting the user to select components to fix. Choose **Configure > Security > Security Audit > One-Step Lockdown**.

The One-Step Lockdown option streamlines device hardening by implementing the AutoSecure feature with a one-click option. To confirm the process, click the **Deliver** button in the confirmation dialog box, shown in [Figure 3-22](#), to deliver the new settings. Notice that the confirmation dialog box also explains how to roll back the new settings using the Security Audit feature. Also, notice that the wizard lists the settings as changed before offering the option to deliver them.

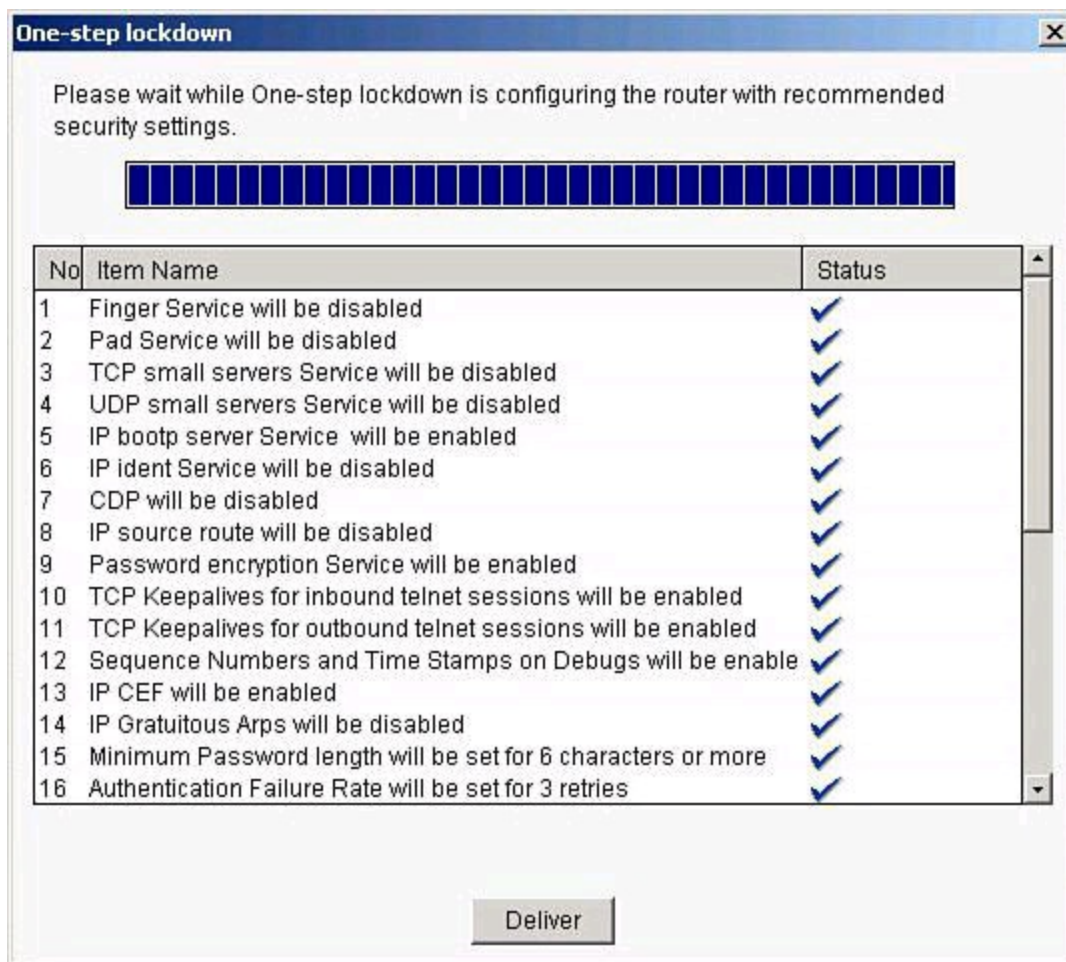


Figure 3-22. One-Step Lockdown

Cisco IOS AutoSecure

The following list provides examples of the Cisco AutoSecure features that are implemented in CCP during the One-Step Lockdown (refer to the CCP documentation on Cisco.com for an updated list of features):

- Disable unnecessary services and interfaces
 - Unused router interfaces
 - BOOTP server
 - Cisco Discovery Protocol
 - Configuration autoloading
 - FTP server
 - TFTP server

- NTP service
- PAD service
- TCP and UDP minor services
- DEC MOP service
- Disable commonly configured management services
 - SNMP
 - HTTP or HTTPS configuration and monitoring
 - DNS
- Ensure path integrity
 - Disable ICMP redirects
 - Disable IP source routing
 - Disable IP directed broadcast
 - Disable gratuitous and proxy ARP
 - Enable firewall on all outside interfaces
- Disable probes and scans
 - Finger
 - ICMP unreachable notifications
 - ICMP mask reply
- Ensure terminal access security
 - Enable SSH
 - Disable IP identification service
 - Enable TCP keepalives
 - Set access class on HTTP and vty
- Ensure authentication integrity
 - Enable password encryption
 - Set minimum password length
 - Set authentication failure rate

AutoSecure is a Cisco IOS feature that, like CCP, lets you more easily configure security features on your router, so that your network is better protected. CCP implements almost all of the configurations that AutoSecure provides.

The following AutoSecure features might not be implemented in CCP:

- **Disabling Network Time Protocol (NTP):** Based on input, AutoSecure will disable NTP if it is not necessary. Otherwise, NTP will be configured with Message Digest 5 (MD5) authentication. CCP does not support disabling NTP.
- **Configuring AAA:** If AAA is not configured, AutoSecure configures local AAA and prompts for configuration of a local username and password database on the router. CCP does not support AAA configuration.

- **Setting Selective Packet Discard (SPD) values:** CCP does not set SPD values.
- **Enabling TCP Intercept feature:** CCP does not enable TCP Intercept feature.
- **Configuring antispoofing ACLs on outside interfaces:** AutoSecure creates three named ACLs used to prevent antispoofing source addresses. CCP does not configure these ACLs.

CCP implements the following Cisco AutoSecure features differently:

- **Disable Simple Network Management Protocol (SNMP):** CCP will disable SNMP, but unlike AutoSecure, it does not provide an option for configuring SNMP version 3.
- **Enable SSH for access to the router:** CCP will enable and configure SSH on crypto Cisco IOS images, but unlike AutoSecure, it will not enable Service Control Point (SCP) or disable other access and file transfer services, such as FTP.

For the details on what is included with CCP, and its detailed functioning (such as whether admin rights are required), read the release notes for the specific version you are working with.

Summary

This chapter had two distinct parts. In the first part, we explored Cisco Network Foundation Protection and discussed, among other things:

- Threats that exploit network availability drive infrastructure protection.
- Cisco Network Foundation Protection implements a divide-and-conquer approach to infrastructure protection.
- Control plane security tools include CoPP and Cisco AutoSecure.
- Management plane security tools include RBAC and AAA services.
- Data plane security tools include ACLs, DHCP snooping, and Dynamic ARP spoofing.

In the second part of this chapter, we saw how Cisco Configuration Professional can be used to implement Cisco Network Foundation Protection, and how CCP provides a comprehensive and easy-to-use GUI tool for router configuration and service integration. Among other things, you learned:

- Features include security, unified communication, license management, application management, and others.
- Communities, templates, and user profiles provide role-based access control and streamline the process of configuring several devices.
- The Security Audit Wizard and the One-Step Lockdown tool are device hardening options.

References

For additional information, refer to these Cisco.com resources:

Cisco Configuration Professional User Guide (Version 2.5),
http://www.cisco.com/en/US/docs/net_mgmt/cisco_configuration_professional/v2_5/olh/cc
 “Cisco Network Foundation Protection (NFP),”
<http://www.cisco.com/en/US/products/ps6642/index.html>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions.](#)”

- 1.** Which device plane processes traffic generated or received by the network device itself, used for the creation and operation of the network, and managed by the CPU in the device route processor?
 - a.** Control plane
 - b.** Data plane
 - c.** Forwarding plane
 - d.** Management plane
- 2.** What is the main difference between Control Plane Policing (CoPP) and Control Plane Protection (CPPr)?
 - a.** No difference exists; both terms define the same functionality.
 - b.** CoPP is an extension of CPPr.
 - c.** CPPr is an extension of CoPP.
 - d.** CoPP relates to rate limiting, while CPPr defines QoS-like filters.
 - e.** CoPP allows more granularity on CPPr rules.
- 3.** Which of the following are components of Control Plane Protection? (Choose two.)
 - a.** Routing protocol authentication
 - b.** SNMP
 - c.** Cisco SecureX
 - d.** Cisco AutoSecure
 - e.** Zone-based firewall
- 4.** Which options are examples of data plane security controls? (Choose two.)
 - a.** RBAC
 - b.** CoPP
 - c.** Port security
 - d.** Spanning Tree Control
 - e.** Access control lists
 - f.** AAA
- 5.** Match the plane with the right traffic type.

Planes	
a. Management	
b. Data	
c. Control	
Traffic	
d. User-generated packets	

- e. Packets used for the creation and operation of the network itself
 - f. Packets used to manage the network
6. Which techniques could be used to mitigate a spoofing attack? (Choose two.)
- a. TTL filtering
 - b. Access class
 - c. Private VLANs
 - d. Access list
 - e. Management plane
 - f. Unicast Reverse Path Forwarding
7. Which management tool is available on the router's flash and can be used to configure some LAN and WAN interfaces of the router, as well as minimal Cisco IOS Software security features?
- a. Cisco AnyConnect
 - b. Cisco Configuration Professional Express
 - c. Cisco Secure Desktop
 - d. CiscoWorks
 - e. Cisco Security Manager
 - f. Cisco Configuration Professional
8. Match the Cisco Configuration Professional component with its definition.
- CCP components
- a. Communities
 - b. Templates
 - c. Profiles
 - d. Wizards
- Component definitions
- e. Parameterized configuration files
 - f. GUI views that allow role-based access control over Cisco Configuration Professional menus and options
 - g. Groups of devices that share common components
 - h. GUI tools to hide the complexity of commands
9. What does the Parameterize button do when creating Cisco Configuration Professional templates?
- a. Creates variables that acquire different values for different devices
 - b. Designates a parameter for use across several communities
 - c. Applies the same parameter to the configurations of all devices where the parameter is applied
 - d. Removes the special value of a given parameter

10. Which of the following is implemented by Cisco Configuration Professional when the One-Step Lockdown option is used?

- a.** Control Plane Protection
- b.** Security Wizard
- c.** AutoSecure
- d.** Packet Tracer

Chapter 4. Securing the Management Plane on Cisco IOS Devices and AAA

This chapter describes how to securely implement the management and reporting features of Cisco IOS devices. More precisely, it discusses the following:

- Technologies used in secure management and reporting, such as syslog, Network Time Protocol (NTP), Secure Shell (SSH), and Simple Network Management Protocol version 3 (SNMPv3).
- Proper password configuration, management, and password recovery procedures and how to safeguard a copy of the operating system and configuration file with the use of authentication, authorization, and accounting (AAA) both locally and on an external database.
- The use and configuration of Cisco Secure Access Control Server (ACS) as an external AAA database
- Secure management and reporting, as well as AAA, from both the command-line interface (CLI) and from Cisco Configuration Professional (CCP).

We all know that secure management of our equipment starts with physical security. Physical controls were discussed in [Chapter 1](#), “[Network Security Concepts and Policies](#).” In this chapter, therefore, we will focus on the secure remote access of our equipment, and proper reporting of events. We will also discuss methods to manage an administrator’s credentials locally or on a centralized server.

Configuring Secure Administration Access

As discussed in [Chapter 3](#), “[Network Foundation Protection and Cisco Configuration Professional](#),” local access to a router usually involves a direct connection to a console port on the Cisco router.

Remote access typically involves allowing Telnet, Secure Shell (SSH), HTTP, HTTPS, or Simple Network Management Protocol (SNMP) connections to the Cisco IOS device from a computer on the same subnet or a different subnet. As an example, in [Chapter 3](#), we saw how to remotely access our Cisco IOS router using CCP.

Note

Though SNMP is discussed sparsely in this section, it will be presented in greater detail later in this chapter in the section “[Implementing Secure Management and Reporting](#).”

It is preferable to allow only local access to the Cisco IOS device because some remote-access protocols, such as Telnet, send the data, including usernames and passwords, to the network device in plaintext.

If remote access is required, it is recommended that you apply one of the following options:

- Establish a dedicated management network as shown in [Figure 4-1](#). The management

network should include only identified administration hosts and connections to a dedicated interface on the router.

- Encrypt all the traffic between the administrator computer and the router.

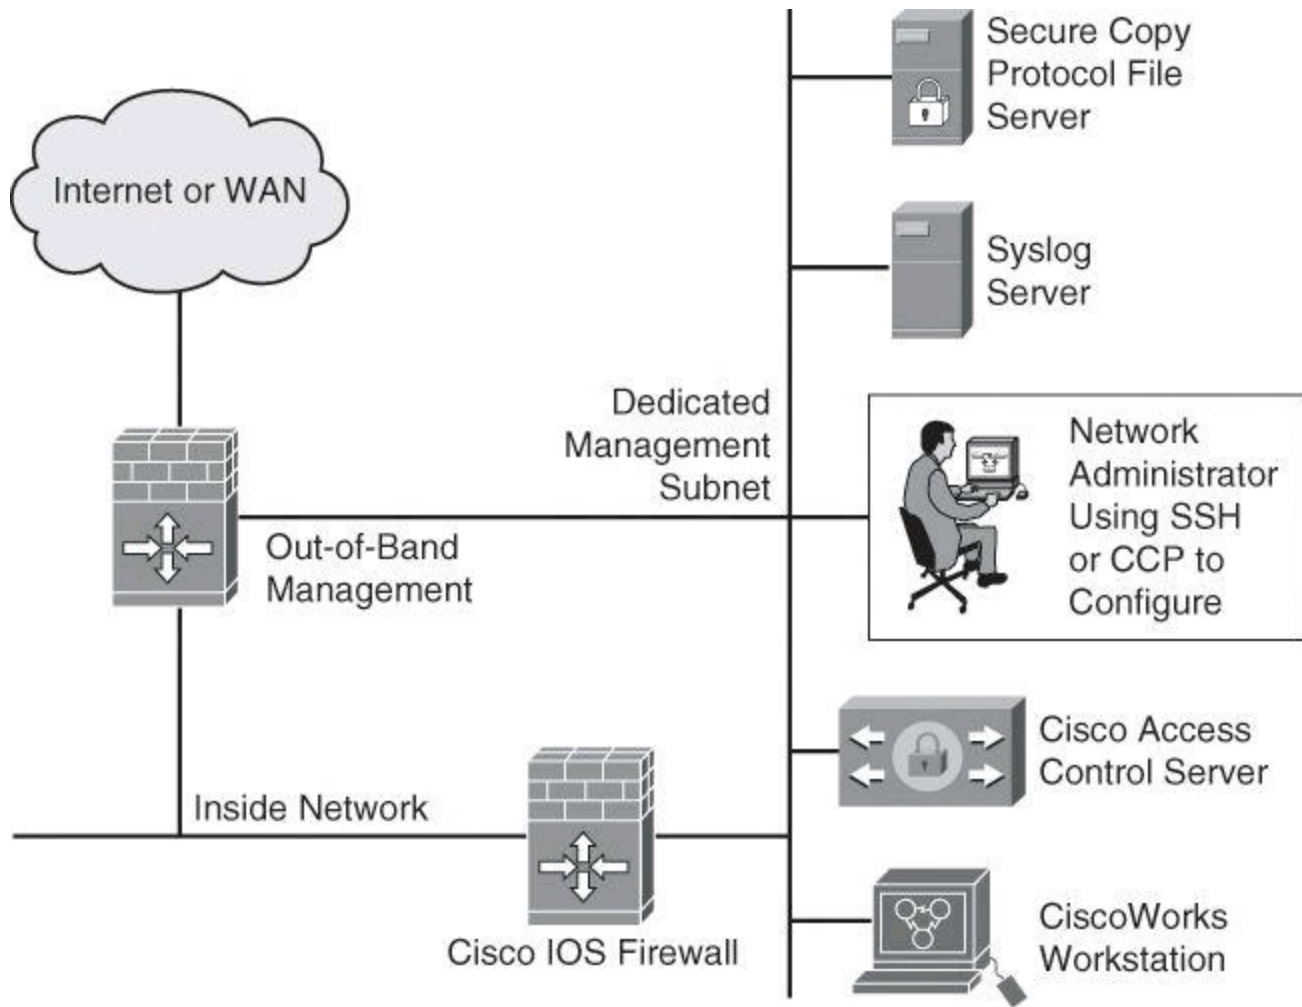


Figure 4-1. Dedicated Management Network

In either case, you can configure a packet filter to allow only the identified administration hosts and preferred protocols to access the router. For example, you can configure the packet filter to permit only SSH requests from the IP address of the administration host to initiate a connection to the routers in the network.

Configuring an SSH Daemon for Secure Management Access

SSH was discussed in [Chapter 3](#), but it's a topic worth revisiting briefly because it's part of providing secure management access to the Cisco router. The SSH daemon is a feature that enables an SSH client to make a secure, encrypted connection to a Cisco router. This connection provides functionality similar to that of an inbound Telnet connection, but it also provides strong encryption to be used with local authentication, thus ensuring secure access and communication to network devices such as routers, switches, firewalls, intrusion prevention system (IPS) sensors, and so forth. While Telnet authentication can be as simple as presenting the right password, SSH always requires presenting a username and a password. The SSH daemon in Cisco IOS Software works with publicly and commercially available SSH clients, such as PuTTY, OpenSSH, or Tera Term. To refresh your knowledge of SSH configuration, here are the steps to configure your Cisco router to support an SSH daemon using the CLI:

Step 1. Configure the IP domain name of your network using the **ip domain-name** *domain-name* command in global configuration mode:

```
Router(config)# ip domain-name cisco.com
```

Note

The domain name and the hostname are used for the generation of the RSA key pairs. If there are any existing key pairs, it is recommended that you overwrite them using the command **crypto key zeroize rsa**

Step 2. Generate keys to be used with SSH by generating the Rivest, Shamir, and Adleman (RSA) keys using the **crypto key generate rsa general-keys modulus** *modulus-size* command in global configuration mode. The modulus determines the size of the RSA key. The larger the modulus, the more secure the RSA key. However, keys with large modulus values take longer to generate, and encryption and decryption operations take longer with larger keys:

```
Router(config)# crypto key generate rsa general-keys  
modulus 1024
```

Note

The minimum recommended key length is modulus 1024. However, as per NIST SP800-313A, a modulus smaller than 2048 bits will be disallowed for federal use after 2013; see <http://csrc.nist.gov/publications/nistpubs/800-131A/sp800-131A.pdf>.

Step 3. Optionally, to display the generated keys, use the **show crypto key mypubkey rsa** command in privileged EXEC mode.

Step 4. Configure the time that the router waits for the SSH client to respond using the **ip ssh time-out** *seconds* command in global configuration mode:

```
Router(config)# ip ssh time-out 120
```

Step 5. Configure the SSH retries (the number of attempts after which the interface will reset) using the **ip ssh authentication-retries** *integer* command in global configuration mode:

```
Router(config)# ip ssh authentication-retries 4
```

Step 6. Enable vty inbound SSH sessions:

```
Router(config)# line vty 0 4  
Router(config-line)# transport input ssh
```

The SSH protocol is automatically enabled when you generate the SSH (RSA) keys. Once the keys are created, you can access the router SSH daemon using your SSH client software.

Tip

If you are using a version of Cisco IOS Software that supports both SSHv1 and SSHv2, by default SSH runs in compatibility mode; that is, both SSHv1 and SSHv2 connections are honored. If you are running Cisco IOS Release 12.3(4)T or later, you can use the **ip ssh version {1 | 2}** command to configure support for only one version of SSH.

The procedure for connecting to a Cisco router SSH daemon varies depending on the SSH client application that you use. Generally, the SSH client passes your username to the router SSH daemon. The router SSH daemon prompts you for the correct password. After the password has been verified, you can configure and manage the router as if you were a standard vty user.

Tip

Cisco routers with Cisco IOS Software Releases 12.1(3)T and later can act as both SSH clients and SSH daemons. This means that you could initiate an SSH client-to-server session from your router to a central SSH daemon system using the **ssh** command. SSH employs strong encryption to protect the SSH client-to-server session. Unlike Telnet, where anyone with a sniffer can see exactly what you are sending to and receiving from your routers, SSH encrypts the entire session.

Many vulnerabilities have been reported for SSH Version 1, such as the following root exploit: <http://www.doecirc.energy.gov/techbull/archive/CIRCTech02-001.shtml>. It is therefore recommended to use SSH Version 2.

Configuring Passwords on Cisco IOS Devices

Strong passwords and similar secrets, such as SNMP community strings, are the primary defense against unauthorized access to your Cisco IOS devices. The first step to secure a device's administrative access is to configure secure system passwords. The best scalable way to handle most passwords is to maintain them on a TACACS+ or RADIUS authentication server, such as the Cisco Secure Access Control Server (ACS) and Cisco Identity Services Engine (ISE). (TACACS+, RADIUS, and Cisco Secure ACS will be covered later in this chapter.) However, routers can have locally configured passwords for privileged access and can have other password information in their configuration files. This section focuses only on configuring local passwords.

The first step to secure Cisco router administrative access is to configure secure system passwords.



Key
Topic

A password can be established on individual lines, such as the console and vty, and for the privileged EXEC mode. Passwords are case sensitive.

By default, the console port does not require a password for console administrative access; however, you should always configure a console port line-level password. As shown in [Example 4-1](#), you can use the **line console 0** command followed by the **login** and **password** subcommands to require login and establish a login password on the console line.

By default, Cisco routers support up to five simultaneous vty (Telnet) sessions. On the router, the vty ports are numbered from 0 through 4. [Example 4-1](#) shows how you can use the **line vty 0 4** command followed by the **login** and **password** subcommands to log in to Telnet sessions. However, in [Chapter 3](#) you were introduced to a better practice using the **login local** command, where an administrator was authenticated against a local database with the global **username** command.

Example 4-1. Configuring the Console and Virtual Terminal Passwords

[Click here to view code image](#)

```
R1 (config)# line console 0
R1 (config-line)# login
R1 (config-line)# password M3rcury$12
R1 (config-line)# exit
R1 (config)# line vty 0 4
R1 (config-line)# login
R1 (config-line)# password v3nus$2012
```

The **enable secret password** global command restricts access to the privileged EXEC mode. You can use the **enable secret** global configuration command to configure the enable secret password. The enable secret password is always hashed inside the router configuration using a Message Digest 5 (MD5) hashing algorithm; it never appears in cleartext. Hashing will be covered in detail in [Chapter 12](#), “[Fundamentals of Cryptography and VPN](#).”

The enable password is also used to enter enable mode, but this weak password configuration is from earlier versions of Cisco IOS Software. By default, the enable password is not encrypted in the router configuration. The **enable password** global configuration command was kept for backward compatibility in case you downgrade the router to a version of Cisco IOS Software that does not support the **enable secret** command. If both **enable password** and **enable secret** are configured, the **enable password password** is ignored and only the **enable secret password** is used.

If you forget the enable password, regardless of whether it was created using the **enable secret** command or the **enable password** command, you have no alternative but to replace it using the Cisco router password recovery procedure specific to your Cisco equipment, which you can find in Cisco Document ID 6130 at http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_tech_note09186a00801746e6.s

With the exception of the enable secret password, all Cisco router passwords are stored in plaintext by default within the router configuration. You can view these passwords with the **show running-config** command. Sniffers can also see these passwords if your TFTP server configuration files traverse an unsecured intranet or Internet connection. If an intruder gains access to the TFTP server where the router configuration files are stored, the intruder is able to obtain these passwords.

As a safeguard against this possible exploit, the **service password-encryption** command encrypts all the passwords (except the previously hashed enable secret password) in the device configuration file, and will encrypt any passwords you set after entering this command until you turn the command off with the **no** form of the command. This method, which uses the Vigenère method (also explained in [Chapter 12](#)), is not as safe as MD5, which is used with the **enable secret** command, but prevents

casual discovery of the device line-level passwords. To configure password encryption, use the **service password-encryption** global command. To remove the **service password-encryption** command, use **no service password-encryption**.

Also by default, Cisco router auxiliary ports do not require a password for remote administrative access. Administrators sometimes use this port to remotely configure and monitor the router using a dialup modem connection. To combat this vulnerability, you can use the **line aux 0** command followed by the **login** and **password** subcommands to require login and establish a login password on an incoming auxiliary line.

Note

If you want to turn off the EXEC process for a specific line, such as on the auxiliary port, use the **no exec** command within the line configuration mode.

Setting Timeouts for Router Lines

Another topic important to the security of the management of the device is the line timeouts. As you learned during your CCNA studies, by default, an administrative interface stays active (and logged in) for 10 minutes after the last session activity, and you should adjust timers using the **exec-timeout** command in line configuration mode for each of the line types used. You can also use Cisco Configuration Professional, introduced in [Chapter 3](#), to configure the exec-timeout for the vty lines.

Configuring the Minimum Length for Router Passwords

Cisco IOS Software Release 12.3(1) and later allow you to set the minimum character length for all router passwords by using the **security passwords** global configuration command. This command provides enhanced security access to the router by allowing you to specify a minimum password length (0 to 16 characters); this eliminates common passwords that are short and prevalent on most networks, such as lab and cisco. [Example 4-2](#) demonstrates the **security passwords** command set for a minimum of ten characters.

Example 4-2. *security passwords* Command

[Click here to view code image](#)

```
R1(config)# security passwords min-length 10
```

This command affects username passwords, enable passwords and enable secret passwords, and line passwords that are created after the command is executed. Existing router passwords remain unaffected.

When creating passwords for Cisco routers, always keep the following rules in mind:

- Establish a minimum of ten characters for a password.
- Passwords can include the following:
 - Any alphanumeric character, but is case sensitive

- A mix of uppercase and lowercase characters
- Symbols and spaces
- Passwords should not use dictionary words.
- Password-leading spaces are ignored, but no spaces after the first character are ignored.
- Decide when and how often the passwords should be changed.

After the **security passwords** command has been enabled, any attempt to create a new password that is less than the specified length fails and results in an error message similar to this message:

```
Password too short - must be at least 10 characters. Password
configuration failed.
```

Enhanced Username Password Security

As discussed briefly in [Chapter 3](#), Cisco routers can maintain a list of usernames and passwords for performing local login authentication. System administrators can choose to use an MD5 hashing mechanism to encrypt a user password. MD5 hashing of passwords is a much better algorithm than the standard type 7 found in the **service password-encryption** command.

MD5 hashing of a Cisco IOS user password is accomplished with the **username secret** command in global configuration mode. Administrators can choose to enter a plaintext password for MD5 hashing by the router (option **0**), or they can enter a previously encrypted MD5 secret (option **5**). The syntax for the **username secret** command is as follows:

```
username name secret { [0] password | 5 encrypted-secret }
```

[Table 4-1](#) shows the parameters of the **username secret** command.

Table 4-1. *username secret* Parameters

Parameter	Description
<i>name</i>	This parameter specifies the username.
0	(Optional) This option indicates that the plaintext password is to be hashed by the router using MD5.
<i>password</i>	This parameter is the plaintext password to be hashed using MD5.
5	This parameter indicates that the encrypted-secret password was hashed using MD5.
<i>encrypted-secret</i>	This parameter is the MD5 encrypted-secret password that is stored as the encrypted user password.

[Example 4-3](#) shows an example of the **username secret** command.

Example 4-3. *username secret* Command

[Click here to view code image](#)

```
R1 (config) # username SecAdmin secret 0 Curium2012
R1 (config) # username SecAdmin secret 5 $1$uyyB$vAWRWP.qFQqb65.KxVxKg1
```

Dissecting the Hashed Password Format

In [Example 4-3](#), **secret 5** means that the password is not in cleartext. The password was hashed with MD5. The hashed password, \$1\$uyyB\$vAWRWP.qFQqb65.KxVxKg1, breaks down as follows:

- 1 appearing between the first dollar sign and second dollar sign indicates that the password was hashed with MD5.
 - uyyB appearing between the second and third dollar signs represents the randomized salt phrase added to the enable secret password prior to hashing.
 - A salt value, typically 48 to 128 bits long and randomly generated, helps to produce a hash result that will not match against a lookup table of hashed dictionary words, because dictionary words have fewer bits.
 - The salt phrase is stored in your configuration file, and the Cisco IOS router can use it along with the enable password you type to generate an MD5 hash value. If the value calculated from using the salt and the password you just typed in the MD5 algorithm matches the hash result stored in the configuration, such as vAWRWP.qFQqb65.KxVxKg1 in [Example 4-3](#), the router knows that the password that was typed is legitimate.
-

Securing ROM Monitor

By default, Cisco IOS routers allow a break sequence during startup that forces the router into ROM monitor mode, which can be used for a password recovery procedure. This procedure, if performed correctly, leaves the router configuration intact. Anyone who gains physical access to the router console port can enter ROM monitor mode, reset the enable secret password, and discover the router configuration. You can mitigate this potential security breach by using the **no service password-recovery** global configuration command, as shown in [Example 4-4](#). This command is a hidden Cisco IOS command and has no arguments or keywords.

Example 4-4. *no service password-recovery* Command

[Click here to view code image](#)

```
R1 (config) # no service password-recovery
WARNING:
Executing this command will disable password recovery mechanism. Do not
execute this
    command without another plan for password recovery.
Are you sure you want to continue? [yes/no]: yes
R1 (config) #
```

Note

To recover a device after the **no service password-recovery** command has been entered,

press the **Break** key within 5 seconds after the image decompresses during the boot. You are prompted to confirm the Break key action. When you confirm the action, the startup configuration is erased, the password recovery procedure is enabled, and the router boots with the factory default configuration. If you do not confirm the Break key action, the router boots normally with the No Service Password-Recovery feature enabled.

Securing the Cisco IOS Image and Configuration Files

Router files are critical to the operations of the router. Operating system images and configuration files are stored in the file system in the router's flash and nonvolatile memory. The Cisco IOS Resilient Configuration feature enables a router to secure and maintain a working copy of the running image and configuration so that those files can withstand malicious attempts to erase the contents of persistent storage (NVRAM and flash storage).

A great challenge for network operators is the total downtime that is experienced after a router has been compromised and its operating software and configuration data are erased from its persistent storage. The operator must retrieve an archived copy (hopefully one is available) of the configuration and a working Cisco IOS image to restore the router. Recovery must then be performed for each affected router, adding to the total network downtime.

The Cisco IOS Resilient Configuration feature is intended to speed up the recovery process. This feature maintains a secure working copy of the router image and the startup configuration at all times. The user cannot remove these secure files. This set of Cisco IOS image and router running configuration files is referred to as the *bootset*.

[Table 4-2](#) describes the key commands that are required to secure the Cisco IOS image and running configuration using the **secure boot-image** command:

```
R1 (config) # secure boot-image  
R1 (config) # secure boot-config
```

Table 4-2. *secure* Commands

Command	Description
<code>secure boot-image</code>	This command enables Cisco IOS image resilience. When turned on for the first time, the running image (as displayed in the <code>show version</code> command output) is secured, and a syslog entry is generated. This command functions properly only when the system is configured to run an image from a disk with an Advanced Technology Attachment (ATA) interface. Images that are booted from a TFTP server cannot be secured. Because this command has the effect of “hiding” the running image, the image file is not included in any directory listing of the disk. If the router is configured to boot with Cisco IOS resilience and an image with a different version of Cisco IOS is detected, a message similar to this is displayed at boot: “ios resilience :Archived image and configuration version 12.2 differs from running version 12.3.”
<code>secure boot-config</code>	This command takes a snapshot of the router running configuration and securely archives it in persistent storage.

Secured files do not appear in the output of a `dir` command that is issued from an executive shell because the Cisco IOS file system prevents the secure files in a directory from being listed. ROM monitor mode does not have any such restriction and can list and boot secured files. Because the running image and running configuration archives are not visible in the output from the Cisco IOS command `dir`, use the `show secure bootset` command to verify the archive existence.

[Example 4-5](#) shows an example of the `show secure bootset` command output. This command is important to verify that the Cisco IOS image and configuration files have been properly backed up and secured.

Example 4-5. `show secure bootset` Command Output

[Click here to view code image](#)

```
R1# show secure bootset
IOS resilience router id FHK085031MD

IOS image resilience version 12.3 activated at 05:00:59 UTC Fri Feb 10
2006
Secure archive flash:c1841-advsecurityk9-mz.123-14.T1.bin type is image
(elf) []
  file size is 17533860 bytes, run size is 17699528 bytes
  Runnable image, entry point 0x8000F000, run from ram

IOS configuration resilience version 12.3 activated at 05:01:02 UTC Fri
Feb 10 2
  006
Secure archive flash:.runcfg-20060210-050102.ar type is config
configuration archive size 4014 bytes
```

Configuring Multiple Privilege Levels

Cisco routers enable you to configure various privilege levels for your administrators. You can configure different passwords to control which administrators have access to the various privilege levels. Configuring various privilege levels is especially useful in a help desk environment where you want certain administrators to be able to configure and monitor every part of the router (level 15), while you want other administrators to only monitor, and not configure, the router (customized levels 2 to 14). There are 16 privilege levels, 0 to 15; level 0 is reserved for the user-level access privileges, levels 1 to 14 are levels you can customize, and level 15 is reserved for enable mode privileges.

To assign privileges to levels 2 to 14, use the **privilege** command from global configuration mode:

```
privilege mode {level level command | reset command}
```

[Table 4-3](#) describes the parameters for this command and [Example 4-6](#) demonstrates its use.

Table 4-3. *privilege* Command Parameters

Parameter	Description
<i>mode</i>	This command argument specifies the configuration mode. Use the Router(config)# privilege ? command to see a complete list of router configuration modes available on your router.
<i>level</i>	(Optional) This keyword enables setting a privilege level with a specified command.
<i>level command</i>	(Optional) This parameter is the privilege level that is associated with a command. You can specify up to 16 privilege levels, using numbers 0 to 15.
<i>reset</i>	(Optional) This keyword resets the privilege level of a command.
<i>command</i>	(Optional) Use this argument when you want to reset the privilege level.

Example 4-6. Configuring Multiple Privilege Levels

[Click here to view code image](#)

```
R1 (config)# privilege exec level 2 ping  
R1 (config)# enable secret level 2 Cariboo2012
```

To assign a password to the custom privilege level, use the command **enable secret level level password** in global configuration mode.

To enter a custom privilege level, use the command **enable level** and enter the password that was assigned to the custom privilege level.

[Example 4-6](#) sets the **ping** command to require privilege level 2 or above access and establishes Cariboo2012 as the secret password for privilege level 2. When you enter the **enable 2** command, as shown in [Example 4-7](#), the router prompts you for the enable secret password for privilege level 2.

Use the **show privilege** command to display the current privilege level, as shown in [Example 4-7](#).

Example 4-7. Using the *enable* level and *show privilege* Commands

[Click here to view code image](#)

```
R1> enable 2
Password: Cariboo2012
R1# show privilege
Current privilege level is 2
```

Configuring Role-Based Command-Line Interface Access

The Role-Based CLI Access feature allows you to create different “views” of router configurations for different users. Views define which commands are accepted from different users and what configuration information is visible to them. With Role-Based CLI Access, you can exercise better control over Cisco networking devices.

Note

Before you create a view, you must enable authentication, authorization, and accounting (AAA) using the **aaa new-model** command or CCP. AAA configurations are covered in the “[Configuring AAA on a Cisco Router Using the Local Database](#)” section of this chapter.

The steps used to configure and confirm a view are as follows:

Step 1. Router> **enable view**

Step 2. Router# **configure terminal**

Step 3. Router(config)# **parser view** *view-name*

Step 4. Router(config-view)# **secret** **5** *encrypted-password*

Step 5. Router(config-view)# **commands** *parser-mode* {**include** | **include-exclusive** | **exclude**} [**all**] [**interface** *interface-name* | *command*]

Step 6. Router(config-view)# **exit**

Step 7. Router(config)# **exit**

Step 8. Router# **enable** [*view-name*]

Step 9. Router# **show parser view** [**all**]

The last two steps allow you to preview the views that you have configured. The next few pages discuss these commands in detail.

The key commands specific to configuring views for Role-Based CLI Access are shown next and are presented again in [Example 4-8](#). When a system is in “root view,” it has all the access privileges of a user who has level 15 privileges. To configure any view for the system, you must be in the root view.

The difference between a user who has level 15 privileges and a root view user is that a root view

user can configure a new view and add or remove commands from the view.

To access the root view, use first the **enable view** command and then the **parser view** command:

```
R1> enable view
R1# configure terminal
R1(config)# parser view view-name
R1(config-view)# secret 0 | 5 view-password
```

[Table 4-4](#) shows the commands and parameters used to access and modify the root view.

Table 4-4. *enable view* and *parser view* Command Parameters

Parameter	Description
enable view	This command puts you in root view, from where you create views and establish view attributes.
config terminal	This command puts you in global configuration mode.
parser view <i>view name</i>	This command creates a view and enters view configuration mode.
secret 0 5 <i>view-password</i>	This command configures a password for this view: secret 0 specifies that an unencrypted password follows. secret 5 specifies that an encrypted secret follows.

Next, you must assign the allowed commands to the selected view. Use the **commands** command in view configuration mode to assign the allowed commands. The syntax for this command is as follows:

[Click here to view code image](#)

```
R1(config-view)# commands parser-mode {include | include-exclusive |
exclude} [all]
[interface interface-name | command]
```

[Table 4-5](#) shows the parameters used with the **commands** command.

Table 4-5. *commands* Command Parameters

Parameter	Description
<i>parser-mode</i>	The mode in which the specified command exists (for example, EXEC mode)
include	Adds a command or an interface to the view and allows the same command or interface to be added to an additional view
include-exclusive	Adds a command or an interface to the view and excludes the same command or interface from being added to all other views
exclude	Excludes a command or an interface from the view
all	A “wildcard” that allows every command in a specified configuration mode that begins with the same keyword or every subinterface for a specified interface to be part of the view
interface <i>interface-name</i>	An interface that is added to the view
<i>command</i>	A command that is added to the view

[Example 4-8](#) displays a complete configuration of a new view, called NetOps.

Example 4-8. Commands to Enable Root View and to Create New Views

[Click here to view code image](#)

```
R1> enable view
Password:
R1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# parser view NetOps
R1(config-view)# secret 0 hardtocrackpw
R1(config-view)# commands exec include ping
R1(config-view)# commands exec include all show
R1(config-view)# commands exec include telnet
R1(config-view)# commands exec include traceroute
R1(config-view)# commands exec include write
R1(config-view)# commands exec include configure
R1(config-view)# commands configure include access-list
R1(config-view)# commands configure include all interface
R1(config-view)# commands configure include all ip
```

To verify a view, use the **enable view** command. Enter the name of the view that you want to verify and provide the password to log in to the view. After you are in the view, use the question mark (?) command to verify that the commands available in the view are correct. [Example 4-9](#) shows the commands only accessible from the NetOps view at the privilege mode and at the configuration mode.

Example 4-9. Verifying Commands Available to the NetOps View

```
R1# enable view NetOps
Password: hardtocrackpw
R1#
Jan 3 13:45:03.887: %PARSER-6-VIEW_SWITCH: successfully set to view
'NetOps'.
R1#?
Exec commands:
  configure Enter configuration mode
  enable Turn on privileged commands
  exit Exit from the EXEC
  ping Send echo messages
  show Show running system information
  telnet Open a telnet connection
  traceroute Trace route to destination
  write Write running configuration to memory, network, or terminal
R1# configure terminal
R1(config)#?
Configure commands:
  access-list Add an access list entry
  do To run exec commands in config mode
  exit Exit from configure mode
  interface Select an interface to configure
  ip Global IP configuration subcommands
```

Implementing Secure Management and Reporting

The previous section of this chapter dealt with securing access to Cisco IOS devices. Before we look at AAA configuration and how it assists with secure access, let's have a look at secure management and reporting, because these features can be useful while troubleshooting AAA configuration.

In this section, we examine the skills necessary to implement secure management and reporting of Cisco IOS devices. The technologies will discuss in this section are as follows:

- Syslog
- Network Time Protocol (NTP)
- Simple Network Management Protocol Version 3 (SNMPv3)

In addition, we will examine some design aspects of a management infrastructure.

Planning Considerations for Secure Management and Reporting

Configuring logging for your Cisco routers is a straightforward operation when your network contains only a few Cisco routers. However, logging and reading information from hundreds of devices can prove to be a challenging proposition and can raise the following issues and considerations:

- What are the most important logs?
- How are important messages separated from routine notifications?
- How do you prevent tampering with logs?

- How do you ensure that time stamps match?
- What log data is needed in criminal investigations?
- How do you deal with the volume of log messages?
- How do you manage all the devices?
- How can you track changes when attacks or network failures occur?

Securing administrative access and device configurations is also a straightforward operation for smaller Cisco router networks. However, managing administrative access and device configurations for many devices can raise questions such as those listed.

Each of these issues is specific to your needs. To identify the priorities of reporting and monitoring, you must get input from management and from the network and security teams. The security policy that you implement should also play a large role in answering these questions.

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

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

Secure Management and Reporting Architecture

In the previous chapter, [Figure 3-4](#) showed a management module with two network segments that are separated by a Cisco IOS router that acts as a firewall and a VPN termination device. In [Chapter 3](#), we introduced the general concept of secure management and reporting, such as the difference between out-of-band (OOB) management and in-band management. Now we will look at some specific guidelines.

Secure Management and Reporting Guidelines

The guidelines for OOB management and in-band management of the architecture are as follows:

- Management guidelines
 - Keep clocks on hosts and network devices synchronized.
 - Record changes and archive configurations.
- OOB management guidelines
 - Provide the highest level of security and mitigate the risk of passing unsecure management protocols over the production network.
- In-band management guidelines to manage or monitor devices:

- Use VPN, SSH, or SSL (HTTPS with CCP) when possible.
- Decide whether the management channel needs to be open at all times.

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

NTP is used to synchronize the clocks of various devices across a network. Synchronization of the clocks within a network is critical for digital certificates and for correct interpretation of events within the syslog data.

Enabling Time Features

Because many things that are involved in the security of your network depend on an accurate date and time stamp, such as security certificates, it is important that the router maintains the correct time.

You can use Cisco Configuration Professional to configure the date and time settings of the router in three ways:

- Synchronize with the local PC clock
- Manually edit the date and time
- Configure NTP

Network Time Protocol

Time synchronization is of essence in secure management and reporting. For management, using time synchronization ensures that digital certificates are used consistently with regard to their expiration and validity. For reporting, time synchronization results in consistent correlation of events across the management infrastructure, and a more effective event and incident detection and analysis.

NTP is a method to synchronize date and time settings for devices on the network. NTP uses User Datagram Protocol (UDP) port 123 and is documented in RFC 1305. Simple Network Time Protocol (SNTP) is a simpler, less secure version of NTP.

When you implement NTP in your network, you can set up your own master clock, or you can use a publicly available NTP server on the Internet. If you implement your own master clock, you should synchronize the private network to Coordinated Universal Time (UTC) via satellite or radio.

You need to be careful when you implement NTP. An attacker can launch a denial of service (DoS) attack by sending bogus NTP data across the Internet to your network in an attempt to change the clocks on network devices, possibly causing digital certificates to become invalid. Further, an attacker could attempt to confuse a network administrator during an attack by disrupting the clocks on network devices. This scenario would make it difficult for the network administrator to determine the order of syslog events on multiple devices.

NTP version 3 (NTPv3) and later support a cryptographic authentication mechanism between NTP peers. You can use this authentication mechanism in addition to access control lists (ACL) that specify which network devices are allowed to synchronize with other network devices, to help mitigate such an attack.

Note

You should weigh the benefits of using clock time from the Internet against the possible risk of doing so and allowing unsecured packets through the firewall. Many NTP servers on the Internet don't require any authentication of peers and don't have a requirement of being accurate. Trustworthy sources of time are Symmetricom (<http://www.symmetricom.com>), Meingerb (<http://www.meinberg-usa.com>), and National Research Council Canada (<http://www.nrc-cnrc.gc.ca/eng/services/time/index.html>).

More information on NTP, such as the level of clock hierarchy known as stratum, can be found online at http://en.wikipedia.org/wiki/Network_Time_Protocol.

Using a Cisco IOS Router as an NTP Server

It is possible to configure your Cisco IOS router as an NTP master, which other appliances will contact to synchronize on. The following commands are used to set the router as an NTP master:

```
router# conf t
router(config)# ntp authenticate
router(config)# ntp trusted-key 99
router(config)# ntp master
router(config)# key chain NTP
router(config-keychain)# key 99
router(config-keychain-key)# key-string PROPERTIME
router(config-keychain-key)# end
```

Using Syslog Logging for Network Security

Syslog is the standard for logging system events. As shown in [Figure 4-2](#), syslog implementations contain two types of systems:

- **Syslog servers:** These systems are also known as log hosts. These systems accept and process log messages from syslog clients.
- **Syslog clients:** Syslog clients are routers or other types of Cisco equipment that generate and forward log messages to syslog servers.

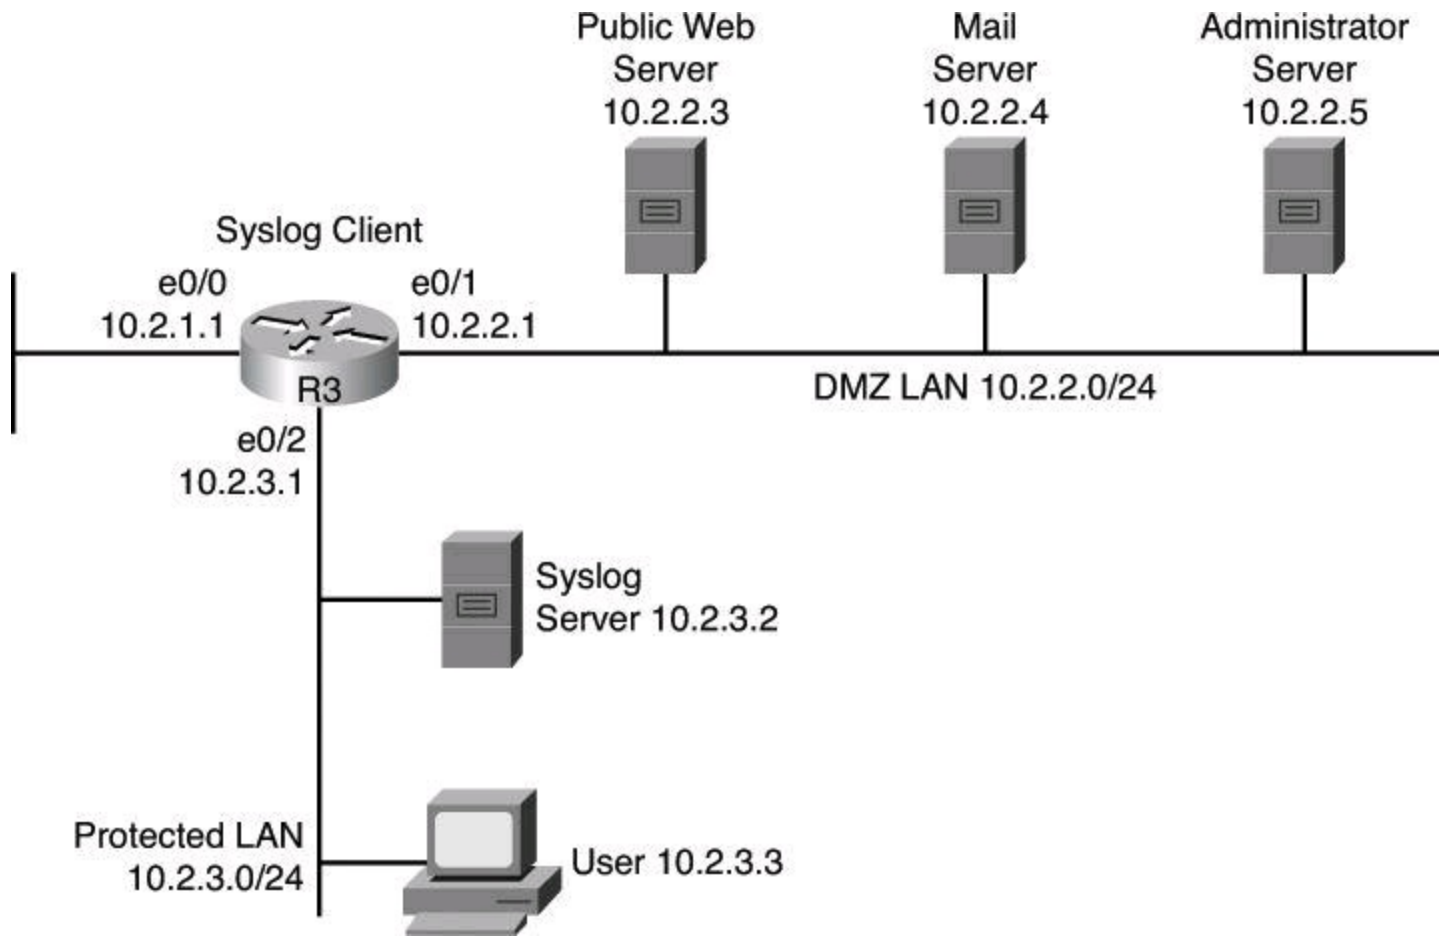


Figure 4-2. Syslog Systems

Note

Performing forensics on router logs can become very difficult if your router clocks are not running the proper time. It is recommended that you use an NTP facility to ensure that all of your routers are operating at the correct time.

If you are not running your own NTP service, you should at least consider synchronizing on an authenticated public NTP service such as the one offered by the National Research Council Canada at http://www.nrc-cnrc.gc.ca/eng/solutions/advisory/calibration/time_frequency.html#authenticated.

Also note that it is recommended to have redundant syslog servers on a network. Many approaches exist on how to use the two systems, such as

- Have the network device send its syslog messages at both servers.
- Have the network device send its syslog messages to only one server but have that server, keeping a copy of the message, forward it immediately to the second server.

Implementing Log Messaging for Security

Implementing a router logging facility is an important part of any network security policy. Cisco routers can log information regarding configuration changes, ACL violations, interface status changes, and many other types of events. Cisco routers can send log messages to several different facilities. You should configure the router to send log messages to one or more of the following items:

- **Console:** Console logging is used when modifying or testing the router while it is connected to the console. Messages sent to the console are not stored by the router and, therefore, are not very valuable as security events.
- **Terminal monitor:** The console port works at 9600 bauds and thus can't simultaneously handle a large quantity of log messages. It is therefore recommended that you remotely connect to the device, using SSH preferably, and that you issue the **terminal monitor** command to receive the log message on this current vty session. This is the subject of the next bullet.
- **Terminal lines:** You can configure enabled EXEC sessions to receive log messages on any terminal lines. Similar to console logging, this type of logging is not stored by the router and, therefore, is valuable only to the user on that line.
- **Buffered logging:** You can direct a router to store log messages in router memory. Buffered logging is a little more useful as a security tool but has the drawback of having events cleared whenever the router is rebooted. Note also that buffers can be written to flash.
- **SNMP traps:** Certain router events can be processed by the router SNMP agent and forwarded as SNMP traps to an external SNMP server. SNMP traps, addressed to destination port UDP 162, are a viable security logging facility but require the configuration and maintenance of an SNMP system.
- **Syslog:** You can configure Cisco routers to forward log messages to an external syslog service, destination port UDP 514. This service can reside on any number of servers, including Microsoft Windows and UNIX-based systems, or since the Cisco Security MARS appliance is at end-of-life status, consider using Cisco Security Manager (CSM) or a Cisco ecosystem partner for Security Information and Event Management System (SIEM), such as Splunk, as a syslog destination. Syslog is the most popular message logging facility because it provides long-term log storage capabilities and a central location for all router messages.

Cisco router log messages fall into one of eight levels, as shown in [Table 4-6](#). The lower the level number, the higher the severity level, as the log messages in the table denote.

Table 4-6. Cisco Router Log Severity Messages

Syslog Level	Definition	Example
0: LOG_EMERG	A panic condition normally broadcast to all users	Cisco IOS Software could not load.
1: LOG_ALERT	A condition that should be corrected immediately, such as a corrupted system database	Temperature too high.
2: LOG_CRIT	Critical conditions; for example, hard device errors	Unable to allocate memory.
3: LOG_ERR	Errors	Invalid memory size.
4: LOG_WARNING	Warning messages	Crypto operation failed.
5: LOG_NOTICE	Conditions that are not error conditions, but should possibly be handled specially	Interface changed state, up or down.
6: LOG_INFO	Informational messages	Packet denied by ACL.
7: LOG_DEBUG	Messages that contain information normally of use only when debugging a program	Packet type invalid.

Note

When entering logging levels in commands, you must specify the level name or the level number.

Cisco router log messages contain three main parts:

- Time stamp
- Log message name and severity level
- Message text

[Figure 4-3](#) shows a syslog entry example for a level 5 syslog message, indicating that someone has configured the router using the vty 0 port.

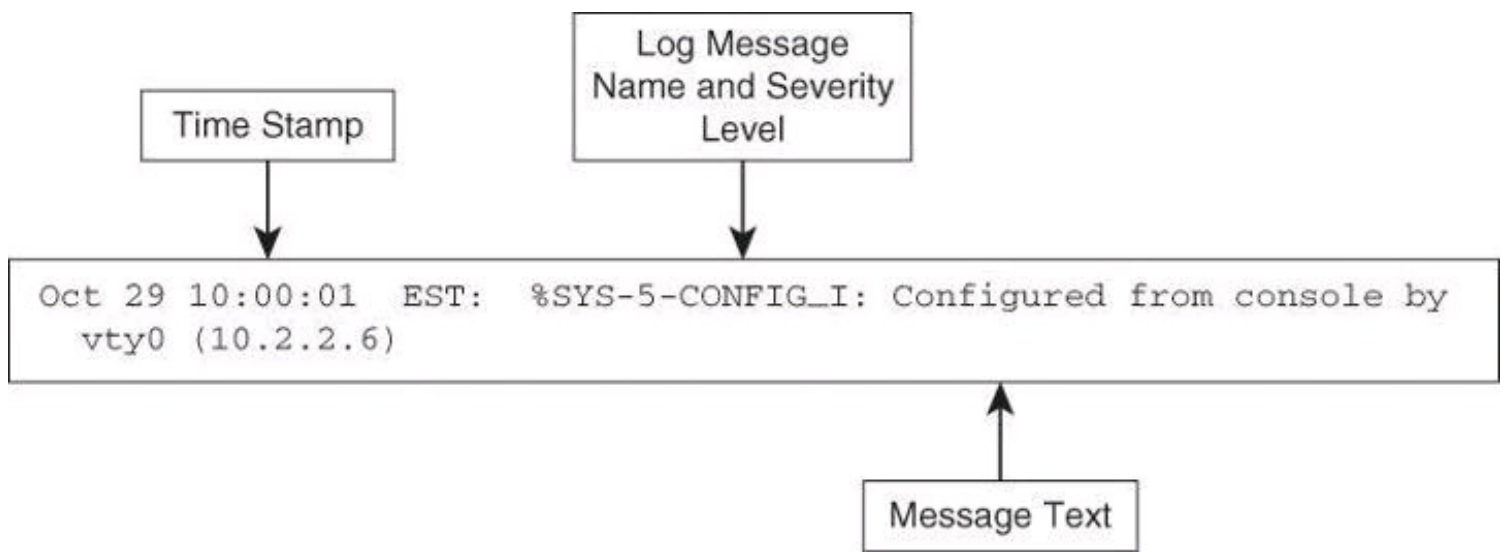


Figure 4-3. Log Message Format

[Figure 4-4](#) illustrates the configuration of syslog settings using Cisco Configuration Professional. Navigate to **Configure > Router > Logging** to enable logging of system messages, and to specify logging hosts where logs can be kept.

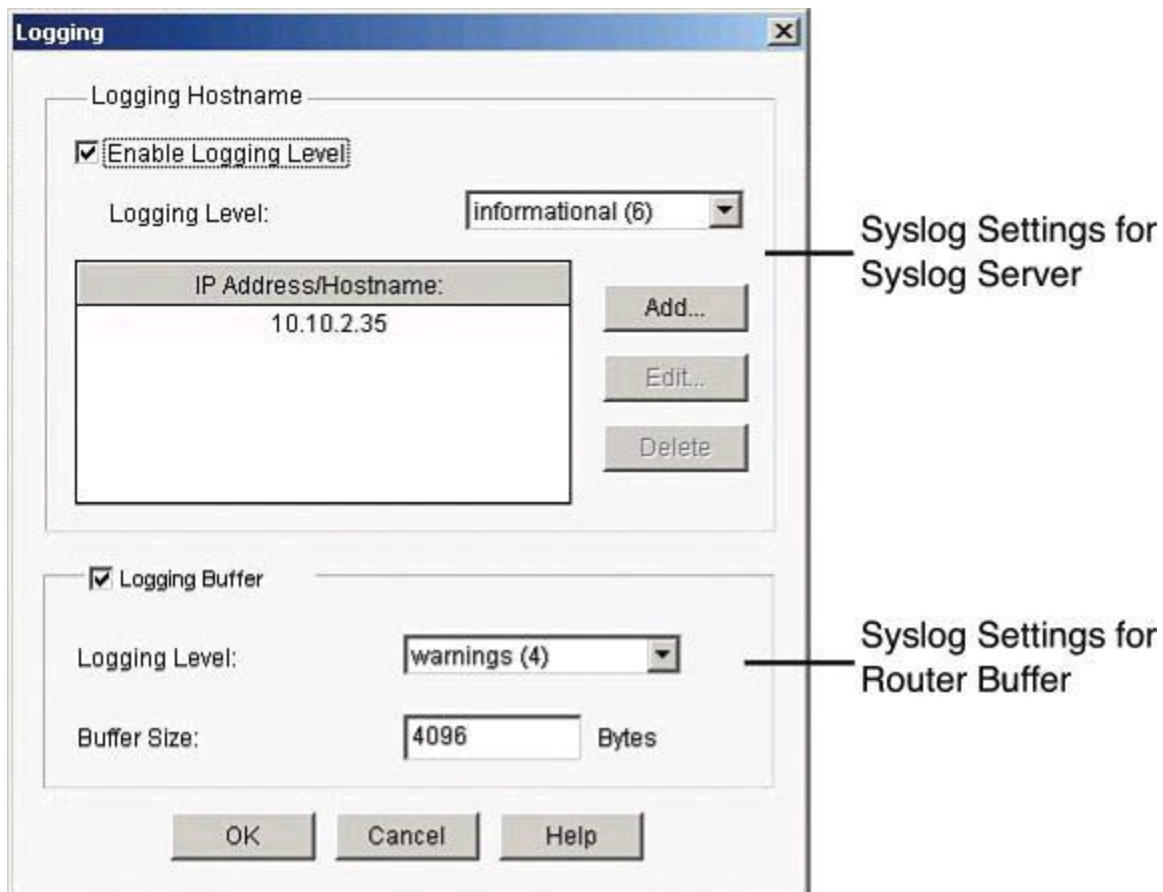


Figure 4-4. Enabling Syslog Logging on CCP

By clicking **Edit**, you can change the syslog settings and specify the level of logging messages that you want to send and to collect, and enter the hostname or IP address of multiple logging hosts.

The logging settings are shown in [Figure 4-4](#). The top of the window allows the configuration of syslog server parameters. The bottom half allows the configuration of the router buffers as receivers of router syslog messages.

When defining the logging level for a syslog server, the log collects all messages from the level you choose plus all messages from lower levels, or the router sends all messages of the level you choose plus all messages of lower levels to the logging hosts. For example, if you choose notifications (5), the log collects or sends messages of levels 0 through 5. Firewall logging messages require a logging level of debugging (7), and application security logging messages require a level of informational (6).

Using SNMP to Manage Network Devices

SNMP is another management option, not only for messaging and notifications via SNMP traps and informs, but also as a general management framework. SNMP was developed to manage nodes, such as servers, workstations, routers, switches, hubs, and security appliances, on an IP network. All versions of SNMP are application layer protocols that facilitate the exchange of management information between network devices. SNMP is part of the TCP/IP protocol suite. SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth.

SNMP Version 1 (SNMPv1) and SNMP Version 2 (SNMPv2) are based on three concepts:

- Managers (network management systems [NMS] installed on servers)
- Agents (installed on managed nodes)
- Management Information Bases (MIB)

In any configuration, at least one manager node runs SNMP management software. Network devices that need to be managed, such as switches, routers, servers, and workstations, are equipped with an SMNP agent software module. The agent is responsible for providing access to a local MIB of objects that reflects the resources and activity at its node.

MIBs

A Management Information Base is a database of the objects that can be managed on a device. The managed objects, or variables, can be set or read to provide information on the network devices and interfaces.

The SNMP manager can retrieve, or “get,” information from the agent, and change, or “set,” information in the agent, as shown in [Figure 4-5](#). Sets can change variables (settings, configuration) in the agent device or initiate actions in devices. A reply to a set indicates the new setting in the device. For example, a set can cause a router to reboot, send a configuration file, or receive a configuration file. SNMP traps enable an agent to notify the management station of significant events by sending an unsolicited SNMP message.

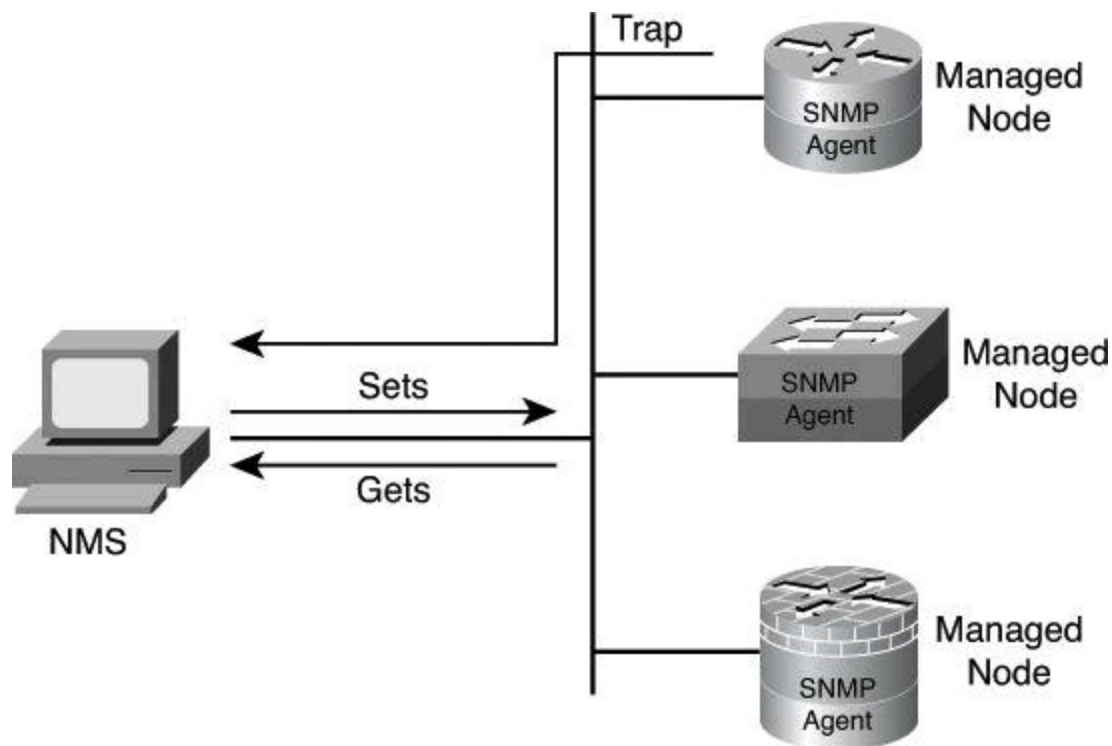


Figure 4-5. SNMPv1/v2 Architecture

The actions of gets and sets are the vulnerabilities that open SNMP to attack.

SNMPv1 and SNMPv2 use a community string to access router SNMP agents. SNMP community strings act like passwords. An SNMP community string is a text string that can authenticate messages between a management station and an SNMP engine:

- If the manager sends one of the correct read-only community strings, it can get information but not set information in an agent.
- If the manager uses one of the correct read-write community strings, it can get or set information in the agent.

In effect, having set access to a router is equivalent to having the enable password of the router.

SNMP agents accept commands and requests only from SNMP systems using the correct community string. By default, most SNMP systems use “public” as a community string. If you configure your router SNMP agent to use this commonly known community string, anyone with an SNMP system is able to read the router MIB. Because router MIB variables can point to things such as routing tables and other security-critical parts of the router configuration, it is extremely important that you create your own custom SNMP community strings.

SNMPv3 Architecture

In its natural evolution, the current version of SNMPv3 addresses the vulnerabilities of earlier versions by including three important services: authentication, privacy, and access control.

SNMPv3 is an interoperable, standards-based protocol for network management. SNMPv3 uses a combination of authenticating and encrypting packets over the network to provide secure access to devices, as shown in [Figure 4-6](#). SNMPv3 provides the following security features:

- **Message integrity:** Ensures that a packet has not been tampered with in transit

- **Authentication:** Determines that the message is from a valid source
- **Encryption:** Scrambles the contents of a packet to prevent it from being seen by an unauthorized source

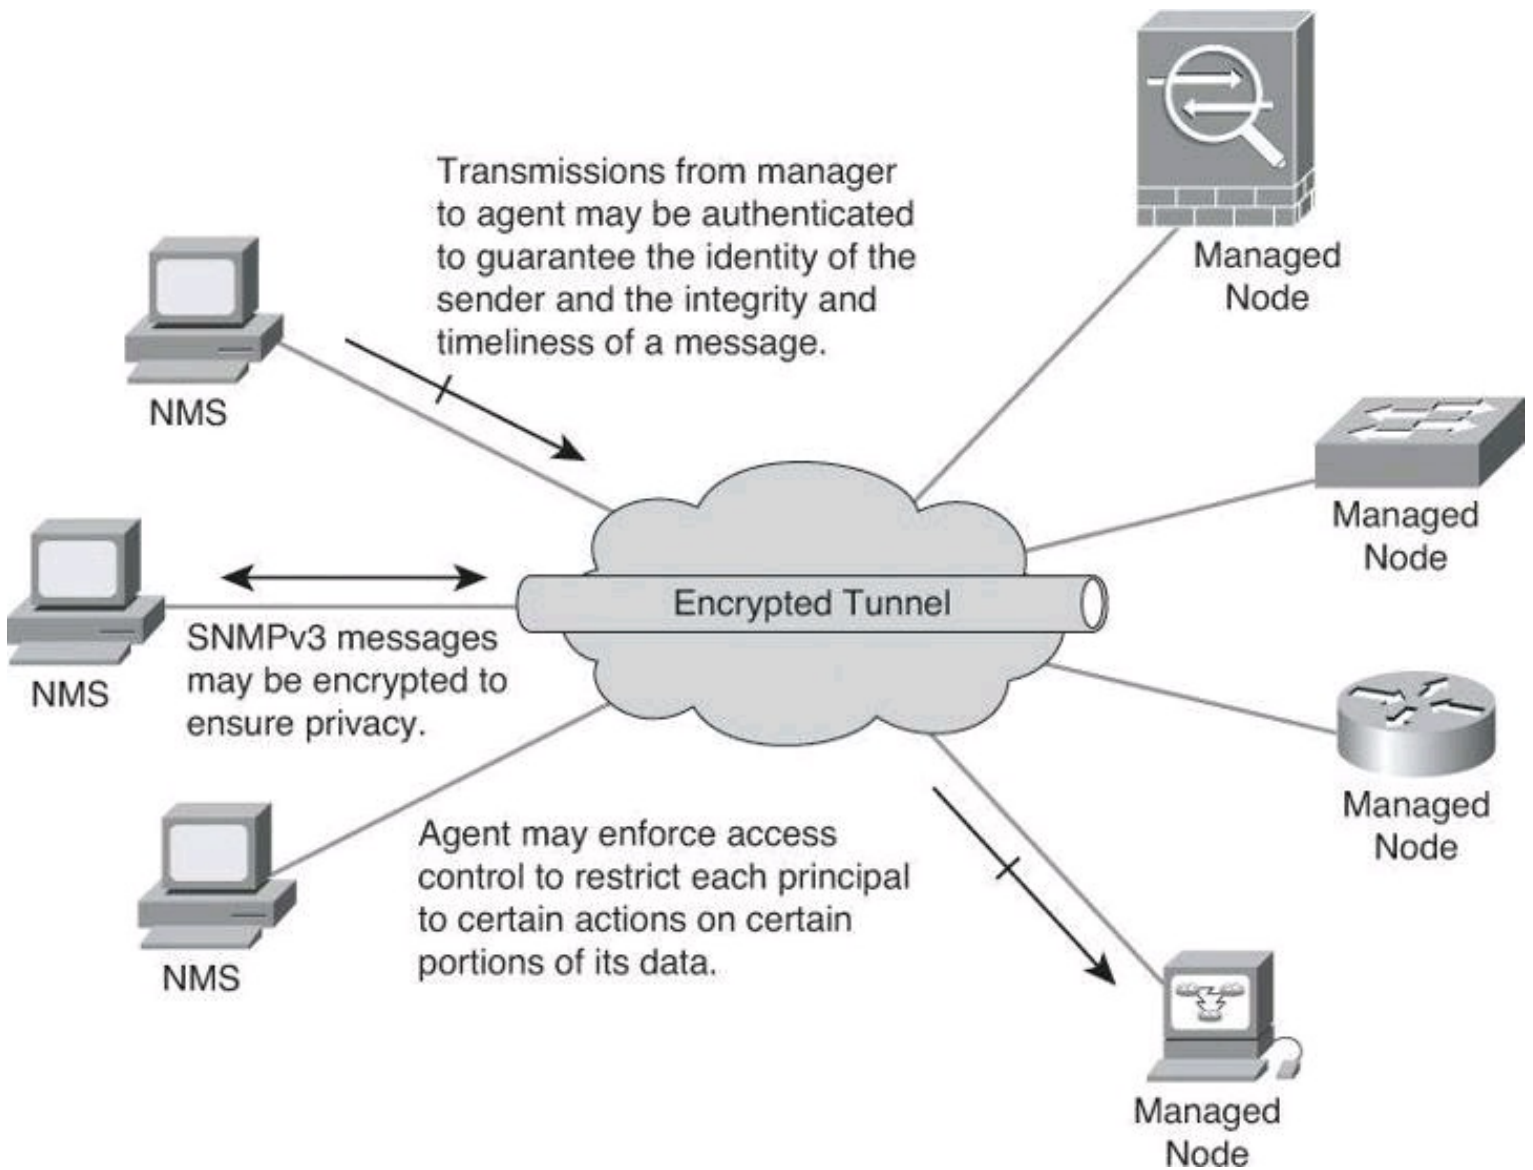


Figure 4-6. SNMPv3 Architecture

SNMPv3 provides for a combination of both security model and security level, which determines the security mechanism that will be used when handling an SNMP packet.

A *security model* is an authentication strategy that is set up for a user and the group in which the user resides. Currently, Cisco IOS Software supports three security models: SNMPv1, SNMPv2c, and SNMPv3. A *security level* is the permitted level of security within a security model. The security level is a type of security algorithm that is performed on each SNMP packet. There are three security levels:

- **noAuth:** This security level authenticates a packet by a string match of the username or community string.
- **auth:** This level authenticates a packet by using either the Hashed Message Authentication Code (HMAC) with Message Digest 5 (MD5) method or the HMAC with Secure Hash Algorithms (SHA) method. Both methods are described in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.

- **Priv:** This level authenticates a packet by using either the HMAC MD5 or HMAC SHA algorithm and encrypts the packet using the Data Encryption Standard (DES), Triple DES (3DES), or Advanced Encryption Standard (AES) algorithm.

Note

Only SNMPv3 supports the auth and priv security levels.

[Table 4-7](#) identifies what the combinations of security models and levels mean. As mentioned earlier, hashing and encryption will be explained in [Chapter 12](#).

Table 4-7. AAA Accounting Using Named Method Lists Procedure

	Level	Authentication	Encryption	What Happens
SNMPv1	noAuthNoPriv	Community string	No	Authenticates with a community string match
SNMPv2c	noAuthNoPriv	Community string	No	Authenticates with a community string match
SNMPv3	noAuthNoPriv	Username	No	Authenticates with a username
SNMPv3	authNoPriv	MD5 or SHA	No	Provides HMAC MD5 or HMAC SHA algorithm for authentication
SNMPv3	authPriv	MD5 or SHA	Yes	Provides HMAC MD5 or HMAC SHA algorithm for authentication; provides DES, 3DES, or AES encryption in addition to authentication

Enabling SNMP Options Using Cisco CCP

[Figure 4-7](#) illustrates the configuration of SNMP settings using Cisco Configuration Professional. Navigate to **Configure > Router > SNMP** to enable SNMP, set SNMP community strings, and enter SNMP trap receiver information.

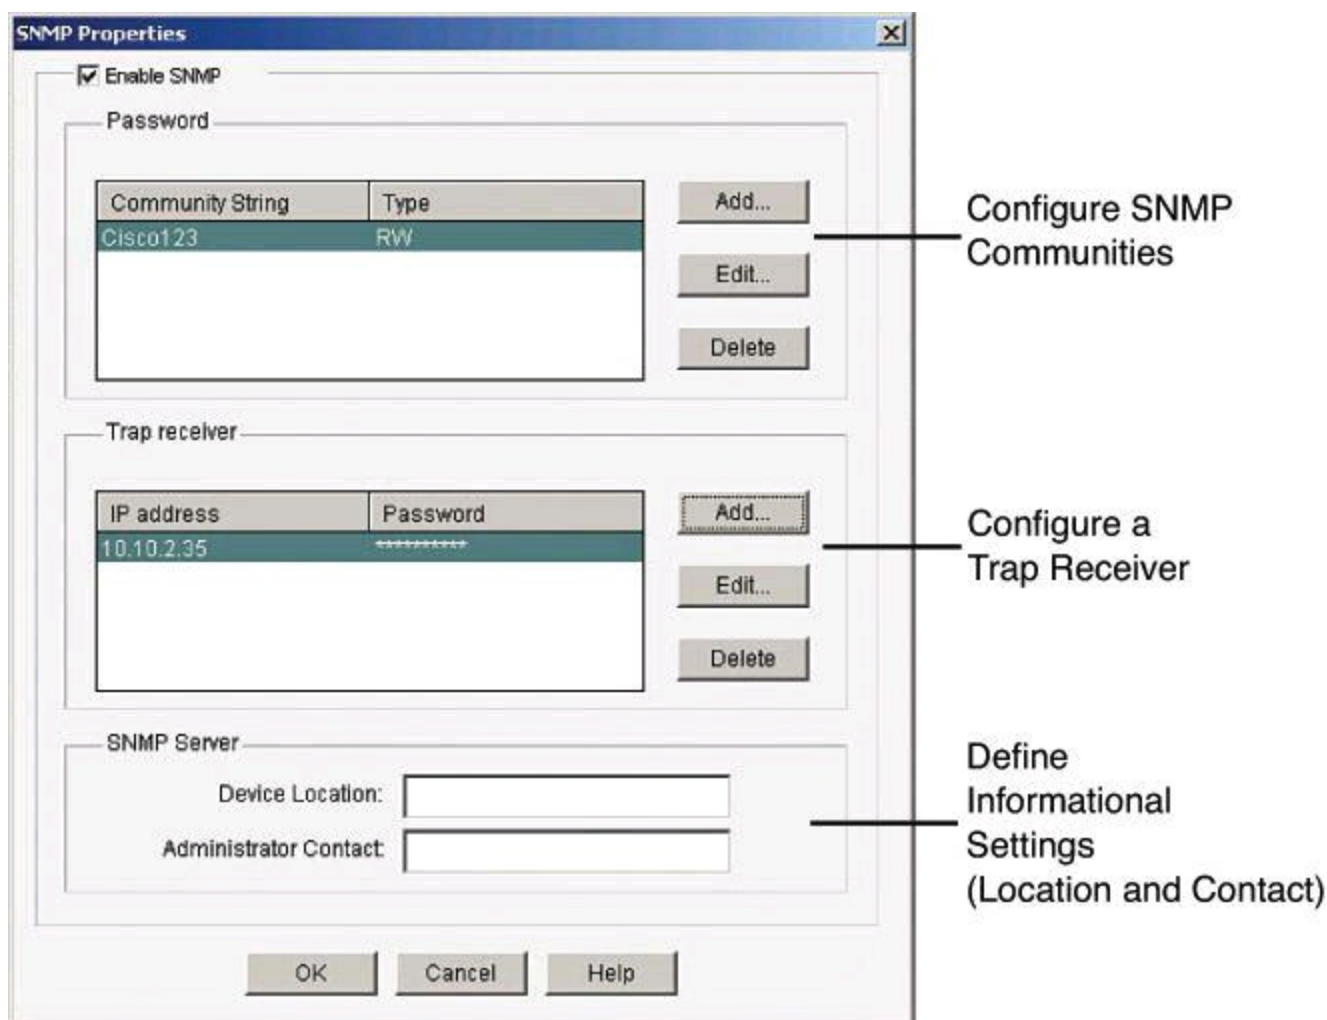


Figure 4-7. Enabling SNMP with CCP

Configuring AAA on a Cisco Router

Previously in this chapter, we discussed how to create a more secure management and reporting environment. A significant part of that environment is AAA. In [Chapter 3](#), we briefly discussed the use of AAA with regard to management plane security. In this section, we will see how AAA can be implemented locally or centrally.

As mentioned, one of the options you have when configuring your network to work with AAA is to use a local username and password database to provide security greater than a simple password. It is likely that smaller organizations will configure AAA to operate locally.

Authentication, Authorization, and Accounting

Access control is the way you control who is allowed access to the access server or router and which services they are allowed to use once they have access. AAA network security services provide the primary framework through which you set up access control on your router. AAA services provide a higher degree of scalability than the line-level and privileged EXEC authentication commands alone.

Network and administrative access security in the Cisco environment, whether it involves campus access or VPN access, is based on a modular architecture that has three functional components:

- **Authentication:** Authentication requires users and administrators to prove that they really are who they say they are. Authentication is established using a username and password,

challenge and response, token cards, and other methods, such as “I am user *student* and my password *validateme* proves it.”

- **Authorization:** After authenticating the user and administrator, authorization services decide which resources the user and administrator are allowed to access and which operations the user and administrator are allowed to perform, such as “User *student* can access host serverXYZ using Telnet.”

- **Accounting and auditing:** Accounting records what the users and administrators actually did, what they accessed, and for how long they accessed it. Accounting keeps track of how network resources are used, such as “User *student* accessed host serverXYZ using Telnet for 15 minutes.”

Two examples of AAA implementation include authenticating remote users that are accessing the corporate LAN through VPN connections (as shown in [Figure 4-8](#)) and authenticating administrator access to the router console port, auxiliary port, and vty ports.

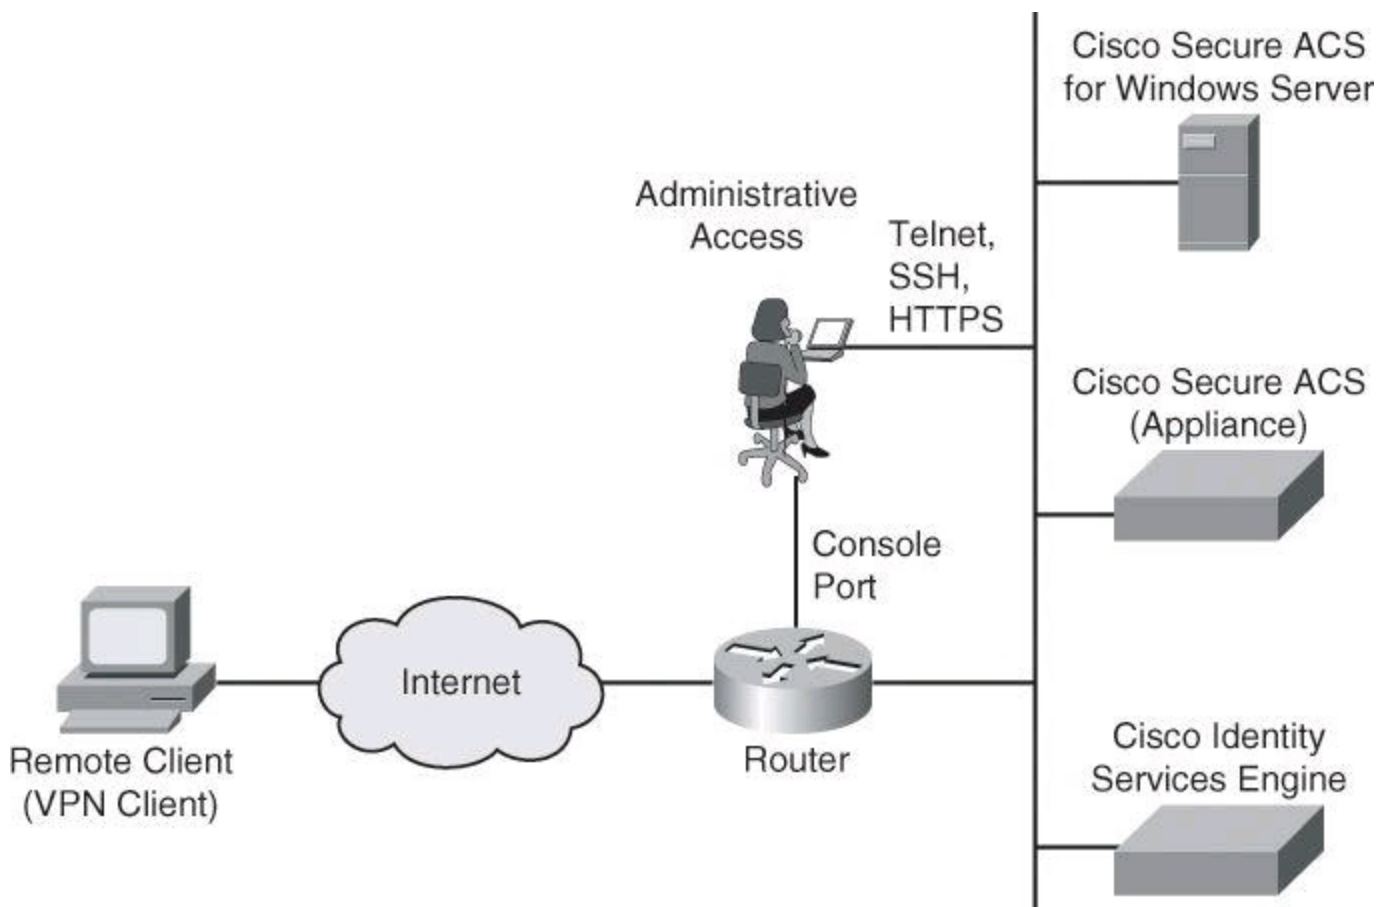


Figure 4-8. Implementing Cisco AAA

Cisco networking products support AAA access control using a local usernames and passwords database or remote security server databases. A local security database was introduced in [Chapter 3](#) with the **username** command, and covered in detail earlier in this chapter with [Table 4-1](#). A remote security server database is a separate server that provides AAA services for multiple network devices and a large number of network users by running RADIUS or TACACS+ protocols.

Note

Omitted from [Figure 4-8](#) is “NAS services.” If your Cisco router accepts dialup

connections, the router is described as a network access server (NAS). With the popularity of VPN access, we don't hear as much nowadays about remote-access dialup services, but they are still around, which explains why [Table 4-8](#) briefly discusses NAS ports.

Table 4-8. Router Access

Access Type	Mode	NAS Ports	Specifications
Remote administrative access	Character (line or EXEC mode)	tty, vty, auxiliary, and console	login, exec, enable
Remote network access	Packet (interface mode)	async, group-async BRI and PRI	ppp, network

Cisco provides many ways to implement AAA services for Cisco routers, all of which are covered in greater detail later in this chapter:

- **Self-contained AAA:** AAA services can be self-contained in the router itself. This form of authentication is also known as local authentication, with the **username** command.
- **Cisco Secure Access Control Server (ACS) for Windows:** AAA services on the router or NAS contact an external Cisco Secure ACS for Microsoft Windows system for user and administrator authentication.

Note

In 2011, Cisco announced the end of life of Cisco Secure ACS for Windows. It is still vastly used and it is still supported, and therefore we are briefly covering it in this book. However, if you are planning a new implementation for external authentication, the next two options are your current Cisco choices.

- **Cisco Secure ACS:** A policy-driven access control system and an integration point for network access control and identity management. The ACS 5.2 software runs either on a dedicated Cisco 1121 Secure Access Control System appliance or on a VMware server.
- **Cisco Identity Services Engine (ISE):** Cisco ISE is the new-generation centralized policy engine for business-relevant policy definition and enforcement. Cisco ISE consolidates features and functions that are found in Cisco Secure ACS and Network Admission Control (NAC) products, to deliver all the necessary services that are required by an enterprise network (AAA, profiling, posture, and guest management) in a single appliance platform. To learn more about ISE, check out TechWiseTV on YouTube for “Fundamentals of ISE.”

Authenticating Router Access

You can use AAA to secure two different types of router access mode. The mode refers to the format of the packets that are requesting AAA services:

- **Character mode:** A user is sending a request to establish an EXEC mode process with

the router, for administrative purposes. The user wishes to get the router prompt to start managing the router as the administrator.

- **Packet mode:** A user is sending a request to establish a dialup connection through the router with a device on the network. The user does not wish to get the router prompt. The sole purpose of connecting to the router is to get access to network resources. This is no longer a frequent scenario because remote users typically connect to organization resources with a VPN connection instead of a dialup connection.

With the exception of accounting commands, all the AAA commands apply to both character mode and packet mode.

For a truly secure network, you must configure the router to secure administrative access and remote LAN network access using AAA services.

[Table 4-8](#) compares the router access modes, port types, and AAA command elements.

The **aaa authentication** command can be used for different authentication settings. [Table 4-9](#) describes common options of the **aaa authentication** command. Many more options are available than those presented here.

Table 4-9. *aaa authentication* Command Parameters

Parameter	Description
enable default	Used to enable AAA authentication to determine if a user can access the privileged command level, referred to in Table 4-8 as remote administrative access, enable and exec modes.
login	Used to set AAA authentication at login, referred to in Table 4-8 as remote administrative access, login. The aaa authentication login command will be covered later in this chapter
ppp	Used to specify one or more AAA authentication methods for use on serial interfaces that are running Point-to-Point Protocol (PPP), referred to in Table 4-8 as remote network access mode.

Configuring AAA Authentication and Method Lists

AAA authentication is based on method lists as its building blocks. A method list is a sequential list describing the authentication methods to be queried in order 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 listed method to authenticate users. If that method fails to respond, the Cisco IOS Software selects the next authentication method that is listed in the method list. This process continues until there is successful communication with a listed authentication method, or all methods that are defined in the method list are exhausted. In [Figure 4-9](#), the default method list showcases two fallback methods. A typical configuration usually includes an AAA server database as the first method and the local database as a fallback method.

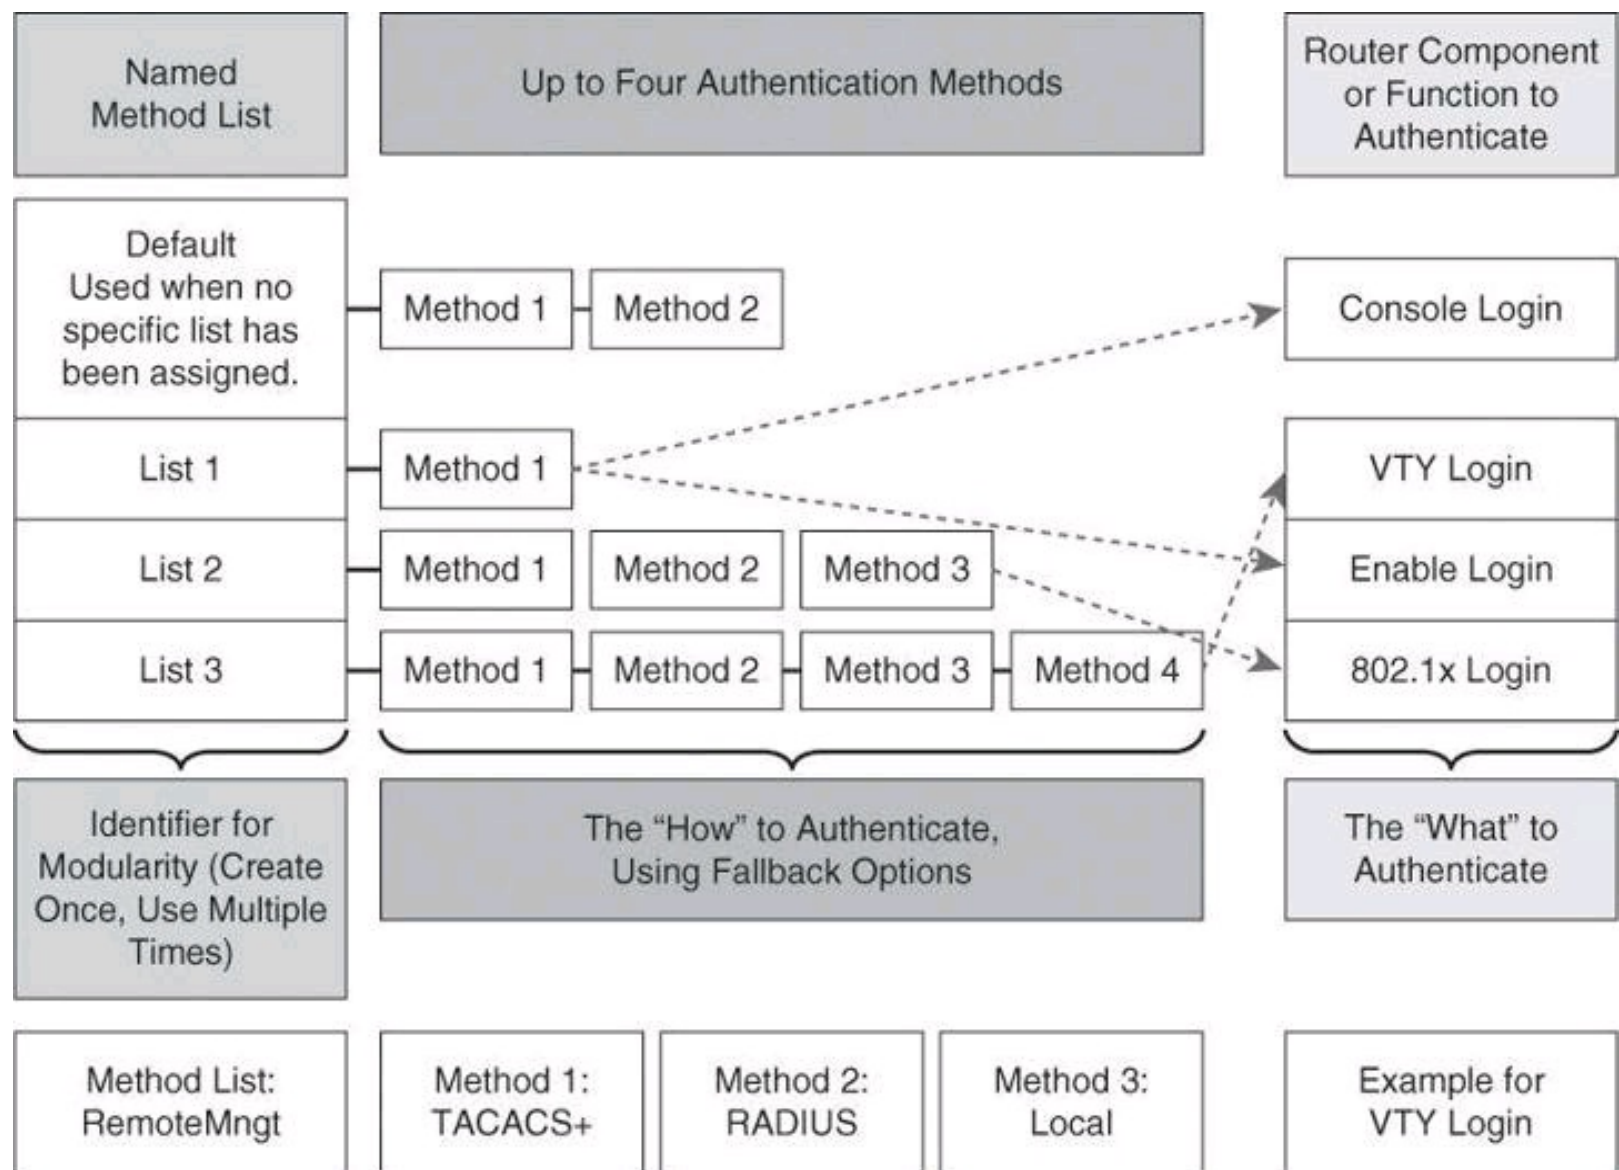


Figure 4-9. Modular Objects in AAA Configuration

It is important to note that the 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—meaning that the security server or local username database responds by denying the user access—the authentication process stops and no other authentication methods are attempted.

Each method list is an individual configuration object. They can be assigned to multiple router components or function to provide authentication services. In [Figure 4-9](#), the authentication methods in List 1 are assigned to both console and enable logins. List 2 methods apply to 802.1X logins, and List 3 methods apply to vty logins. Any other component not assigned a method list will use the methods on the default list.

At the bottom of [Figure 4-9](#), we have an example where a method list called RemoteMngt is created and assigned to vty logins. When an administrator connects to the device with an SSH session, the administrator's credentials will first be checked against a TACACS+ server. If there is no communication with the TACACS+ server, the administrator's credentials will be presented to the RADIUS server, and if that communication fails as well, the administrator's credentials will be checked against the local database. The difference between the TACACS+ and RADIUS protocols

will be discussed later in this chapter.

Although it is not common to see this level of complexity and granularity for different access methods on the same router, the modular configuration that is based on method lists allows for this level of granularity and flexibility.

Configuring AAA on a Cisco Router Using the Local Database

If you have only a few devices that provide remote access to your network for a limited number of users, you can store username and password security information locally on the Cisco devices. This is referred to as local authentication on a local security database. The following are local authentication characteristics:

- Used for small networks
- Stores usernames and passwords in the Cisco router
- Users authenticate against the local security database in the Cisco router
- Does not require an external database

The system administrator must populate the local security database by specifying username and password profiles for each user that might log in. This local authentication is difficult to scale because

- There is limited persistent storage on network devices (running and startup configs).
- If the same credentials are used to access multiple network devices, those credentials require manual replication on concerned devices.

Local authentication typically works as follows, as shown in [Figure 4-10](#):

1. The client establishes a connection with the router.
2. The router prompts the user for a username and password.
3. The router authenticates the username and password in the local database. The user is authorized to access the network based on information in the local database.




Figure 4-10. Implementing Authentication Using Local Services

Configuring AAA Local Authentication

To configure AAA local services to authenticate administrator access (character mode access) or network access (packet mode) that includes VPN access, follow these general steps:

- Add usernames and passwords to the local router database for users who need administrative access to the router.
- Enable AAA globally on the router, or confirm that it is already enabled.

- Configure AAA parameters on the router. These parameters include authentication policies at a minimum, using method lists for the desired access type. Authorization and accounting policies can optionally be configured.
- Confirm and troubleshoot the AAA configuration.

The first step to configure AAA services for local authentication is to create users. Local user accounts can be configured by navigating to **Configure > Router > Router Access > User Accounts/View** and clicking Add, as shown in [Figure 4-11](#). This window enables you to define accounts and passwords that will enable users to authenticate when logging in to the router using HTTP, Telnet, PPP, or some other means. Privilege levels and CLI views can also be configured using this option by adding a new user or editing existing users.

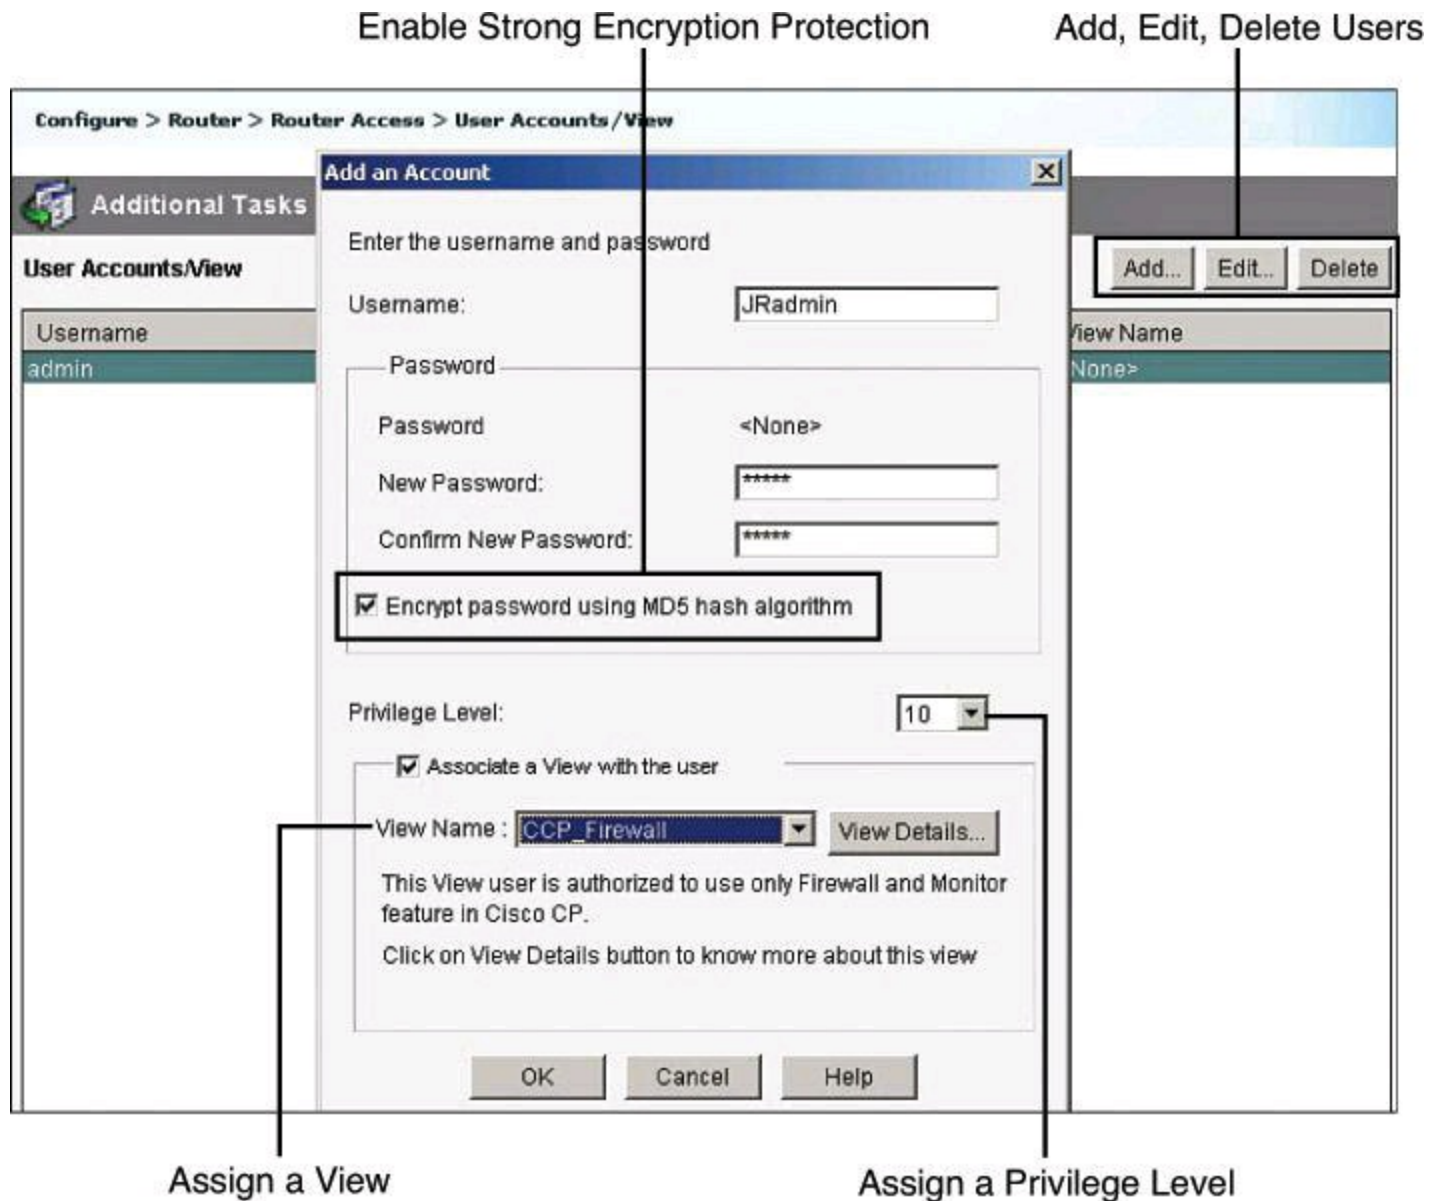


Figure 4-11. Configuring Local User Accounts Using CCP

The following are the detailed steps for adding a local user account:

Step 1. Choose **Configure > Router > Router Access > User Accounts/View**.

Step 2. Click **Add** to add a new user.

Step 3. In the Add an Account dialog box, enter the username and password in the appropriate fields to define the user account.

Step 4. From the Privilege Level drop-down list, choose **15** unless you have defined lesser privilege levels, as shown in [Figure 4-11](#).

Step 5. If you have defined views, you can check the **Associate a View with the User** check box and choose from the View Name drop-down list a view that you want to associate with this user.

Step 6. Click **OK**.

Enabling AAA Authentication Policy

Authentication policies for local access can be configured by navigating to **Configure > Router > AAA > AAA Summary**. From this window, shown in [Figure 4-12](#), you can enable AAA globally by clicking the **Enable AAA** button and then clicking **Yes** in the Enable AAA dialog box that appears. As a precaution and a first level of protection, Cisco Configuration Professional will enable authentication and authorization for the vty lines and authentication for the console line using the local database of the router. A warning message lets the administrator know about this default policy, which is also displayed in the summary window when the commands are delivered to the router. Other authentication policies can be created using authentication method lists.

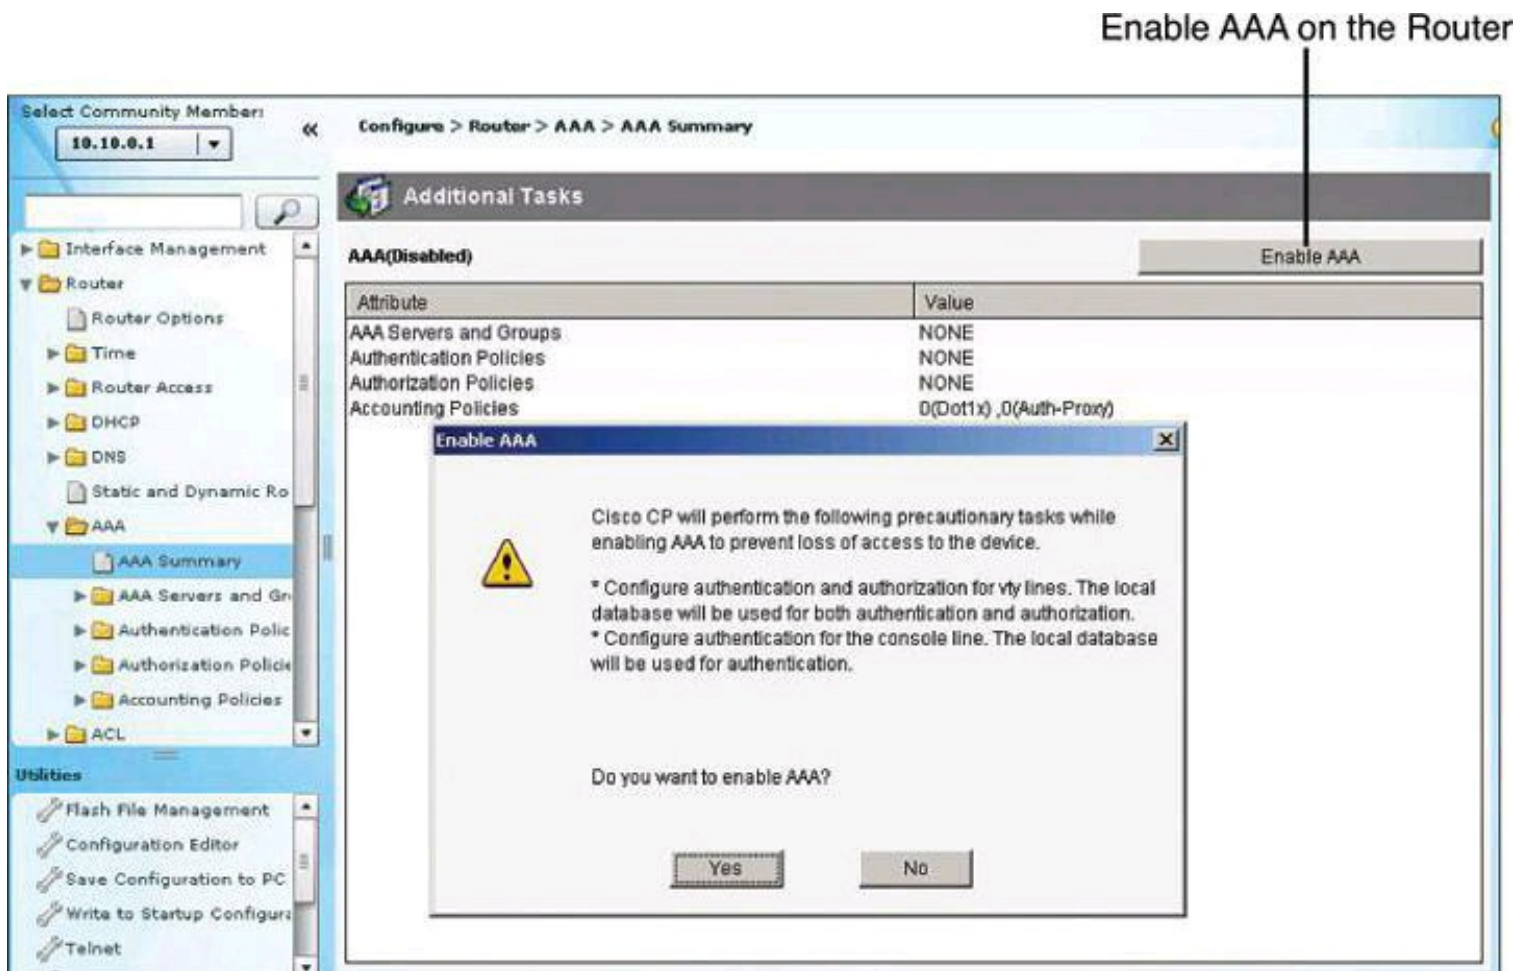


Figure 4-12. Enabling and Disabling AAA Using CCP

[Figure 4-13](#) shows the resulting CCP configuration following the enabling of AAA.

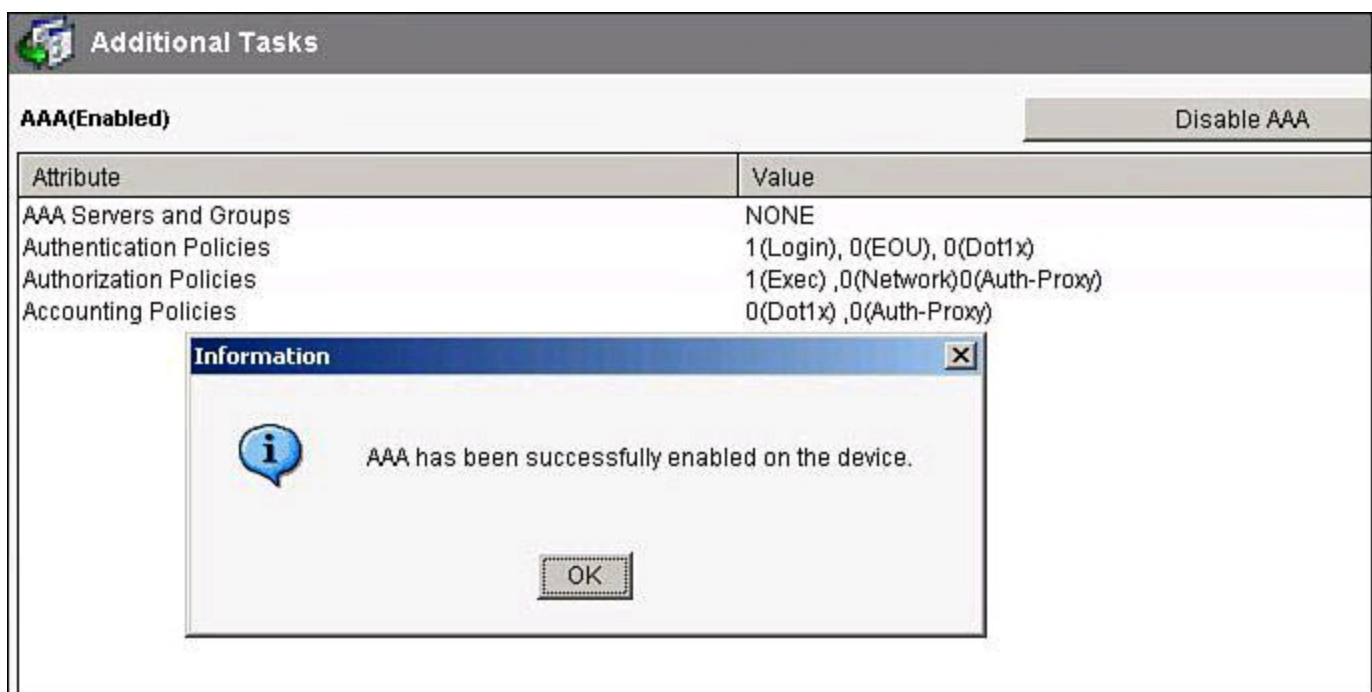


Figure 4-13. AAA Summary After AAA Is Enabled

Configuring Method Lists

Authentication method lists can be created or edited by navigating to **Configure > Router > AAA > Authentication Policies > Login**. The default method list is shown in the background window of [Figure 4-14](#), as well as how to add a new list or edit an existing list.

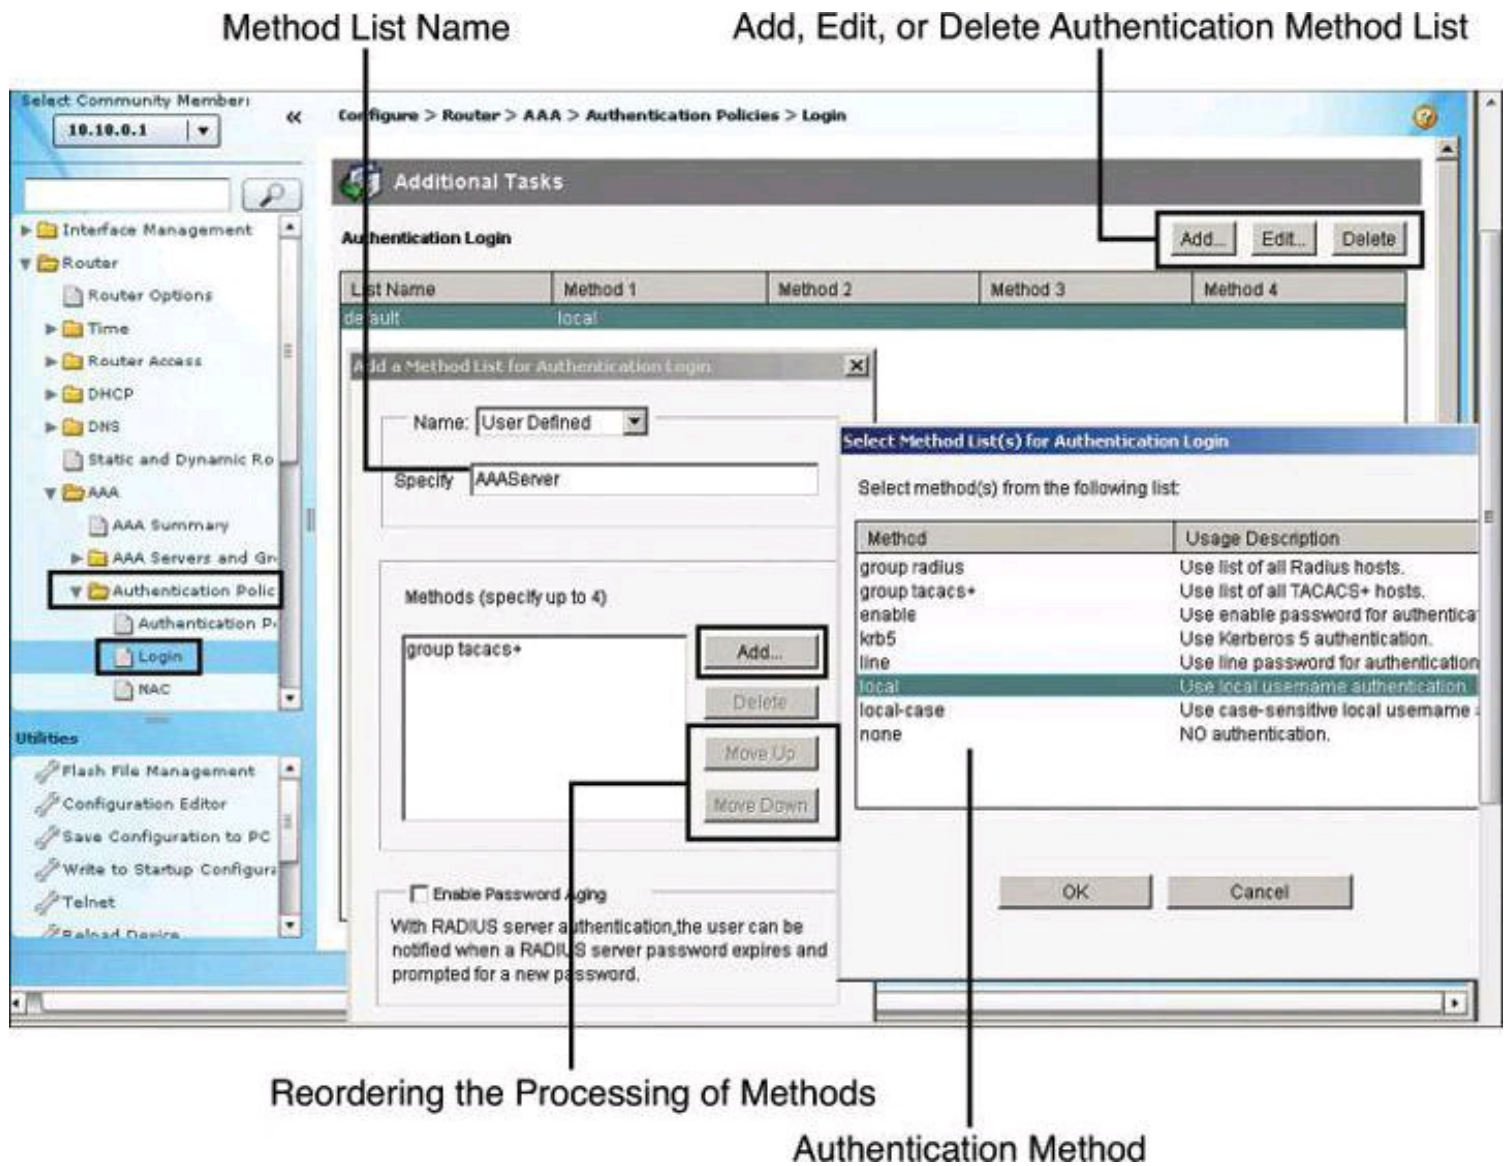


Figure 4-14. Configuring AAA Authentication Method Lists

The method list requires a name so that you can later assign the list to the router function requiring authentication (management access, in our example). This name is specified in the Add a Method List for Authentication Login dialog box, shown on the left in [Figure 4-14](#).

The authentication methods comprise a list, which you can modify by clicking Add to add a new method to the list. The resulting Select Method List(s) for Authentication Login dialog box, shown on the right in [Figure 4-14](#), displays the available authentication methods and allows you to select among these methods. You can specify up to four methods and place them in the list in the order in which you want the router to use them. As explained earlier in conjunction with [Figure 4-9](#), the router will attempt the first method in the list. If the authentication request receives a PASS or a FAIL response, the router does not query further. If the router does not receive a response by using the first method, it uses the next method in the list, and continues to the end of the list until it receives a PASS or a FAIL response. Each method in the list represents, then, a backup or fallback option to the previous method.

You can change the order of processing of the method list in the Add a Method List for Authentication Login dialog box. Click Move Up to move a method up the list. Click Move Down to move a method further down the list. The method “none” will always be last in the list. No other method in the list can be moved below it. This is a Cisco IOS restriction. Cisco IOS will not accept any method name after the method name “none” has been added to a method list. Method name “none”

should be used with care, because it removes the requirement of having to provide authentication.

AAA Authentication CLI Configuration

[Table 4-10](#) describes common options of the **aaa authentication login** command. More options are available than those presented here.

Table 4-10. *aaa authentication login* Command Parameters

Parameter	Description
enable	Uses the enable password for authentication.
line	Uses the line password for authentication.
local	Uses the local username database for authentication.
exec	Creates a method list that provides accounting records about user EXEC terminal sessions on the NAS, including username, date, and start and stop times.
none	Uses no authentication
group-radius	Uses the list of all RADIUS servers for authentication.
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the aaa group server radius command or aaa group server tacacs+ command.

[Example 4-10](#) shows a CLI configuration for local authentication, using the following commands:

- **aaa new-model:** Enables AAA.
- **aaa authentication login default local:** Defines the default method list for login authentication using the local database.
- **aaa authentication login *list name*:** Defines a custom method list for login authentication using the local database, using the enable password as a fallback option. In [Example 4-10](#), the list name is MGT-ACCESS.
- **username:** Adds usernames and passwords to the local security database.
- **login authentication:** Assigns a method list to an access line; applied to the console line in [Example 4-10](#).

Maximum Number of Login Attempts and Login Delay

To further secure administrative access to the router, you can specify the maximum number of failed AAA login attempts that can occur before an account is locked out, using the **aaa local authentication attempts max-fail** command in global configuration mode, as shown in [Example 4-10](#).

Not used in [Example 4-10](#) is the **login delay** command. The **aaa local authentication**

attempts max-fail command differs from the **login delay** command in how it handles failed attempts. The **aaa local authentication attempts max-fail** command locks the user account if the authentication fails. This account stays locked until it is cleared by an administrator. The **login delay** command introduces a delay between failed login attempts without locking the account.

Example 4-10. AAA CLI Configuration Example with Local Authentication

[Click here to view code image](#)

```
R1(config)# aaa new-model
R1(config)# aaa local authentication attempts max-fail 10
R1(config)# aaa authentication login default local
R(config)# aaa authentication login MGT-ACCESS local enable
R1(config)# enable secret SnowyDay2012
R1(config)# username admin privilege 15 view root secret sanfran2012
R1(config)# username FWadmin privilege 10 view CCP_Firewall secret
1StopUn0w
R1(config)# line con 0
R1(config-line)# login authentication MGT-ACCESS
R1(config-line)# end
R1# debug aaa authentication
```

To display information on AAA authentication, use the **debug aaa authentication** command in privileged EXEC command mode. [Example 4-11](#) contains debug output for an unsuccessful AAA authentication followed by a successful authentication on the router console, using the local user database defined in [Example 4-10](#).

Example 4-11. Troubleshooting Using the *debug aaa authentication* Command

[Click here to view code image](#)

```
User Access Verification

Username: wrongusername
Password:

Feb 11 11:06:47.971: AAA/BIND(0000001B): Bind i/f
Feb 11 11:06:47.971: AAA/AUTHEN/LOGIN (0000001B): Pick method list 'MGT-ACCESS'
Feb 11 11:06:48.223: AAA/AUTHEN/ENABLE(0000001B): Processing request
action LOGIN
Feb 11 11:06:48.223: AAA/AUTHEN/ENABLE(0000001B): Done status
GET_PASSWORD
Feb 11 11:06:49.231: AAA/AUTHEN/ENABLE(0000001B): Processing request
action LOGIN
Feb 11 11:06:49.235: AAA/AUTHEN/ENABLE(0000001B): Done status FAIL - bad
password
% Authentication failed
```

```
Username: admin
Feb 11 11:06:51.239: AAA/AUTHEN/LOGIN (0000001B): Pick method list 'MGT-
ACCESS'
Password:

R1>
```

AAA on a Cisco Router Using Cisco Secure ACS

Cisco Secure Access Control Server (ACS) provides a centralized identity networking solution and simplified user management experience across all Cisco devices and security management applications. This section describes Cisco Secure ACS and its uses, the requirements for installing Cisco Secure ACS for Windows, the Cisco Secure ACS installation procedure, and its configuration for router AAA services.

Cisco Secure ACS Overview

Local implementations of AAA, as previously explained, do not scale well. Most corporate environments have multiple Cisco routers and NASs with multiple router administrators and hundreds or thousands of users needing access to the corporate LAN. Maintaining local databases for each Cisco router and NAS for this size of network is not feasible.

To solve this challenge, you can use one or more Cisco Secure ACS systems (servers or engines) to manage the entire user and administrative access needs for an entire corporate network using one or more databases. External AAA systems, such as the Cisco Secure ACS for Windows, Cisco Secure ACS appliance, or Cisco Identity Services Engine (ISE), communicate with Cisco routers and NASs using the TACACS+ or RADIUS protocols to implement AAA functions. This allows you to make changes to user accounts and passwords in a centralized place (the ACS server) and have all the Cisco routers and NASs in your network access this information.

[Figure 4-15](#) shows the following steps of the authentication and authorization process using an external Cisco Secure ACS system to provide AAA services to a network:

Step 1. The client establishes a connection with the router.

Step 2. The router prompts the user for a username and password.

Step 3. The router passes the username and password to the Cisco Secure ACS (server or engine).

Step 4. The Cisco Secure ACS authenticates the user. The user is authorized to access the router (administrative access) or the network based on information found in the Cisco Secure ACS database.




Figure 4-15. Implementing Authentication Using External Servers

Cisco Secure ACS, in its appliance and software form factors, is a highly scalable, high-performance ACS that operates as a centralized RADIUS and TACACS+ server that provides the following features:

- Extends access security by combining authentication, user access, and administrator access with policy control within a centralized identity networking solution
- Allows greater flexibility and mobility, increased security, and user-productivity gains
- Enforces a uniform security policy for all users regardless of how they access the network
- Reduces the administrative and management burden involved in scaling user and network administrator access to the network

Cisco Secure ACS uses a central database, which allows it to centralize the control of all user privileges and enables hundreds or thousands of access points throughout the network to reference those privileges. Cisco Secure ACS provides detailed reporting and monitoring capabilities of user behavior, access connections, and device configuration changes. This feature has become extremely important for organizations trying to comply with Sarbanes-Oxley Act regulations. Cisco Secure ACS supports a broad variety of access connections, including wired and wireless LAN, dialup, broadband, content, storage, VoIP, firewalls, switches, and VPNs.

You can leverage the Cisco Secure ACS framework to control administrator access and configuration for all the network devices in your network that support RADIUS and TACACS+. The following are some of the advanced features of Cisco Secure ACS:

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

Cisco Secure ACS is an important component of the Cisco Identity Based Networking Services (IBNS) architecture. Cisco IBNS is based on port-security standards such as IEEE 802.1x and Extensible Authentication Protocol (EAP), and extends security from the perimeter of the network to every connection point inside the LAN. You can deploy new policy control, such as per-user quotas, VLAN assignments, and ACLs, within this new architecture because of the extended capabilities of Cisco switches and wireless access points to query Cisco Secure ACS over the RADIUS protocol. Please note that IBNS is being superseded by TrustSec. Have a look at TechWiseTV's "Fundamentals of TrustSec" on YouTube to learn more.

Cisco Secure ACS is also an important component of Cisco Network Admission Control (NAC). Cisco NAC is an industry initiative sponsored by Cisco that uses the network infrastructure to enforce security-policy compliance on all devices seeking to access network computing resources, thereby limiting damage from viruses and worms. With NAC, customers can choose to allow network access only to compliant and trusted endpoint devices (for instance, PCs, servers, and personal digital assistants [PDA]) and can restrict the access of noncompliant devices. Cisco NAC is part of the Cisco Self-Defending Network initiative and is the foundation for enabling NAC on Layer 2 and Layer 3 networks. Future phases extend endpoint and network security interoperation to include dynamic incident-containment capabilities. This innovation enables compliant system elements to report misuse emanating from rogue or infected systems during an attack. Thus, infected systems can be dynamically quarantined from the rest of the network to significantly reduce virus, worm, and blended-threat propagation.

Cisco Secure ACS is a powerful access control server with many high-performance and scalability features for any organization growing its WAN or LAN. The following lists the main benefits of Cisco Secure ACS:

- **Ease of use:** A web-based user interface simplifies and distributes the configuration for user profiles, group profiles, and Cisco Secure ACS configuration.
- **Scalability:** Cisco Secure ACS is built to support large networked environments with support for redundant servers, remote databases, and database replication and backup services.
- **Extensibility:** LDAP authentication forwarding supports the authentication of user profiles that are stored in directories from leading directory vendors, including Sun, Novell, and Microsoft.
- **Management:** Microsoft Windows Active Directory support consolidates Windows username and password management and uses the Windows Performance Monitor for real-time statistics viewing.
- **Administration:** Different access levels for each Cisco Secure ACS administrator and the ability to group network devices together make it easier and more flexible to control the enforcement and changes of security policy administration over all of the devices in a network.
- **Product flexibility:** Because Cisco IOS Software has embedded support for AAA, Cisco Secure ACS can be used across virtually any NAS that Cisco sells (the Cisco IOS Software release must support RADIUS or TACACS+). Cisco Secure ACS is now available in an appliance format or in a virtual machine format. Until version 4.x, in addition to being

offered as an appliance, it was also available as software call Cisco Secure ACS for Windows and as a small appliance called Cisco Secure ACS Express, mentioned later in this chapter.

- **Integration:** Tight coupling with Cisco IOS routers and VPN solutions provides features such as Multichassis Multilink PPP and Cisco IOS Software command authorization.
- **Third-party support:** Cisco Secure ACS offers token server support for any one-time password (OTP) vendor that provides an RFC-compliant RADIUS interface, such as RSA, PassGo, Secure Computing, Vasco, or CryptoCard.
- **Control:** Cisco Secure ACS provides dynamic quotas to restrict access based on the time of day, network use, number of logged sessions, and the day of the week.

Cisco Secure ACS for Windows

The last available version of Cisco Secure ACS running on Windows Server was version 4.2. This 4.2 version was also available in two additional form factors: VMware ESX Server, and as a dedicated appliance. In 2011, Cisco announced the end of life of Cisco Secure ACS 4.2 and thus the end of ACS for Windows Server. The new 5.2 version, covered later in the chapter, does not support ACS for Windows Server. Version 5.2 is only available as a VM or as a dedicated appliance. Because Cisco Secure ACS for Windows is still vastly used, we are briefly covering this option even though it has reached end-of-life status.

Cisco Secure ACS 4.2 for Windows must meet certain minimum hardware and operating system requirements, covered briefly in the next section. For third-party software requirements, such as web browsers and Java, consult the release notes.

The server that will be running Cisco Secure ACS must meet the following minimum hardware requirements:

- Pentium IV processor that is 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 800×600 pixels

Cisco Secure ACS 4.2 for Windows supports the English-language versions of the following Microsoft Windows operating systems:

- Windows 2000 Server, with Service Pack 4 installed
- Windows 2000 Advanced Server, with the following conditions:
 - Service Pack 4 installed
 - Without Microsoft Clustering Service installed
 - No other features specific to Windows 2000 Advanced Server enabled, such as Terminal Services
- Windows Server 2003 Service Pack 1, Enterprise Edition or Standard Edition
- Windows Server 2003, R2, Standard Edition
- Windows Server 2003, Service Pack 2

- Windows Server 2003, R2, Service Pack 2

Note

ACS for Windows supports the multiprocessor feature on dual-processor computers. Cisco Secure ACS 4.2 supports the Japanese Windows Server 2003.

You can apply the Windows service packs before or after installing Cisco Secure ACS. If you do not install a required service pack before installing Cisco Secure ACS, the Cisco Secure ACS installation program might warn you that the required service pack is not present. If you receive a service pack message, continue the installation, and then install the required service pack before starting user authentication with Cisco Secure ACS.

Cisco Secure ACS Express

Now at end-of-life status, Cisco Secure ACS Express was an entry-level RADIUS and TACACS+ AAA server for retail branch locations, enterprise branch offices, and small businesses that have fewer than 350 users and 50 devices. It also ran on a Windows platform.

Cisco Secure ACS: Appliance and Virtual Machine

The Cisco Secure Access Control System (the non-Windows version) showcases a physical or virtual appliance form factor. It is a policy-based security server that also provides standards-compliant AAA services to your network. Cisco Secure ACS requires version 5.x of the Cisco Secure ACS software, and facilitates the administrative management of devices and applications manufactured by Cisco and other vendors.

As a dominant enterprise network access control platform, Cisco Secure ACS serves as an integration point for network access control and identity management. Two footprints are available—a Linux-based hardware appliance running the Cisco Secure ACS software, as shown in [Figure 4-16](#), or a virtual server image for VMware ESX environments.



Figure 4-16. Cisco 1121 Secure Access Control System

Cisco Secure ACS provides a rule-based policy model that allows you to control network access based on dynamic conditions and attributes. The rule-based policy is designed to meet complex access policy needs.

Within the greater context of two major AAA protocols—RADIUS and TACACS+—Cisco Secure ACS provides the following basic areas of functionality:

- Under the framework of the RADIUS protocol, Cisco Secure ACS controls the wired and wireless access by users and host machines to the network and manages the accounting of

the network resources used.

- Cisco Secure ACS supports multiple RADIUS-based authentication methods, including the following:
 - Password Authentication Protocol (PAP)
 - Challenge Handshake Authentication Protocol (CHAP)
 - Microsoft CHAP version 1 (MS-CHAPv1)
 - Microsoft CHAP version 2 (MS-CHAPv2)
- Cisco Secure ACS also supports many members of the Extensible Authentication Protocol (EAP) family of protocols:
 - EAP-MD5
 - Lightweight EAP (LEAP)
 - Protected EAP (PEAP)
 - EAP-Flexible Authentication via Secure Tunneling (EAP-FAST)
 - EAP-Transport Layer Security (EAP-TLS)
- In association with PEAP or EAP-FAST, Cisco Secure ACS also supports EAP-MSCHAPv2 and EAP-Generic Token Card (EAP-GTC).
- Under the framework of the TACACS+ protocol, Cisco Secure ACS helps to manage Cisco and other vendor network devices such as switches, wireless access points, routers, and gateways. It also helps to manage services and entities such as dialup, VPN, and firewall.

Cisco Secure ACS is the point in your network that identifies users and devices that try to connect to your network. This identity establishment can occur directly by using the Cisco Secure ACS internal identity repository for local user authentication or indirectly by using external identity repositories. For example, Cisco Secure ACS can use Active Directory as an external identity repository to authenticate a user and grant the user access to the network.

Cisco Secure ACS provides advanced monitoring, reporting, and troubleshooting tools that help you administer and manage your Cisco Secure ACS deployments.

Cisco Secure ACS does the following:

- Enforces access policies for VPN and wireless users
- Provides simplified device administration

Cisco Identity Services Engine

Cisco Identity Services Engine (ISE) is a next-generation identity and access control policy platform that enables enterprises to enforce compliance, enhance infrastructure security, and simplify service operations. Its unique architecture allows enterprises to gather real-time contextual information from networks, users, and devices to make proactive governance decisions by enforcing policy across the network infrastructure. Cisco ISE is an integral component of the Cisco TrustSec solution that helps secure and govern borderless networks. This rack-mounted server runs its own software, which is different from the traditional Cisco Secure ACS software.

Cisco ISE, which uses the same platform as Cisco 1121 (shown previously in [Figure 4-16](#)),

provides a highly powerful and flexible attribute-based access control solution that combines AAA, posture, profiling, and guest management services on a single platform. Administrators can centrally create and manage access control policies for users and endpoints in a consistent fashion, and gain end-to-end visibility into everything that is connected to the network. Cisco ISE automatically discovers and classifies endpoints, provides the right level of access based on identity, and enforces endpoint compliance by checking the posture of a device. Cisco ISE also provides advanced enforcement capabilities, including Security Group Access (SGA) by using security group tags (SGT) and security group ACLs (SGACL).

Cisco ISE is supported in hardware and software footprints, meaning that there are three hardware appliances for small, midsize, and large scenarios, and the ESX-based virtual server footprint.

ISE

At the time of writing, ISE does not yet support TACACS+; it supports only RADIUS.

Refer to TechWiseTV on YouTube for a great video introducing ISE fundamentals:
<http://www.youtube.com/watch?v=sel1F7mKdtI>.

TACACS+ and RADIUS Protocols

The Cisco Secure ACS family of products supports both RADIUS and TACACS+ protocols, which are the two predominant AAA protocols that are used by Cisco security appliances, routers, and switches for implementing AAA.

Cisco Secure ACS 5.2 supports RADIUS for network access control and TACACS+ for network device access control. There are several differences between TACACS+ and RADIUS, leading to their deployment in different scenarios. TACACS+ is the more secure option because it encrypts the complete transaction between the AAA client and the AAA server. It also separates authentication from authorization, making it suitable for device administration, where users are authenticated once, and device administration is then authorized granularly without the need for constant authentication. RADIUS showcases a robust application programming interface (API) for custom accounting, but it combines authentication and authorization requests. This makes it suitable for network access control, but not for device administration.

TACACS+

TACACS+ is a Cisco enhancement to the original TACACS protocol. Despite its name, TACACS+ was designed from the ground up and is therefore incompatible with any earlier version of TACACS. TACACS+ has been submitted to the Internet Engineering Task Force (IETF) as a draft proposal.

TACACS+ provides separate message types for AAA services. Because TACACS+ separates authentication and authorization, it is possible to use TACACS+ authorization and accounting while using another method of authentication.

The extensions to the TACACS+ protocol provide more types of authentication requests and response codes than were in the original specification. TACACS+ offers multiprotocol support, such as IP and AppleTalk. Normal TACACS+ operation encrypts the entire body of the packet for more secure communications and uses TCP port 49.

RADIUS

RADIUS is an open IETF standard AAA protocol for applications such as network access or IP mobility that was developed by Livingston Enterprises. RADIUS works in both local and roaming situations and is commonly used for accounting purposes. RADIUS is currently defined by RFCs 2865, 2866, 2867, and 2868.

The RADIUS protocol hides the passwords during transmission between the NAS and RADIUS server, even with the PAP protocol, using a rather complex operation that involves MD5 hashing and a shared secret. However, the rest of the packet is sent in plaintext.

RADIUS combines authentication and authorization as one process. Once users are authenticated, they are authorized as well. RADIUS uses UDP port 1645 or 1812 for authentication and UDP port 1646 or 1813 for accounting.

In addition, RADIUS is widely used by VoIP service providers. It is used to pass login credentials of a Session Initiation Protocol (SIP) endpoint (such as a broadband phone) to a SIP registrar using digest authentication, and then to a RADIUS server using RADIUS. RADIUS is also a common authentication protocol that is used by the 802.1x security standard.

The Diameter protocol is the planned replacement for RADIUS. Diameter is more secure than RADIUS because it uses the Stream Control Transmission Protocol (SCTP) or TCP rather than UDP. It also provides for failover procedures, and offers a transition path for current RADIUS implementations.

Comparing TACACS+ and RADIUS

There are several differences between TACACS+ and RADIUS, as described in the following list and summarized in [Table 4-11](#).

Table 4-11. TACACS+/RADIUS Comparison

	TACACS+	RADIUS
Functionality	Separates AAA	Combines authentication and authorization
Standard	Mostly Cisco supported	Open/RFC
Transport protocol	TCP	UDP
CHAP	Bidirectional	Unidirectional
Protocol support	Multiprotocol support	No ARA, no NetBEUI
Confidentiality	Entire packet encrypted	Password encrypted
Customization	Provides authorization of router commands on a per-user or per-group basis	Has no options to authorize router commands on a per-user or per-group basis
Accounting	Limited	Extensive

- **Functionality:** TACACS+ separates AAA functions according to the AAA architecture,

allowing modularity of the security server implementation. RADIUS combines authentication and authorization, but separates accounting, thus allowing less flexibility in implementation than TACACS+.

- **Standard:** TACACS+ is a standard that is used mostly by Cisco customers. RADIUS is an open industry standard.
- **Transport protocol:** TACACS+ uses TCP. RADIUS uses UDP, which was chosen for the simplification of client and server implementations; however, it makes the RADIUS protocol less robust and requires the server to implement reliability measures such as packet retransmission and timeouts.
- **Challenge and response:** TACACS+ supports bidirectional challenge and response as used in CHAP between two routers. RADIUS supports unidirectional challenge and response from the RADIUS security server to the RADIUS client.
- **Protocol support:** TACACS+ provides more complete dialup and WAN protocol support. RADIUS does not support AppleTalk Remote Access (ARA), NetBIOS Extended User Interface (NetBEUI), NetWare Access Server Interface (NASI), and X.25 packet assembler/disassembler (PAD) connections.
- **Confidentiality:** TACACS+ encrypts the entire packet body of every packet. RADIUS encrypts only the password attribute portion of the Access-Request packet, which makes TACACS+ more secure.
- **Customization:** The flexibility that is provided in the TACACS+ protocol allows many things to be customized on a per-user basis or per-group basis, including which commands a user can execute on a router. RADIUS lacks this flexibility, and therefore many features that are possible with TACACS+ are not possible with RADIUS.
- **Accounting:** TACACS+ accounting includes a limited number of information fields. RADIUS accounting can contain more information than TACACS+ accounting records, which is the key strength of RADIUS over TACACS+.

AAA on a Cisco Router Using an External Database

Previously in this chapter, you learned how to configure a Cisco router to perform local AAA authentication and authorization. Now let's see how to configure the Cisco router to perform external AAA authentication, authorization, and accounting. By the way, AAA accounting can be performed only using an external database.

Configuration Steps for AAA Using an External Database

Follow the steps presented in this section to configure a Cisco IOS router for AAA using Cisco Secure ACS as an external AAA server. Notice that the configuration is modular and object oriented, where AAA servers are tied together into AAA server groups for ease of management. This allows you to configure multiple servers that will act as backups of each other in a high-availability fashion.

The following steps follow the same configuration approach as the previous example where the local router database was used. AAA needs to be enabled globally on the router, and method lists are configured for authentication and, optionally, for authorization and accounting. When configuring Cisco Secure ACS as an external server, the primary method is TACACS+ or RADIUS, instead of the local database.

Step 1. Globally enable AAA.

Step 2. Configure AAA servers and groups.

Step 3. Enable AAA globally on the router.

Step 4. Configure authentication policies using method lists.

Step 5. (Optional) Configure authorization policies using method lists.

Step 6. (Optional) Configure accounting policies using method lists.

Step 7. Verify the AAA configuration.

Earlier, we learned how to turn on AAA globally with CCP or using the **aaa new-model** configuration command.

AAA Servers and Groups

The next step in our external AAA configuration, shown in [Figure 4-17](#), is to configure AAA server objects. Navigate to **Configure > Router > AAA > AAA Servers and Groups > Servers** to view a snapshot of the information about the AAA servers that the router is configured to use. The IP address, server type, and other parameters are displayed for each server.

Add TACACS+ or RADIUS Server

Global Settings for All Servers

Select Community Member: 10.10.0.1

Configure > Router > AAA > AAA Servers and Groups > Servers

Additional Tasks

AAA Servers Global Settings... Add... Edit...

Server IP	Server Type	Timeout	Parameters
-----------	-------------	---------	------------

Add AAA Server

Server Type: RADIUS

Server IP or Host: 10.10.2.35

Authorization Port: 1645 Accounting Port: 1646

Server-Specific Setup (Optional)

Timeout (seconds):

Configure Key

Current Key: <NONE>

New Key: *****

Confirm Key: *****

OK Cancel Help

Configure IP Address, Timeout, and Key

Figure 4-17. Configuring AAA Server in CCP

Click **Add** to add a TACACS+ or RADIUS server type, configuring the IP address for each server, as shown in [Figure 4-17](#). You can configure timeouts and keys for TACACS+ servers, and define custom authorization and accounting ports for RADIUS servers.

The command equivalent to [Figure 4-17](#) is as follows:

```
R1(config)# radius-server host 10.0.1.10 auth-port 1645 acct-port 1646
key 0 cisco
```

In the preceding command, **cisco** is the RADIUS key.

After you create the servers, you can group them into a server group. Server groups provide the opportunity to modularly combine AAA servers, just defined, in high-availability scenarios. You can add server groups, as shown in [Figure 4-18](#), by navigating to **Configure > Router > AAA > AAA Servers and Groups > Groups**. If you click **Add** to add a new group, you can choose the servers from the list and add them to the group for redundancy.

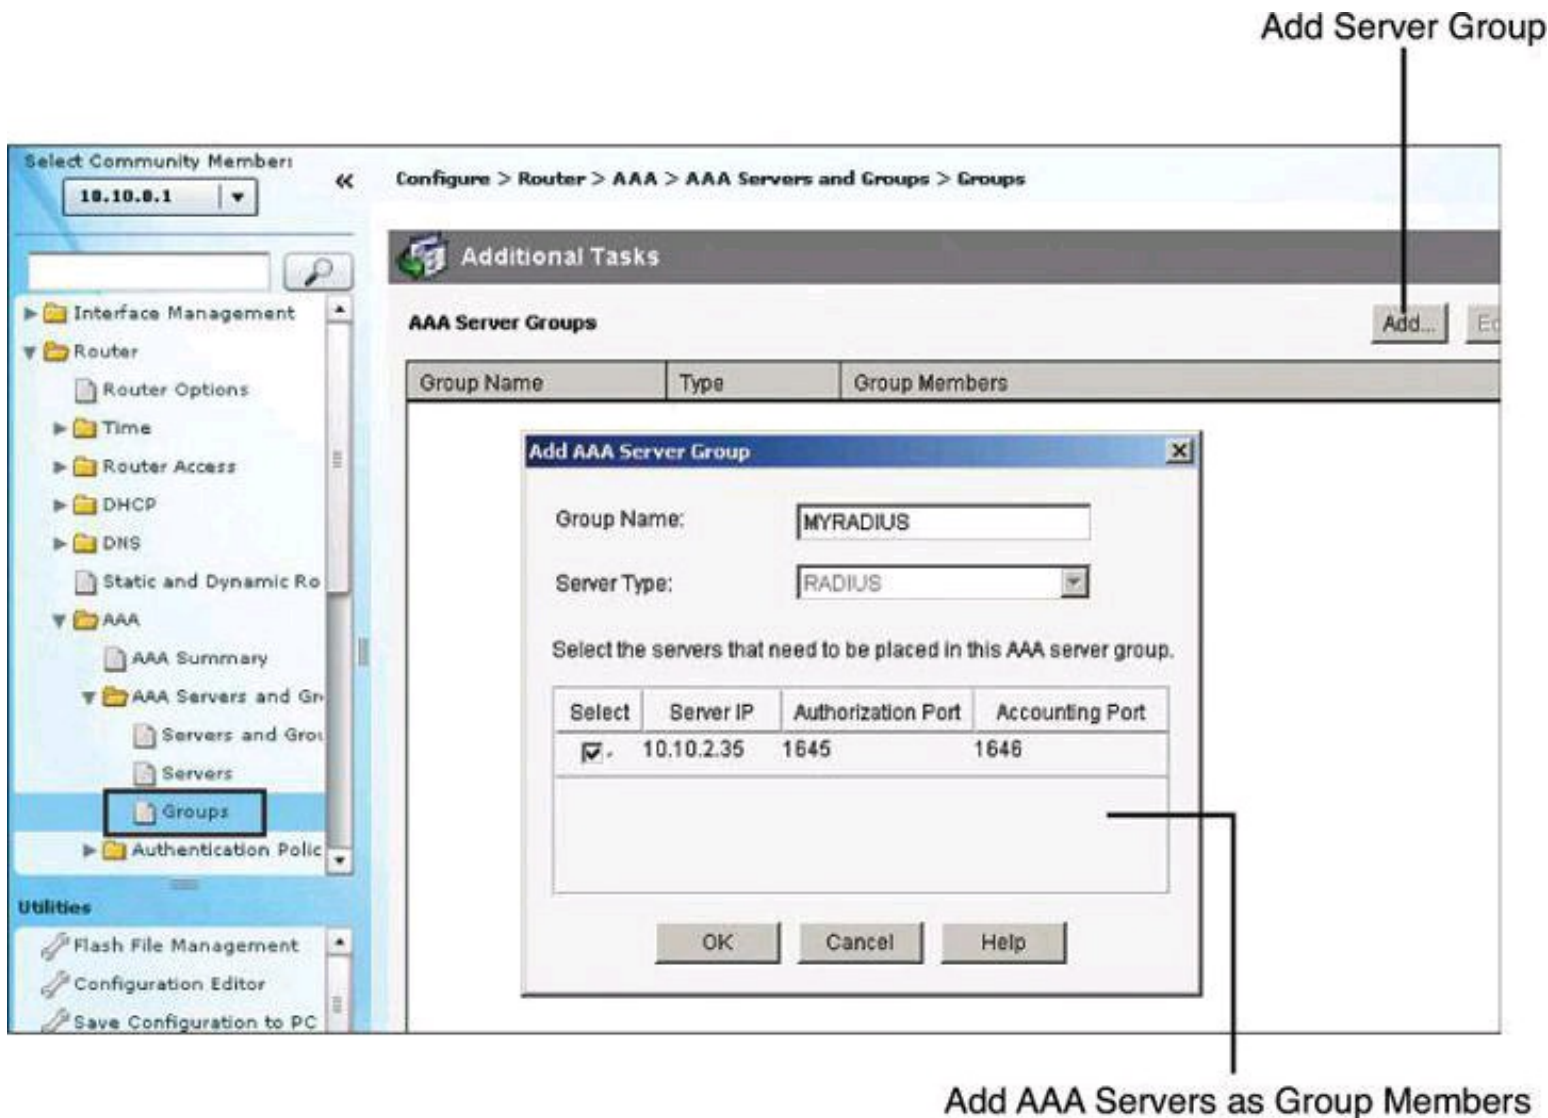


Figure 4-18. Configuring AAA Server Group in CCP

The command equivalent to [Figure 4-18](#) is as follows:

```
R1(config)# aaa group server radius MYRADIUS
R1(config-sg-radius)# server 10.0.1.10 auth-port 1645 acct-port 1646
```

Note

In [Figure 4-18](#), in the Add AAA Server Group dialog box, the Server Type field is grayed out. With the current version of CCP, version 2.6, only the RADIUS server group can be created. However, you can create a TACACS+ server group at the CLI with the following configuration commands:

```
R1(config)# aaa server group tacacs+ MYTACACS
R1(config-sg-tacacs+)# server 10.0.1.67
```

AAA Authentication Method Lists

The rest of the configuration follows the familiar approach of defining named method lists, or selecting the default method list, to include the AAA servers as an option for authentication, authorization, and accounting.

AAA must be enabled globally in order to define AAA method lists. To start AAA, click the button titled **Enable AAA** found by navigating to **Configure > Router > AAA > AAA Summary**. This was demonstrated earlier in [Figure 4-12](#).

Once AAA is enabled globally, you can navigate to **Configure > Router > AAA > Authentication Policies > Login** and click **Add** to add new authentication method lists, as shown in [Figure 4-19](#).

Define a Name for the Method List Add a New List or Edit Existing Lists

Additional Tasks

Authentication Login

List Name	Method 1	Method 2	Method 3	Method 4
default	local			

Add a Method List for Authentication Login

Name: User Defined

Specify: AAAServer

Methods (specify up to 4)

group MYRADIUS	Add...
local	Delete
	Move Up
	Move Down

Enable Password Aging

With RADIUS server authentication, the user can be notified when a RADIUS server password expires and prompted for a new password.

Select Method List(s) for Authentication Login

Select method(s) from the following list:

Method	Usage Description
group radius	Use list of all Radius hosts.
group tacacs+	Use list of all TACACS+ hosts.
group MYRADIUS	Use Server-group name
enable	Use enable password for authentication.
kerb5	Use Kerberos 5 authentication.
line	Use line password for authentication.
none	NO authentication.

OK Cancel

Add Local Method as a Fallback Mechanism to the Group RADIUS Select TACACS+ or RADIUS Methods

Figure 4-19. Configuring Authentication Method Lists

The command equivalent to [Figure 4-19](#) is as follows:

```
R1(config)# aaa authentication login AAAServer group MYRADIUS local
```

Earlier in the chapter, [Figure 4-14](#) showed the Cisco Configuration Professional window that is used to define AAA method lists and authentication methods. When using external AAA servers such as Cisco Secure ACS, you simply select TACACS+ or RADIUS server groups as your main method. It is recommended to add fallback methods, most typically the local router database, in case communication to the AAA servers fails, also shown in [Figure 4-14](#).

AAA Authorization Policies

Authorization and accounting rules follow a similar approach using method lists to define the location of the authorization permissions and the accounting logs.

Because the TACACS+ protocol allows you to separate authentication from authorization, you can configure a router to restrict the user to be able to perform only certain functions after successful authentication. You can configure authorization for both character mode (EXEC authorization) and packet mode (network authorization).

To configure the router to use the Cisco Secure ACS server for authorization, you must create a user-defined authorization method list or edit the default authorization method list. The default authorization method list is automatically applied to all interfaces except those that have a user-defined authorization method list explicitly applied. A user-defined authorization method list overrides the default authorization method list.

Tip

To avoid locking yourself out of the router, make sure you configure authorization on the Cisco Secure ACS server *before* you configure the router for authorization.

Also as a precaution, consider logging on to the router console in privilege mode before starting the authorization configuration.

It is also recommended that you do *not* enable authorization on the backdoor port (for example, the console port) until you have appropriately confirmed the alternative authentication method, such as login via the AAA server, is working. In absence of console access, you may also simply open a second, concurrent Telnet/SSH session to confirm that you have appropriate access to the device, while still keeping the original session, used to log in, active. Once you have access to the device through the second vty session, you should save the configuration. A final resort, albeit not often desired, is to do a “reload in 5” and answer NO when asked to save the configuration. If the authentication/authorization mechanism fails, it will reload the router without saving in 5 minutes (adjustable time).

[Figure 4-20](#) illustrates using an authorization method list for authorizing the use of the CLI with TACACS+. Notice that the local database is again used as a fallback option.

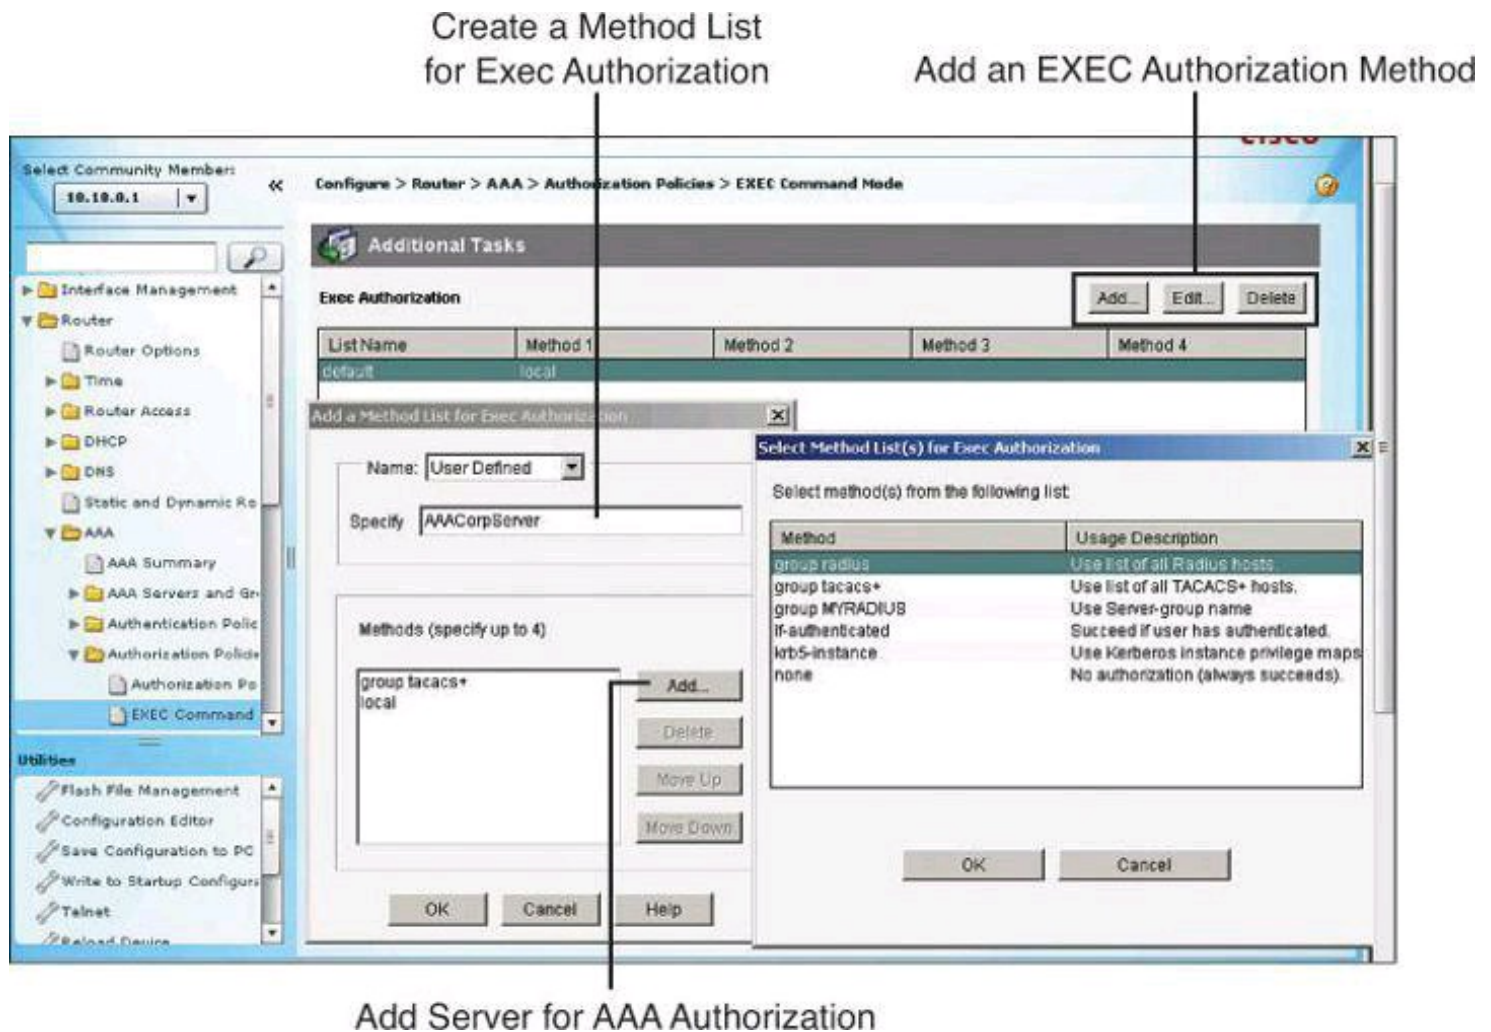


Figure 4-20. Configuring Authorization Method Lists

AAA Accounting Policies

Cisco Secure ACS serves as a central repository for accounting information, essentially tracking events that occur on the network. Each session that is established through Cisco Secure ACS can be fully accounted for and stored on the server. This stored information can be very helpful for management, security audits, capacity planning, and network-usage billing.

Like authentication and authorization method lists, method lists for accounting define the way accounting will be performed and the sequence in which these methods are performed. The default accounting method list is automatically applied to all interfaces except those that have a named accounting method list explicitly defined. A defined accounting method list overrides the default accounting method list.

AAA supports six different types of accounting: network, connection, exec, system, command, and resource.

Currently, AAA accounting can be configured only via the CLI. Interestingly, CCP v2.6 has an AAA Accounting Policies window. However, it does not have Add, Edit, and Delete buttons, so we can't add accounting policies using the GUI. It must be done at the CLI, and even then, it still doesn't appear in the AAA Accounting Policies window in CCP v2.6.

To configure AAA accounting using named method lists, use the commands shown in [Table 4-12](#) while in global configuration mode.

Table 4-12. AAA Accounting Using Named Method Lists Procedure

Step	Command	Notes
1	Router(config)# aaa accounting { system network exec connection commands level } { default <i>list-name</i> } { start-stop stop-only none } [<i>method1</i> [<i>method2...</i>]]	This command creates an accounting method list and enables accounting. The argument <i>list-name</i> is a character string used to name the list you are creating as an account methods. To edit the default method list, use the default parameter. <i>level</i> refers to a specific privilege level, such as level 15. <i>method</i> refers to local , group tacacs+ , group , and so on.
2	Router(config)# line [aux console tty vty] <i>line-number</i> [<i>ending-line-number</i>] -Or- Router(config)# interface <i>interface-type</i> <i>interface-number</i>	This command enters the line configuration mode or interface configuration mode for the lines or interface to which you want to apply the accounting method list.
3	Router(config-line)# accounting { arap commands level connection exec } { default <i>list-name</i> } -Or- Router(config-if)# ppp accounting { default <i>list-name</i> }	This command applies the accounting method list to a line or set of lines, or to an interface or set of interfaces.

Named accounting method lists are specific to the indicated type of accounting. The AAA accounting parameters shown in [Table 4-13](#) describe the types of accounting that can use named accounting method lists. The **aaa accounting** command enables you to specify how much information to record for accounting. [Table 4-13](#) also describes the AAA accounting record types.

Table 4-13. AAA Accounting Command Parameters

Parameter	Description
system	Performs accounting for all system-level events not associated with users, such as reloads. Note that when system accounting is used and the accounting server is unreachable at system startup time, the system will not be accessible for approximately 2 minutes.
network	This parameter creates a method list to enable accounting for all network-related service requests, including SLIP, PPP, PPP NCP, and ARAP protocols.
exec	This parameter creates a method list that provides accounting records about user EXEC terminal sessions on the NAS, including username, date, and start and stop times.
connection	This parameter creates a method list that provides accounting information about all outbound connections made from the NAS.
commands	This parameter creates a method list that provides accounting information about specific, individual EXEC commands associated with a specific privilege level.
default	This parameter uses the listed accounting methods that follow this keyword as the default list of methods for accounting services.
start-stop	This parameter instructs the TACACS+ server to send a start accounting notice at the beginning of the requested event and a stop accounting notice at the end of the event.
stop-only	This parameter instructs the TACACS+ server to send a stop record accounting notice at the end of the requested user process.
none	This parameter instructs the TACACS+ server to stop all accounting activities on this line or interface.

Note

System accounting provides information about all system-level events, such as when the system reboots or when accounting is turned on or off. System accounting does not use named method lists. For system accounting, you can define only the default method list.

AAA Configuration for TACACS+ Example

[Example 4-12](#) shows the resulting running configuration of a router that has been configured for TACACS+ services using Cisco Configuration Professional and CLI commands. This example is unrelated to the RADIUS configuration we have been doing in the preceding sections.

Example 4-12. Example of AAA Configuration for TACACS+

[Click here to view code image](#)

```
aaa new-model
!
aaa authentication login TACACS_SERVER group tacacs+ local
aaa authorization exec default group tacacs+
aaa authorization network default group tacacs+
aaa accounting exec default start-stop tacacs+
aaa accounting network default start-stop tacacs+
aaa accounting commands 15 default stop-only group tacacs+
!
!
tacacs-server host 10.0.1.11
tacacs-server key ciscosecure
!
line vty 0 4
  login authentication TACACS_SERVER
```

The following is an explanation of the commands displayed in [Example 4-12](#):

- **aaa new-model**: Enables AAA.
- **aaa authentication login TACACS_SERVER group tacacs+ local**: Defines an AAA login policy entitled TACACS_SERVER that uses TACACS+ as the first authentication method and the local database as a second method if TACACS+ is unavailable. Note that the command could have been entered as **aaa authentication login TACACS_SERVER tacacs+ local**, but it would have appeared in the config as **aaa authentication login TACACS_SERVER group tacacs+ local**.
- **aaa authorization exec default group tacacs+**: Defines an AAA authorization policy that uses TACACS+ for access to an EXEC prompt. Note that the command could have been entered as **aaa authorization exec default tacacs+**, but it would have appeared in the config as **aaa authorization exec default group tacacs+**.
- **aaa authorization network default group tacacs+**: Defines an AAA authorization policy that uses TACACS+ for network access. Note that the command could have been entered as **aaa authorization network default tacacs+**, but it would have appeared in the config as **aaa authorization network default group tacacs+**.
- **aaa accounting exec default start-stop group tacacs+**: Defines an AAA accounting policy that uses TACACS+ for logging both start and stop records for user EXEC terminal sessions. Note that the command could have been entered as **aaa accounting exec default start-stop tacacs+**, but it would have appeared in the config as **aaa accounting exec default start-stop group tacacs+**.
- **aaa accounting network default start-stop tacacs+**: Defines an AAA accounting policy that uses TACACS+ for logging both start and stop records for all network-related service requests. Note that the command could have been entered as **aaa accounting network default start-stop tacacs+**, but it would have appeared in the config as **aaa accounting network default start-stop group tacacs+**.
- **aaa accounting commands 15 default stop-only group tacacs+**: Defines a default commands accounting method list, where accounting services are provided by a TACACS+ security server, set for privilege level 15 commands with a stop-only restriction. Note that

the command could have been entered as **aaa accounting commands 15 default stop-only tacacs+**, but it would have appeared in the config as **aaa accounting commands 15 default stop-only group tacacs+**.

- **tacacs-server host 10.0.1.11**: Configures the IP address of the TACACS+ server.
- **tacacs-server key ciscosecure**: Configures an encryption key of ciscosecure to be used when communicating with the TACACS+ server.
- **line vty 0 4**: Enters line configuration mode for vty 0 through vty 4.
- **login authentication TACACS_SERVER**: Applies the AAA authentication policy named TACACS_SERVER to all five vty lines.

Troubleshooting TACACS+

Because we are still dealing with how to configure the router, let's review the debug commands available to us. Later in this chapter, you learn the configuration that must be done on the Cisco Secure ACS server itself.

In [Example 4-11](#), we saw the use of the **debug aaa authentication** command to get a high-level view of login activity. When the TACACS+ protocol is used on the router, you can also use the **debug tacacs** command for more detailed debugging information.

[Example 4-13](#) shows sample output from the **debug tacacs** command for a TACACS+ login attempt that was unsuccessful, as indicated by the status FAIL.

Example 4-13. *debug tacacs* Command for an Unsuccessful TACACS+ Login Attempt

[Click here to view code image](#)

```
Router# debug tacacs

13:53:35: TAC+: Opening TCP/IP connection to 192.168.60.15 using source
192.48.0.79
13:53:35: TAC+: Sending TCP/IP packet number 416942312-1 to
192.168.60.15 (AUTHEN/
START)
13:53:35: TAC+: Receiving TCP/IP packet number 416942312-2 from
192.168.60.15
13:53:35: TAC+ (416942312): received authen response status = GETUSER
13:53:37: TAC+: send AUTHEN/CONT packet
13:53:37: TAC+: Sending TCP/IP packet number 416942312-3 to
192.168.60.15 (AUTHEN/
CONT)
13:53:37: TAC+: Receiving TCP/IP packet number 416942312-4 from
192.168.60.15
13:53:37: TAC+ (416942312): received authen response status = GETPASS
13:53:38: TAC+: send AUTHEN/CONT packet
13:53:38: TAC+: Sending TCP/IP packet number 416942312-5 to
192.168.60.15 (AUTHEN/
CONT)
13:53:38: TAC+: Receiving TCP/IP packet number 416942312-6 from
192.168.60.15
13:53:38: TAC+ (416942312): received authen response status = FAIL
```

To display information from the TACACS+ helper process, use the **debug tacacs events** command in privileged EXEC mode. To disable debugging output, use the **no** form of this command.

Note

Use the **debug tacacs events** command with caution, because it can generate a substantial amount of output.

Also, because console ports communicate at 9600 bauds, it is recommended to send debug output to the syslog server and enter **no logging console** on the router.

[Example 4-14](#) shows sample output from the **debug tacacs events** command. The example shows the opening and closing of a TCP connection to a TACACS+ server, the bytes read and written over the connection, and the TCP status of the connection.

Example 4-14. *debug tacacs events* Command Output

[Click here to view code image](#)

```
Router# debug tacacs events
```

```
%LINK-3-UPDOWN: Interface Async2, changed state to up
00:03:16: TAC+: Opening TCP/IP to 192.168.58.104/1049 timeout=15
00:03:16: TAC+: Opened TCP/IP handle 0x48A87C to 192.168.58.104/1049
00:03:16: TAC+: periodic timer started
00:03:16: TAC+: 192.168.58.104 req=3BD868 id=-1242409656 ver=193
handle=0x48A87C
  (ESTAB)
expire=14 AUTHEN/START/SENDAUTH/CHAP queued
00:03:17: TAC+: 192.168.58.104 ESTAB 3BD868 wrote 46 of 46 bytes
00:03:22: TAC+: 192.168.58.104 CLOSEWAIT read=12 wanted=12 alloc=12
got=12
00:03:22: TAC+: 192.168.58.104 CLOSEWAIT read=61 wanted=61 alloc=61
got=49
00:03:22: TAC+: 192.168.58.104 received 61 byte reply for 3BD868
00:03:22: TAC+: req=3BD868 id=-1242409656 ver=193 handle=0x48A87C
(CLOSEWAIT)
  expire=9
AUTHEN/START/SENDAUTH/CHAP processed
00:03:22: TAC+: periodic timer stopped (queue empty)
00:03:22: TAC+: Closing TCP/IP 0x48A87C connection to
192.168.58.104/1049
00:03:22: TAC+: Opening TCP/IP to 192.168.58.104/1049 timeout=15
00:03:22: TAC+: Opened TCP/IP handle 0x489F08 to 192.168.58.104/1049
00:03:22: TAC+: periodic timer started
00:03:22: TAC+: 192.168.58.104 req=3BD868 id=299214410 ver=192
handle=0x489F08
  (ESTAB)
expire=14 AUTHEN/START/SENDPASS/CHAP queued
00:03:23: TAC+: 192.168.58.104 ESTAB 3BD868 wrote 41 of 41 bytes
```

```
00:03:23: TAC+: 192.168.58.104 CLOSEWAIT read=12 wanted=12 alloc=12
got=12
00:03:23: TAC+: 192.168.58.104 CLOSEWAIT read=21 wanted=21 alloc=21
got=9
00:03:23: TAC+: 192.168.58.104 received 21 byte reply for 3BD868
00:03:23: TAC+: req=3BD868 id=299214410 ver=192 handle=0x489F08
(CLOSEWAIT)
    expire=13
AUTHEN/START/SENDPASS/CHAP processed
00:03:23: TAC+: periodic timer stopped (queue empty)
```

Note

The TACACS messages are intended to be self-explanatory to IT service personnel only.

Deploying and Configuring Cisco Secure ACS

Cisco Secure ACS uses a modular approach to define AAA policies. Cisco Secure ACS 5.2 offers a new paradigm. In previous versions of ACS, the policy is group based. A user inherits network privileges and restrictions based on the group to which the user belongs. As an example, Mary from Marketing is limited to a single remote-access session, while Indy from IT is provided with up to three concurrent remote-access sessions. ACS 5.2 provides much more granular policies with its rule-based approach (not to be confused with Role-Based CLI Access). With rules-based policies, the ACS 5.2 looks at the environment used by the user: for example, the user is Indy, connecting at 2:45 p.m., from network device Switch B, port 19, and so forth. So, ACS 5.2 looks at not only who the user is, but also from which network device the user's session is initiating, what time of day the user is connecting, and so forth. In the next two sections, we will examine how prior versions of Cisco Secure ACS operate, and then we will see how the new ACS 5.2 works.

Evolution of Authorization

Cisco made a significant change in authorization when it introduced Cisco ACS 5.2

Before: Group-Based Policies

As mentioned, earlier versions of Cisco Secure ACS base policy on membership in user groups. The user groups define the access restrictions and permissions for the users who are members of the group.

Using the group-based approach, as shown in [Figure 4-21](#), when a user requests access, the credentials of the user are authenticated with an identity store. Once authenticated, the user is associated with the appropriate user group. The authorization rules in this environment are part of the group itself, including restrictions such as time of day or point of access into the network.

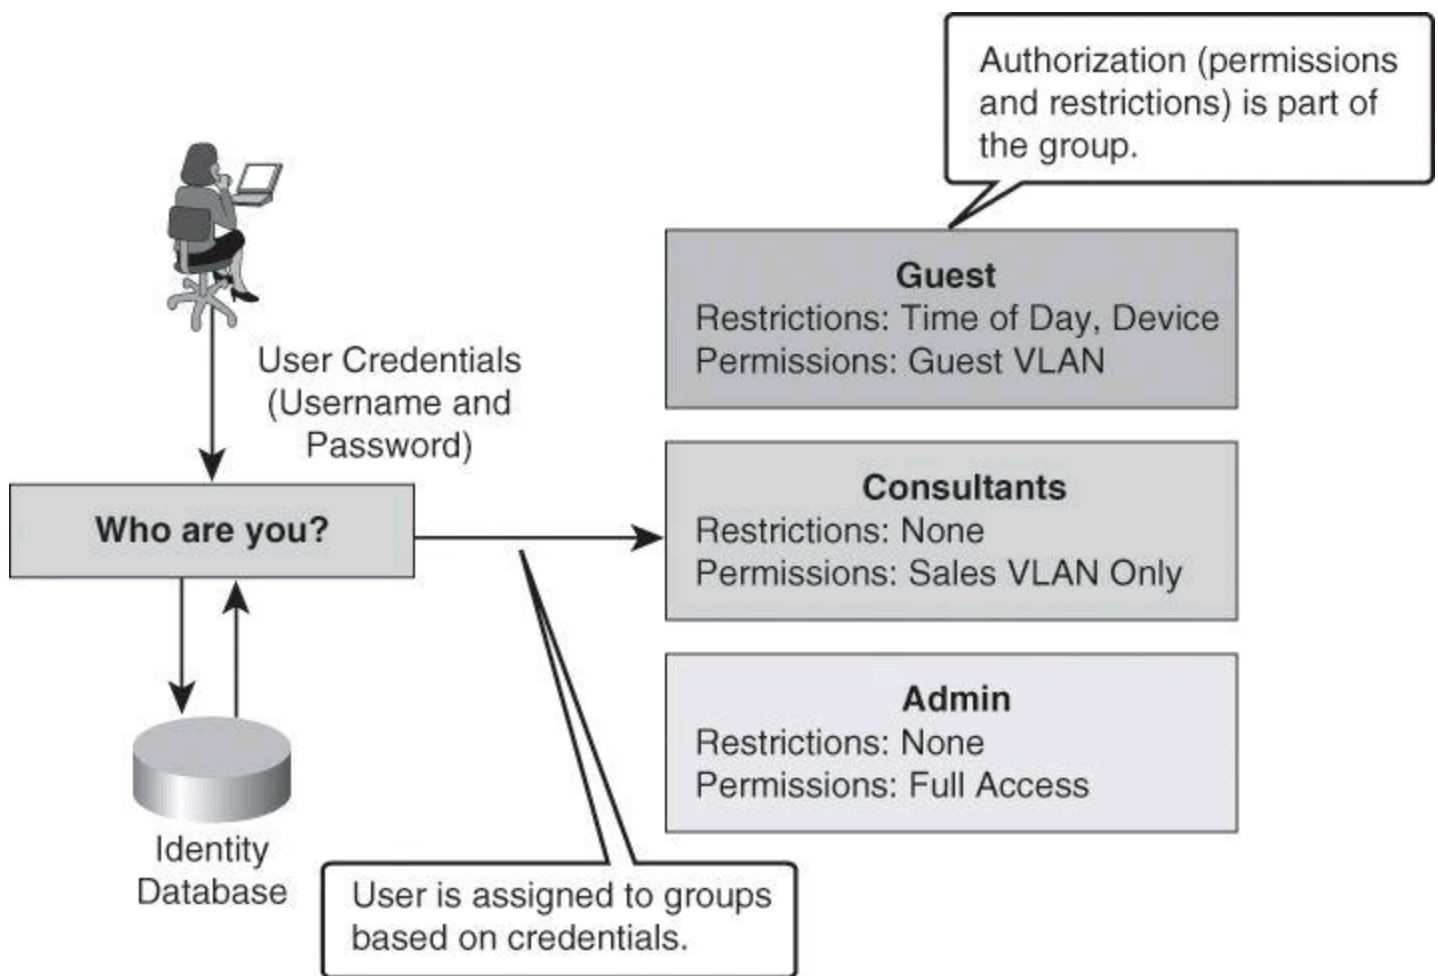


Figure 4-21. Previous Version of Cisco Secure ACS Using Group-Based Policies

Using group-based policies, permissions are set based primarily on the identity-based association of a user with a single group. Because authorization is tied to the user group, all members of the user group have the same access restrictions and permissions all the time.

This type of authorization is suitable for simple policies, in which identity is the dominant or only condition, but it results in lack of flexibility in the presence of other requirements or conditions that define the permissions of the user.

The trends in network connectivity today call for universal access from multiple locations, at multiple times, using multiple devices, as exemplified in [Figure 4-22](#). For example, you might want to grant an employee full access when they are on campus, but restricted access when they are working remotely over the weekend. The group-based policy approach results in complex access policies and rules because you must consider the user's choice of device (laptop versus PDA), the security posture of the device at the time of connection, and many other conditions.

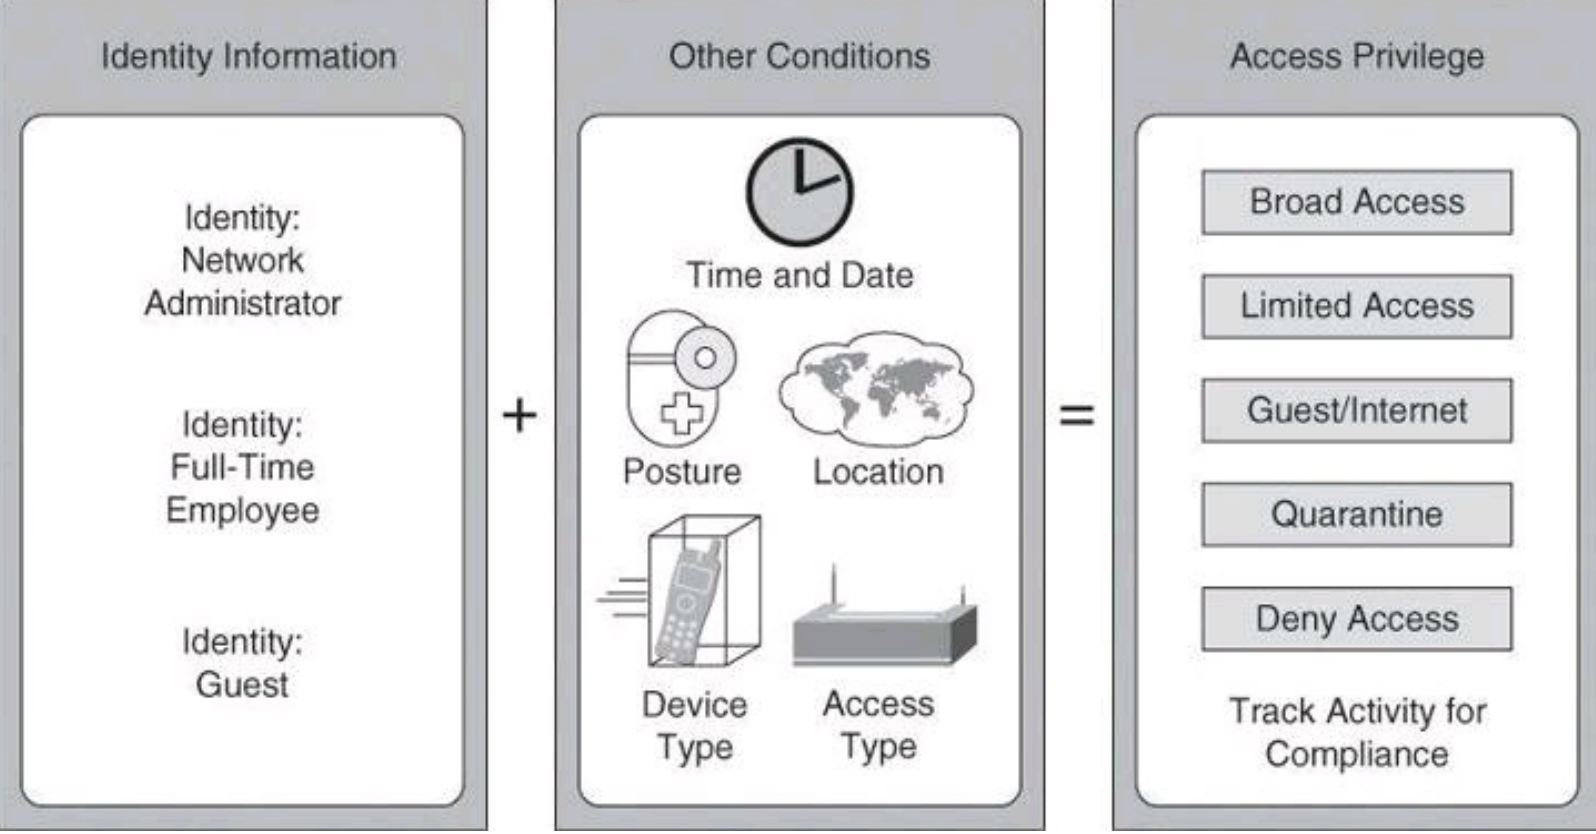


Figure 4-22. Complexity of a Mobile Workforce and Borderless Networks

Now: More Than Just Identities

To manage situations such as those shown in [Figure 4-22](#), a different approach is needed—one that grants permissions on conditions other than identity alone. Using Cisco Secure ACS, starting with version 5.x, group-based policies are replaced by rule-based policies, which provide a more flexible approach that can match the variety of access conditions that are found in current networks.

Using the rule-based approach, instead of relying on the group to contain all the information, identity classification is separated from other conditions or restrictions, as shown in [Figure 4-23](#). Permissions are also decoupled from the group, resulting in the group acting as an identity classification only and no longer containing access permissions and restrictions.

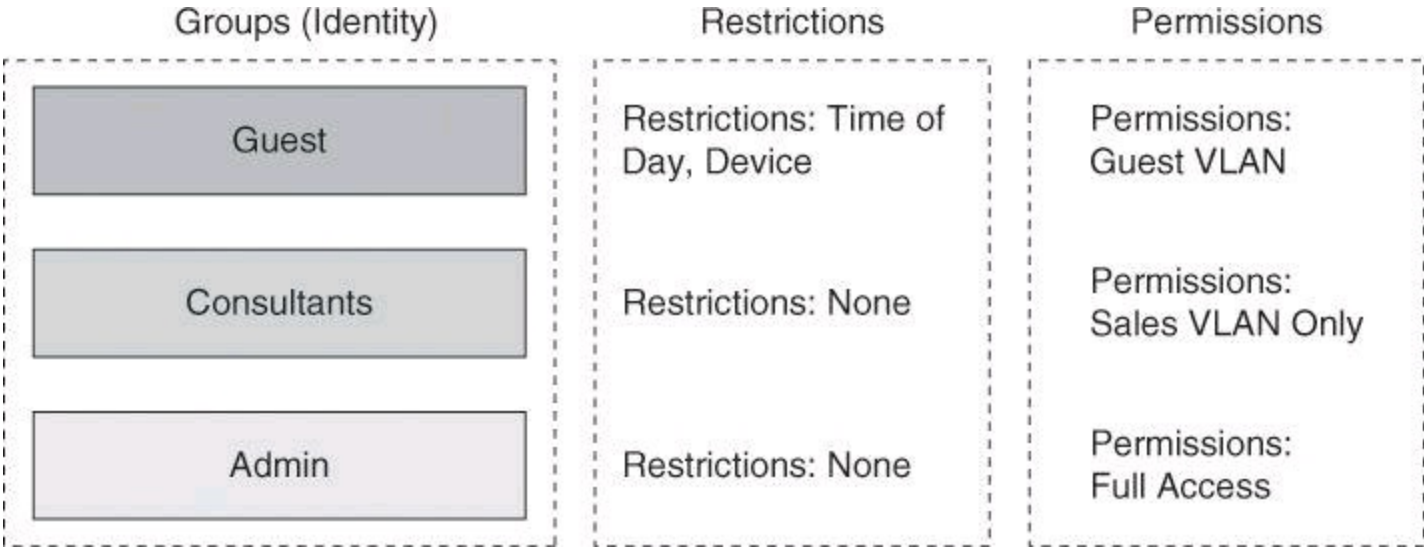


Figure 4-23. New in Cisco Secure ACS 5.2: Rule-Based Policies

In addition to an expanded concept of identity, which is defined by the users and group membership, the idea of restrictions was also expanded to consider other session and environment attributes, such as access location, access type, time, date, end-station health, and so forth.

These attributes become a contextual set of conditions that determines which permissions to apply to the network access request. In other words, the previous approach of associating permissions directly at the user and group level is replaced by a new approach in which the permissions are a separate object called an *authorization profile*. In addition, contextual conditions can be evaluated separately, eventually resulting in different sets of permissions for the same user, based on the different access conditions presented in [Figure 4-24](#).

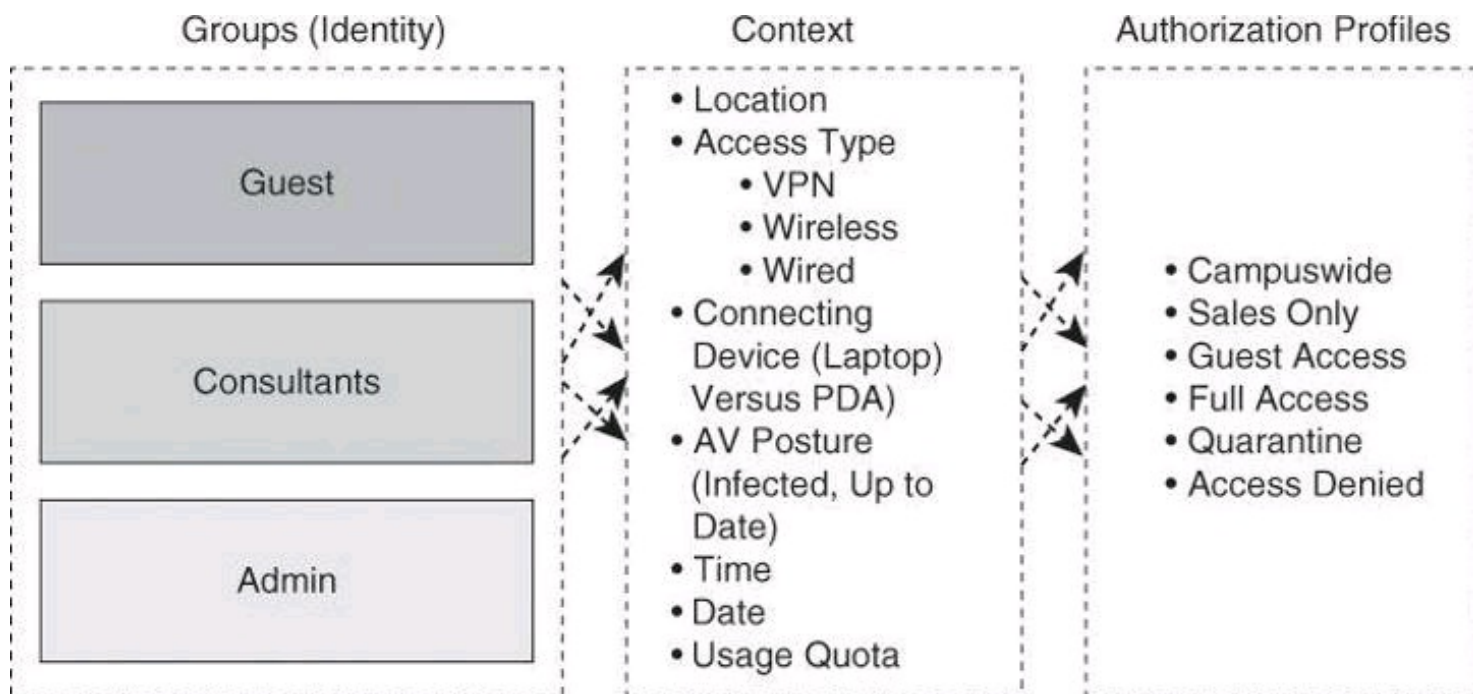


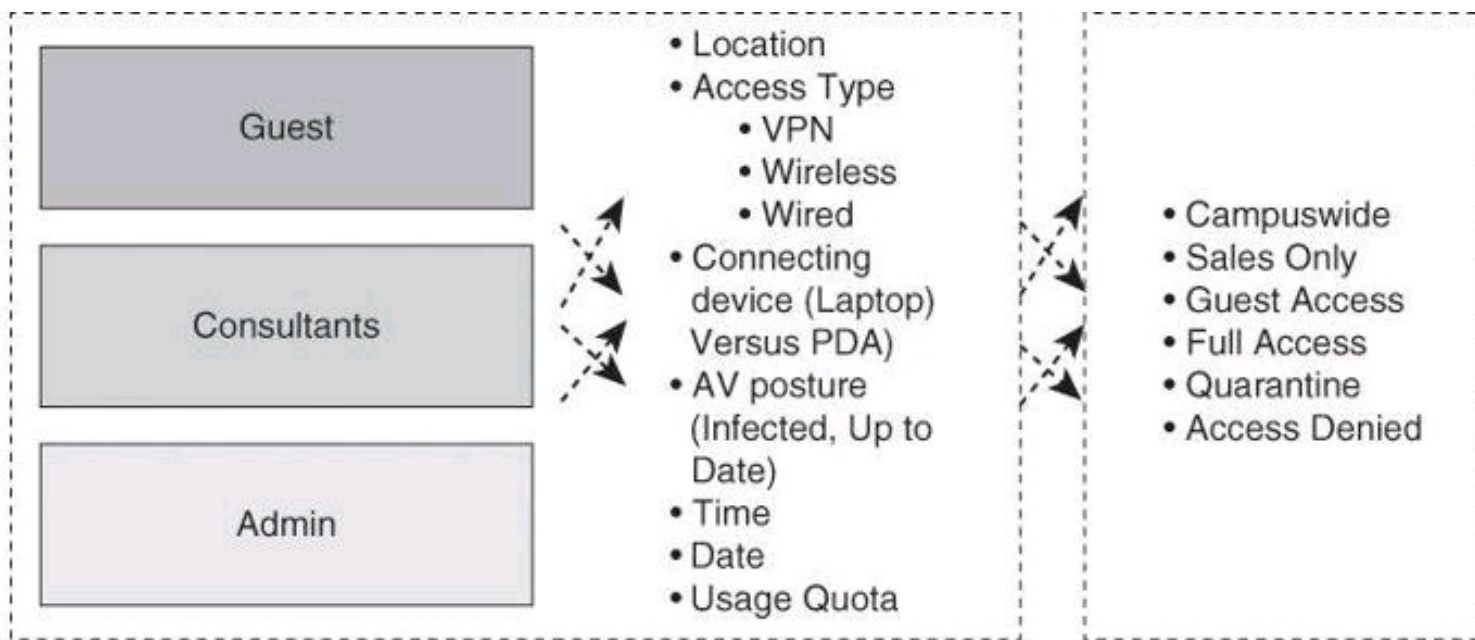
Figure 4-24. Context-Aware Authorization Profiles

Dynamic Access Policies on Cisco ASA

For those readers familiar with Cisco ASA, you will relate [Figure 4-24](#) to the dynamic access policies (DAP) used in VPN settings, where the ASA assigns a policy to an incoming connection based on many criteria, including not only the user's identity, but also how the computer is used to connect, whether the computer is a corporate asset, whether the computer has antivirus protection, and so forth.

Rule-Based Policies

Authorization is now determined using a set of logical rules, which resemble the structure of if-then-else conditions, that better match the variety of network access conditions and resulting permissions. [Figure 4-25](#) illustrates this logic. If access conditions match a certain set of attributes for the user or group, then apply a set of permissions. Using rule-based policies, the conditions can be any combination of available session attributes, including, but certainly not limited to, identity.



RULES: If <condition> then <apply profile>

Policy =
Set of Rules

Conditions			Profile
Group	Time	Device	
Consultant	Weekdays	Corp Laptop	Production VLAN
Consultant	Weekends	Home PC	Guest VLAN
Admin	All	Corp.	Production and Admin VLANs

Figure 4-25. Rule-Based Policies with Cisco Secure ACS 5.2

As previously mentioned, the permissions in Cisco Secure ACS v5.x are defined in another modular object known as an authorization profile. Several profiles can be created, and multiple sets of conditions can arrive at the same profile. For instance, a user can acquire the Guest role if the user does not exist in the identity database, but a user can also acquire this role if the user is valid but tries to access the network from an unknown location, using a smartphone, over the weekend.

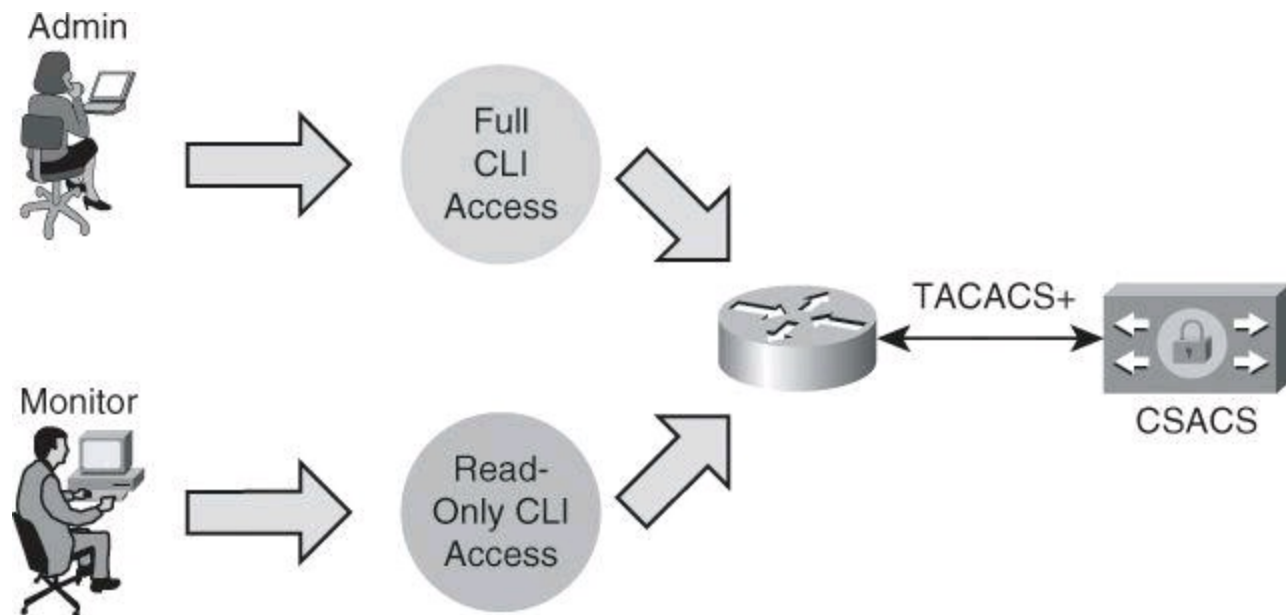
Authorization profiles can contain permissions in multiple dimensions; for instance, VLANs, downloadable ACLs (dACL), quality of service (QoS) settings, RADIUS attributes, and others.

The example in [Figure 4-25](#) highlights the flexibility of the rule-based policies. Notice that a single classification of users can get different authorizations based on more than just the identity attribute. The benefits of this are that the rule conditions can be any combination of the access session attributes, and that the assignment of permissions does not have to be based on a specific identity at all. It is a set of conditions, including identity, that defines access control.

Configuring Cisco Secure ACS 5.2

Configuration of Cisco Secure ACS 5.2 will be explained using the following scenario, illustrated

in [Figure 4-26](#), which focuses on authorization policies for device administration. In our scenario, Cisco Secure ACS will be used as the policy server to grant the admin user full access to the router CLI, at privilege level 15. Similarly, the monitor user will obtain privilege level 1 for read-only access to the router's CLI. The system uses TACACS+ between the router and the policy server. Obviously, this scenario does not present the full potential of Cisco Secure ACS 5.x, which is beyond the scope of this book. The goal in this book, which focuses on Cisco IOS security, is to show you a basic configuration of Cisco Secure ACS 5.2 when it is configured for basic interaction with your network devices.



- Grant administrator user full access to router CLI at privilege level 15.
- Grant monitor user read-only access to router CLI at privilege level 1.

Figure 4-26. Rule-Based Policies Scenario

Configuring Authorization Policies for Device Administration

The rules-based approach is easily configured using the Cisco Secure ACS GUI. The following steps are required to build a device administration access policy using Cisco Secure ACS 5.2. This policy will define the type of CLI access that different user roles have to a router or group of routers.

The general steps used to configure Cisco Secure ACS for device administration are as follow:

- Step 1.** Add network devices to Cisco Secure ACS so the router appears as an AAA client.
- Step 2.** Define identity groups and identity the store.
- Step 3.** Configure access services to process requests.
- Step 4.** Create the identity policy.
- Step 5.** Create the authorization policy.

Adding Network Devices to Cisco Secure ACS

The first step in adding a network device to Cisco Secure ACS is to define network device groups and add routers and network devices to the groups. This step allows Cisco Secure ACS to know which devices will send AAA requests.

In Cisco Secure ACS, you can define network device groups (NDG), which are sets of devices. These NDGs provide logical grouping of devices—for example, by device location or type—which you can use in policy conditions, as shown in [Figure 4-27](#).

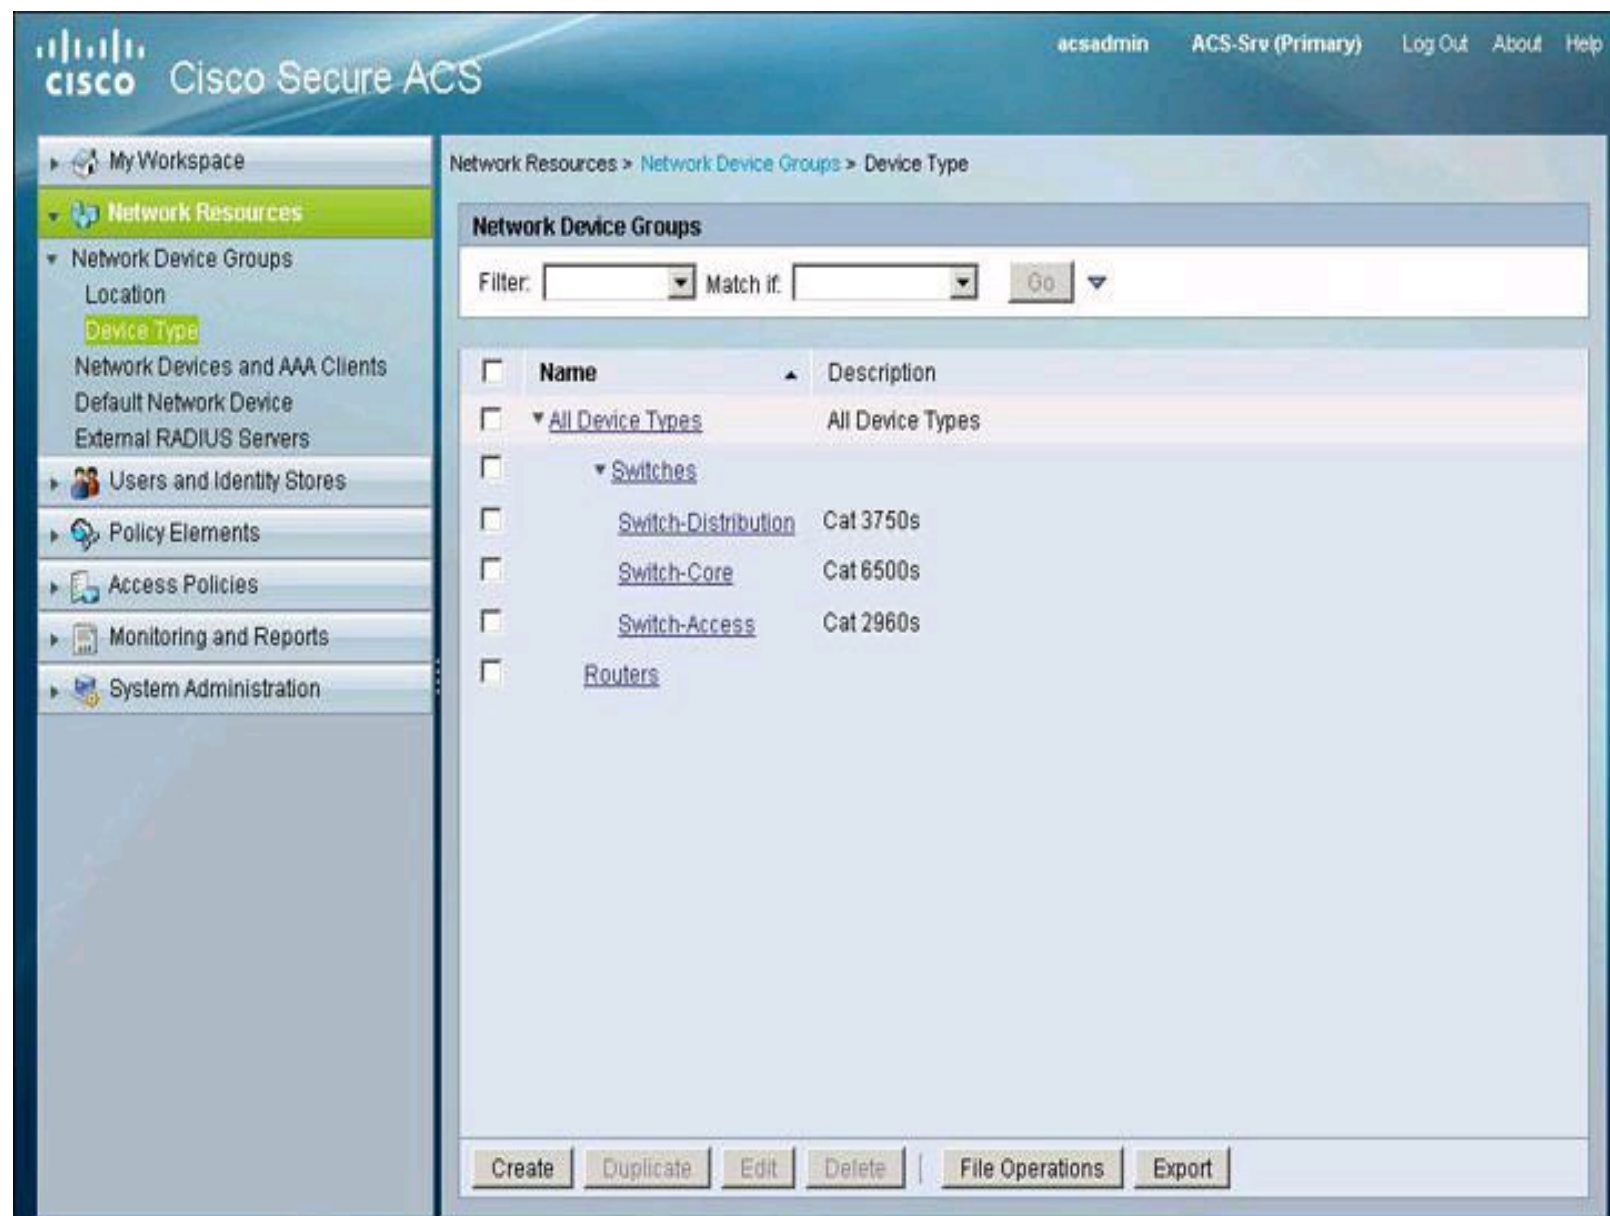


Figure 4-27. Cisco Secure ACS and the Logical Grouping of Devices

When Cisco Secure ACS receives a request for a device, the NDGs that are associated with that device are retrieved and compared against those in the policy table. With this method, you can group multiple devices and assign them the same policies. For example, you can group all devices in a specific location and assign to them the same policy.

The device group hierarchy is the hierarchical structure that contains the NDGs. The hierarchy provides further flexibility to create groups of devices that follow a functional classification, in addition to a geographic or location-based classification.

The access the configuration options for this step, navigate to **Network Resources > Network Device Groups** in the menu on the left side of the screen. As you can see in [Figure 4-27](#), there are two types of NDG: one based on location and one based on device type. In this example, we are creating device type groups.

[Figure 4-28](#) illustrates the task of adding a network device to a group under **Network Resources > Network Devices and AAA Clients**. The network device definition can be associated with a specific IP address or a subnet mask, where all IP addresses within the subnet can access the network. The device definition includes the association of the device to NDGs. You also configure here whether the device uses TACACS+ or RADIUS.

The screenshot shows the Cisco Secure ACS web interface. The left sidebar contains a navigation menu with items like My Workspace, Network Resources, Network Device Groups, Location, Device Type, Network Devices and AAA Clients (highlighted), Default Network Device, External RADIUS Servers, Users and Identity Stores, Policy Elements, Access Policies, Monitoring and Reports, and System Administration. The main content area is titled 'Network Resources > Network Devices and AAA Clients > Create'. The form includes the following fields and options:

- Name:** IOS-FW
- Description:** Corporate 2811 Router
- Network Device Groups:**
 - Location:** All Locations (with a 'Select' button)
 - Device Type:** All Device Types:Routers (with a 'Select' button)
- IP Address:** Single IP Address selected, IP: 10.10.0.1
- Authentication Options:**
 - TACACS+:** Checked. Shared Secret: TACACS-key. Single Connect Device checked. Legacy TACACS+ Single Connect Support selected. TACACS+ Draft Compliant Single Connect Support unselected.
 - RADIUS:** Unchecked. Shared Secret: (empty). CoA port: 1700. Enable KeyWrap unselected. Key Encryption Key: (empty). Message Authenticator Code Key: (empty). Key Input Format: ASCII unselected, HEXADECIMAL selected.

At the bottom of the form are 'Submit' and 'Cancel' buttons.

Figure 4-28. Adding Network Devices to Cisco Secure ACS

Defining Users, Identity Groups, and Internal Identity Stores

The second step is to define the user and identity stores. In our scenario, the internal Cisco Secure ACS store is used to validate users and credentials. The next step, then, is to define users and groups in the local Cisco Secure ACS database. This can be accomplished by navigating to **Users and Identity Stores > Internal Identity Stores > Users** and clicking **Create** to add a new user, as illustrated in [Figure 4-29](#). This figure shows the creation of a new user in the internal identity store. The user is assigned to an existing group.

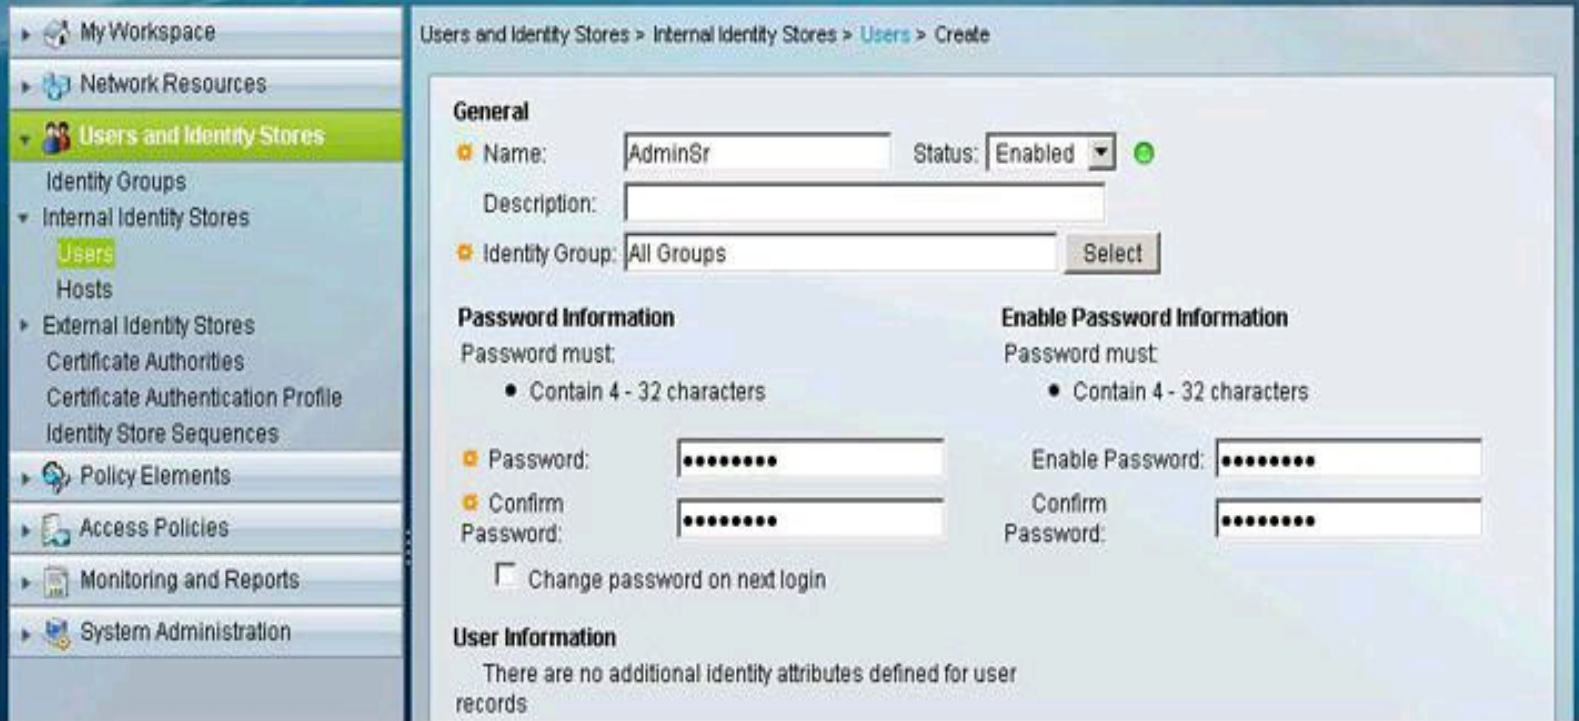


Figure 4-29. Creating Users in Identity Store

Configuring Access Services

The third step is to configure the access services that will define the CLI access policy. In Cisco Secure ACS 5.x, a policy is a set of rules that the server uses to evaluate an access request and return a decision. Cisco Secure ACS 5.x organizes the sequence of independent policies (a policy workflow) into an access service, which it uses to process an access request. You can create multiple access services to process different kinds of access requests, such as for device administration or network access.

When you create an access service, you define the type of policies and policy structures that it contains; for example, policies for device administration or network access, as shown in [Figure 4-30](#). The default access services use and match RADIUS for network access policies, while using and matching TACACS+ for device administration policies. This means that, in this example, there is no need to change the default settings; we are using TACACS+ as the AAA protocol between the router and the Cisco Secure ACS server.

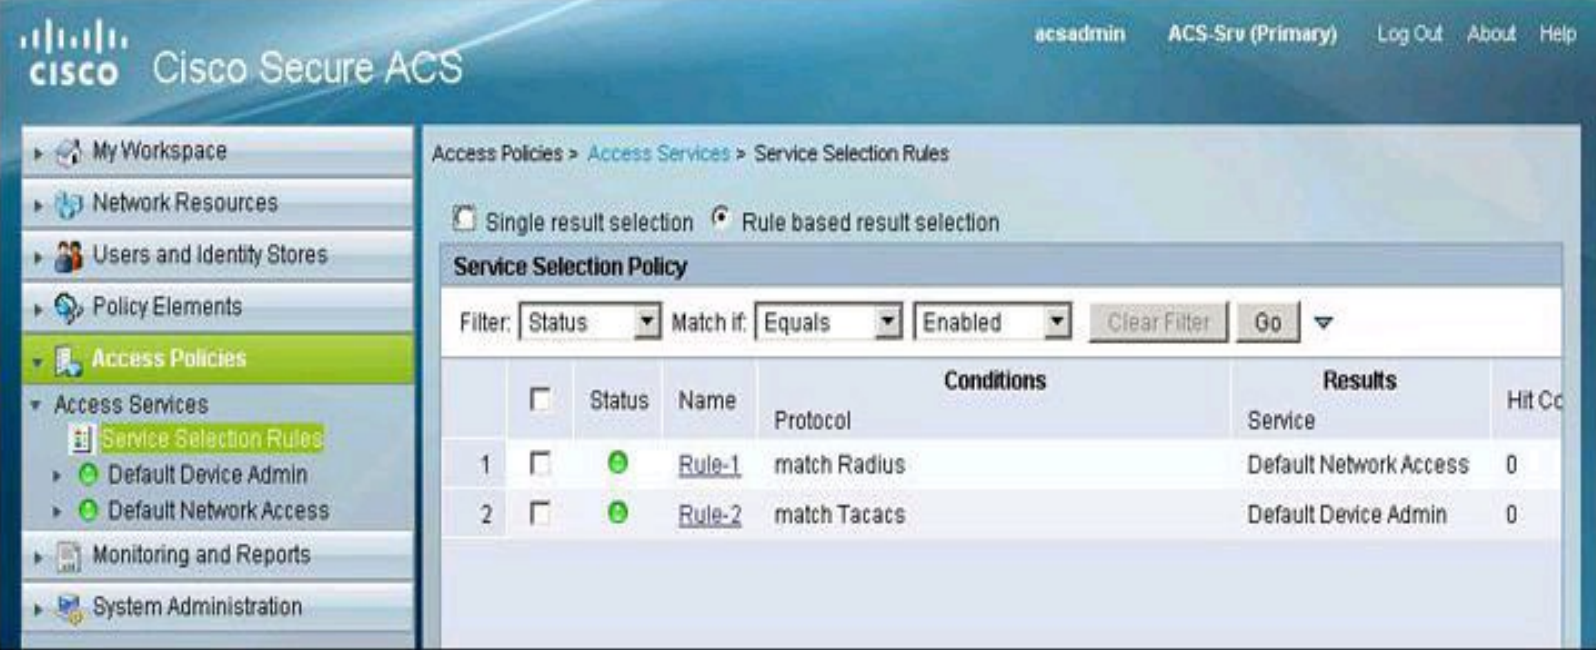


Figure 4-30. Configuring Access Services

Configuring the Identity Policy

The fourth step is to configure the identity policy, which defines how requests are authenticated. The navigation path for this option is **Access Policies > Access Services > Default Device Admin > Identity**. In our scenario, the access service uses the local Cisco Secure ACS store for router device administration. This means that the Default Device Admin Access Service is to be configured with the Internal Users store as an identity source, as shown in [Figure 4-31](#). This is also the default setting. Other options include external databases, such as LDAP and Active Directory.

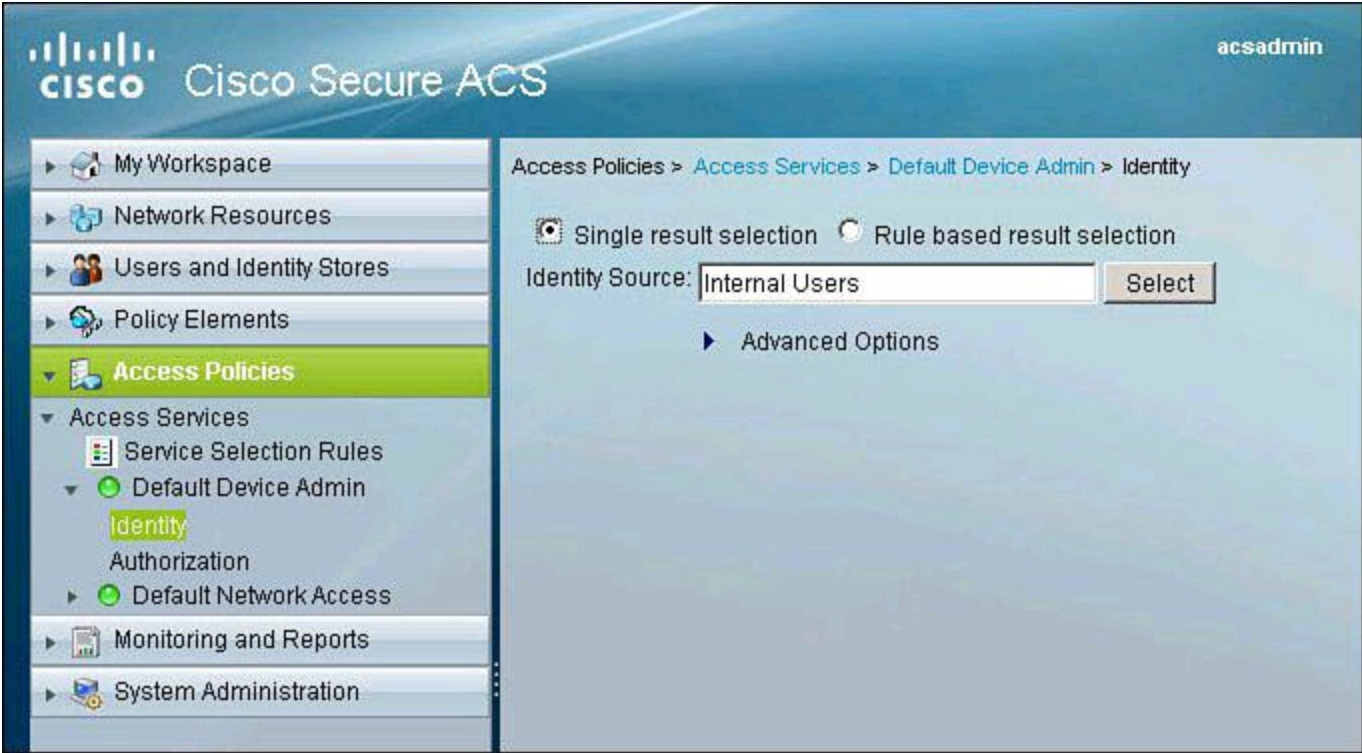


Figure 4-31. Configuring Identity Policy

Configuring the Authorization Policy

The fifth step configures the authorization policy, which defines the actual rules for Role-Based CLI Access and device administration, in our scenario shown in [Figure 4-32](#). The navigation path is **Access Policies > Access Services > Default Device Admin > Authorization**. You then click **Create** to add a new authorization policy.



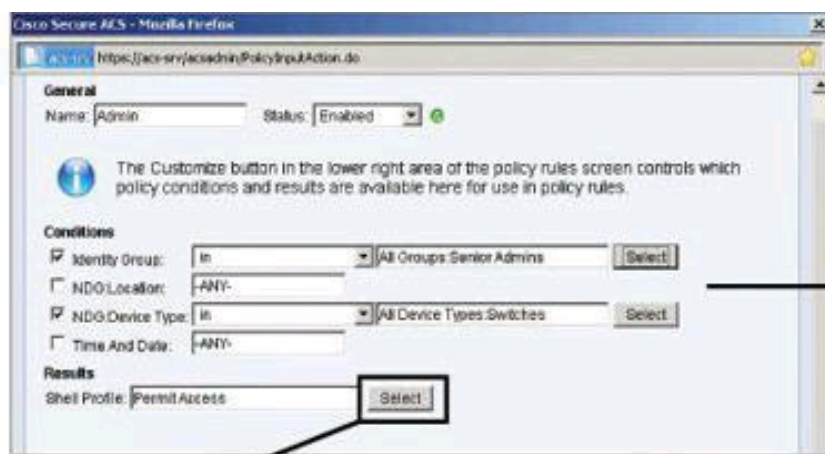
Figure 4-32. Configuring Authorization Policy

Notice the familiar framework of policies that are based on a set of conditions and a profile to apply if those conditions match an access request.

When you create an access service and select a service policy structure for Device Administration, ACS automatically creates a shell/command authorization policy. You can then create and modify policy rules.

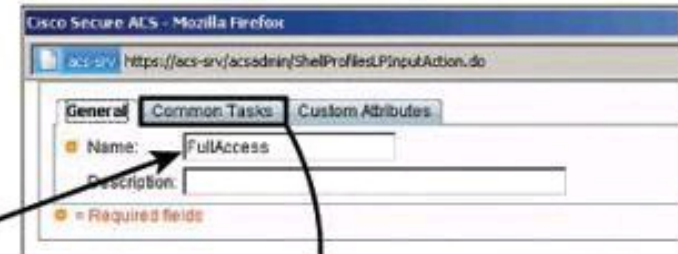
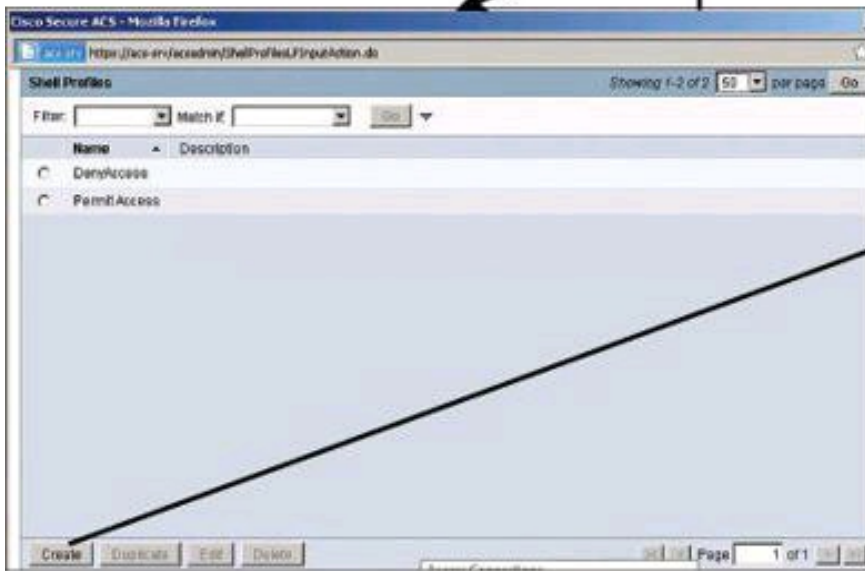
The web interface supports the creation of multiple command sets for device administration. With this capability, you can maintain a smaller number of basic command sets. You can then choose the command sets in combination as rule results, rather than maintaining all the combinations themselves in individual command sets.

[Figure 4-33](#) illustrates the creation of an authorization rule for the admin user. Notice the Conditions section, where you can define the matching conditions that would trigger the policy (in this example, membership to the Admin group and an NDG of Routers). Also notice the Results section, with its corresponding Select button. This option allows you to create customized tasks or attributes to be granted to user admin. In this example, we configure a privilege level 15 for this user.



If <conditions>...

...then apply this profile



using this privilege level.

Figure 4-33. Configuring an Authorization Rule for Admin User

Figure 4-34 shows the final configuration of authorization policies for both users: Admin, obtaining the shell profile FullAccess, which grants privilege level 15, and Monitor, obtaining the shell profile ReadOnly, which grants privilege level 1.



Figure 4-34. Final Configuration of Authorization Policies for Users

Summary

The following topics were discussed in this chapter:

- Role-based access control, use of strong authentication and cryptography, and deployment of NTP are among the recommended practices in secure management and reporting.
- Role-based access control can be accomplished on the router CLI using views.
- Syslog and SNMP provide management notification options for Cisco routers.
- AAA services can use the local router database or a remote AAA server database.
- Method lists provide redundancy and fallback options for AAA configuration using Cisco Configuration Professional.
- CLI commands can be used to verify the correct configuration of the router.
- Cisco Secure ACS can be used as an external AAA database for authentication, authorization, and accounting.
- Two popular AAA protocols, TACACS+ and RADIUS, can be used by network devices to communicate with an external AAA server.
- AAA deployment can be completed by configuring network elements to use AAA authentication and authorization.

References

For additional information, refer to these resources:

Cisco.com Resources

“Cisco Configuration Professional,” <http://www.cisco.com/go/ccp>

“Cisco

ISE

Fundamentals,”

http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5712/ps11637/ps11195/ise_funda

“Cisco Secure Access Control System,” <http://www.cisco.com/go/acs>

“Password Recovery Procedures,”

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_tech_note09186a008

User Guide for the Cisco Secure Access Control System 5.2,

http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_system/5.2/user

RFCs

RFC 1305, “Network Time Protocol (Version 3) Specification, Implementation and Analysis,” <http://www.faqs.org/rfc/rfc1305.html>

RFC 2571, “An Architecture for Describing SNMP Management Frameworks,” <http://www.ietf.org/rfc/rfc2571.txt>

Secure Shell

Wikipedia. “Secure Shell,” http://en.wikipedia.org/wiki/Secure_Shell

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions.](#)”

1. Which port is used by NTP?

- a. UDP 49
- b. TCP 69
- c. UDP 123
- d. TCP 1211

2. Which two statements are true about passwords?

- a. The console port does not require a password for initial console administrative access.
- b. Passwords cannot include symbols and spaces.
- c. You can use either the enable password or the enable secret password to enter the enable mode when both are configured on a router.
- d. The enable secret password is always hashed inside the router configuration using a MD5 hashing algorithm.

3. What is the purpose of the 5 in the following command: **username Admin secret 5 \$1\$feb0\$a104Qd9UZ./Ak00KTggPD0?**

- a. The encrypted password will be converted to cleartext in the running configuration using MD5.
- b. The user Admin has privilege level 5 in this router.
- c. The user Admin has five attempts to log in to this router.
- d. The encrypted-secret password was hashed with MD5.

4. What is the purpose of the command **parser view?**

- a. It creates different views of router configuration for different users.
- b. It provides a root view of all commands.
- c. It configures a password for a view.
- d. It is used to enter commands accessible for that view.

5. Which command would you use to enable Cisco IOS resilience?

- a. **secure boot-ios**
- b. **secure boot-config**
- c. **secure boot-image**
- d. **secure boot-flash**

6. Which command is used to enable AAA on a router?

- a. **aaa enable**
- b. **aaa authentication new-model**
- c. **aaa tacacs+**
- d. **aaa new-model**

7. Which of the following is not supported by Cisco Secure ACS?

- a. StreetTalk Directory Services
- b. Directory Services
- c. One-time passwords
- d. ODBC databases

8. Which portion of a RADIUS packet is sent in plaintext?

- a. All portions
- b. All portions except for the password
- c. All portions except for the CHAP password
- d. All portions except for the username

9. Which of the following protocols is automatically enabled once RSA keys have been generated?

- a. Telnet
- b. SDM
- c. SSH
- d. HTTPS

10. What is significantly different with Cisco Secure ACS 5.x compared to an older version of Cisco Secure ACS?

- a. Users are authenticated locally but authorized on a remote AAA server.
- b. Permissions are granted on conditions other than identity alone.
- c. The desktop is first checked against a set of policies prior to presenting the user with a login prompt.

d. Different users can share one single AAA policies.

Chapter 5. Securing the Data Plane on Cisco Catalyst Switches

In this chapter, you learn that, like routers, both Layer 2 and Layer 3 switches have their own set of network security requirements. 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 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 much damage. In this chapter, you will examine various Layer 2 attacks and strategies to mitigate them. Topics covered in this chapter include the following:

- An introduction to fundamental switching concepts, starting with the building blocks of VLANs and trunking
- An introduction to other building blocks of switching technology, including Spanning Tree Protocol for high availability
- A revisit and further explanation of security threats that exploit vulnerabilities in the switching infrastructure
- A description of how to plan and develop a strategy for protecting the data plane
- A description of the Spanning Tree Protocol Toolkit found on Cisco IOS routers that prevents STP operations from having an impact on the security posture
- A review of port security and how to configure it, to illustrate security controls that are aimed at mitigating MAC spoofing and other threats

Note

Prior to covering Layer 2 data plane security, this chapter includes an overview of related technologies such as VLANs, trunking, and spanning tree. For greater details on these topics, refer to the third edition of *Cisco Press CCNA ICND2 640-816 Official Cert Guide*.

Overview of VLANs and Trunking

A virtual LAN (VLAN) is a logical broadcast domain that can span multiple physical LAN segments. Within the switched internetwork, VLANs provide segmentation and organizational flexibility. You can design a VLAN structure that lets you group stations that are segmented logically by functions, project teams, and applications without regard to the physical location of the users, as shown in [Figure 5-1](#). You can assign each port of a switch to only one VLAN, adding a layer of security. Ports in a single VLAN share broadcasts, while ports in different VLANs do not share broadcasts. Containing broadcasts within a VLAN improves the overall performance of the network.




Figure 5-1. A Building LAN Segmented with VLANs

Within the switched internetwork, VLANs provide segmentation and organizational flexibility. Using VLAN technology, you can group switch ports and their connected users into logically defined communities, such as coworkers in the same department, a cross-functional product team, or diverse user groups sharing the same network application.

A VLAN can exist on a single switch or span multiple switches. VLANs can include stations in a single building or in a multiple-building infrastructure. VLANs can also connect across WANs.

A Cisco Catalyst switch operates in a network similar to a traditional bridge. Each VLAN that you configure on the switch implements address learning, forwarding and filtering decisions, and loop avoidance mechanisms as if the VLAN were a separate physical bridge.

A Cisco Catalyst switch implements VLANs by restricting traffic forwarding to destination ports that are in the same VLAN as the originating ports. So when a frame arrives on a switch port, the switch must retransmit the frame to only the ports that belong to the same VLAN. In essence, a VLAN that is operating on a switch limits transmission of unicast, multicast, and broadcast traffic. Traffic originating from a particular VLAN floods to only the other ports in that VLAN, creating a broadcast domain.

What is a VLAN? It is a single broadcast domain. It is a logical network, thus it is a subnet.

**Key
Topic**

Trunking and 802.1Q

By default, a switch port carries the traffic for only a single VLAN. This is called an access port. If two interconnected switches need to exchange frames from more than one VLAN, their interconnecting ports are configured as trunks, instead of access ports. Therefore, a trunk port can

carry traffic for multiple VLANs, as shown at the top of [Figure 5-2](#).

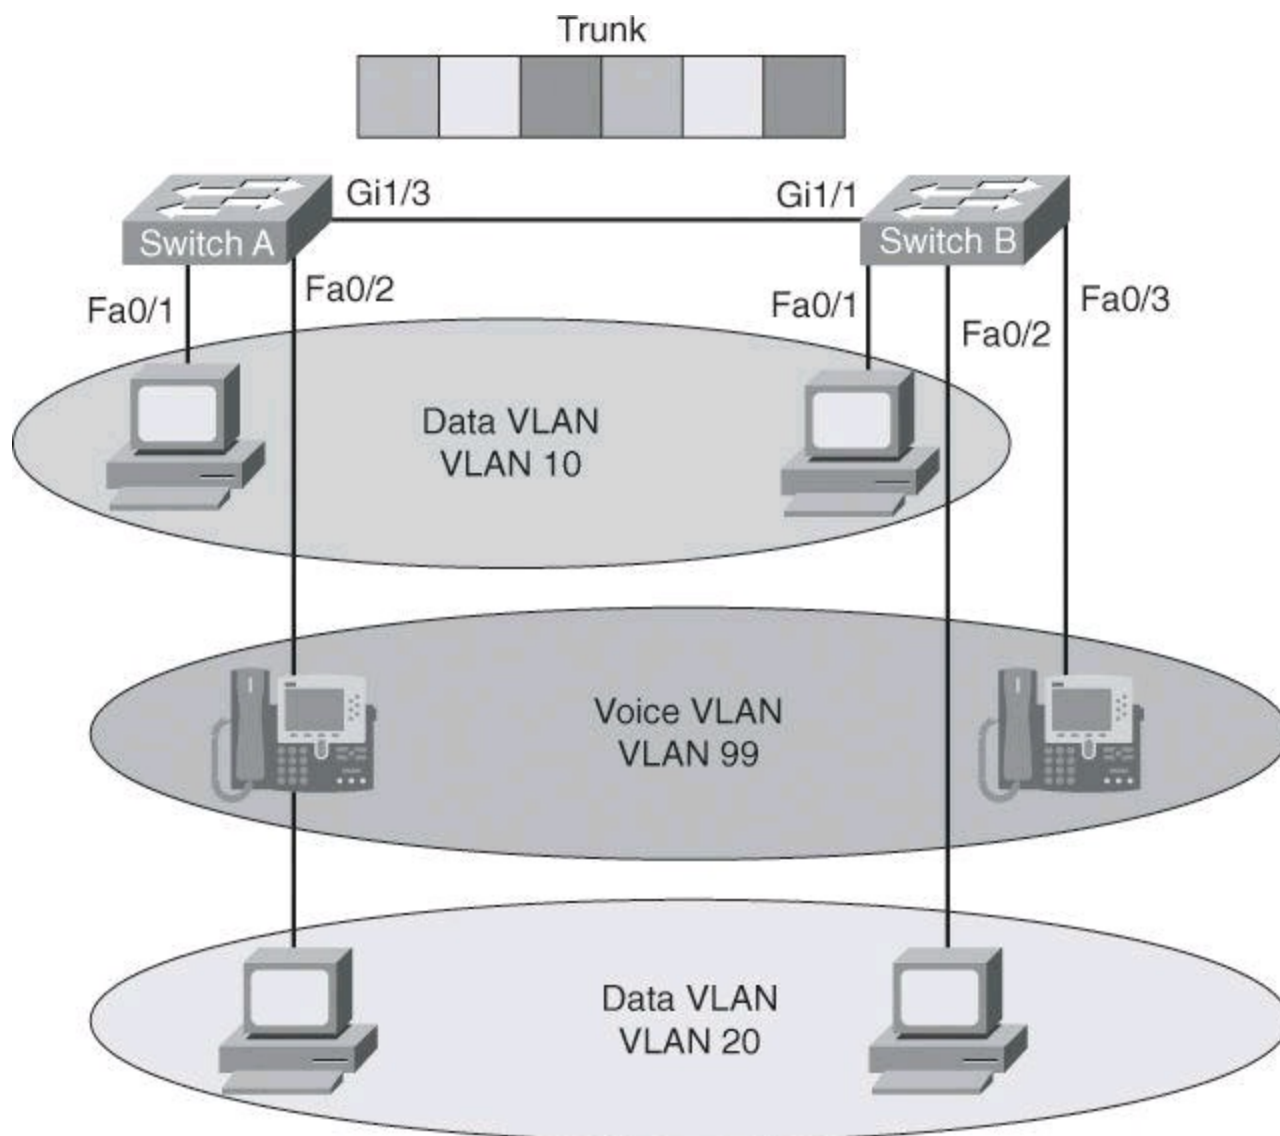


Figure 5-2. Trunk Ports Carry Traffic for Multiple VLANs

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. A trunk is meant to carry the traffic of multiple VLANs over that single link and allow you to extend the VLANs across an entire network.

Trunk Versus Multi-VLAN Access Port

As just mentioned, a port that carries traffic for a single VLAN is called an access port, and a port that carries traffic for multiple VLANs is called a trunk port. However, a clarification is needed: a trunk is a port that carries traffic for multiple *data VLANs*. If a port carries traffic for both a data VLAN and a voice VLAN, it is not called a trunk: it is called a *multi-VLAN access port*. Voice VLANs are not considered data VLANs per se. Therefore, a port carrying voice VLAN traffic and data VLAN traffic is not technically carrying traffic for two data VLANs, and thus is not a trunk.

Looking at [Figure 5-2](#), notice that connected in port Fa0/2 of Switch A is an IP Phone, from which is hanging a PC. This IP Phone has a built-in *mini* switch. The IP Phone's built-in switch, connected to port Fa0/2, passes to Switch A traffic from the PC (VLAN 20 traffic) and traffic generating by the IP Phone itself when the user places a call (VLAN 99 traffic).

Port Fa0/2 on Switch A is said to be a multi-VLAN access port, and not a trunk. By comparison, port Gi1/2 on the same switch is a trunk because it is carrying traffic for more than one data VLAN: it carries traffic for data VLAN 10 and data VLAN 20, in addition to traffic for voice VLAN 99.

A trunk can be made of one or multiple interfaces combined. Interfaces belonging to a trunk require special encapsulation, which indicates the origin VLAN number of the traffic sent across the trunk. This encapsulation process will be discussed later in this section.

Cisco supports IEEE 802.1Q encapsulation. Ethernet interfaces support different trunking modes: you can configure an interface as trunking or nontrunking, or have it negotiate trunking with the neighboring interface; you can also configure which mode it will operate in and if it will be a trunk or an access port.

Besides the widely popular 802.1Q trunk encapsulation standard, Cisco has its own trunking protocol called Inter-Switch Link (ISL). However, it's not supported by all Cisco switch models.

Interfaces configured as 802.1Q ports are assigned to a trunk. All ports on a trunk are in a native VLAN, which will be covered in more detail later in this section. Every 802.1Q port is assigned an identifier value that is based on the native VLAN ID (VID) of the port (the default is VLAN 1). All untagged frames are assigned to the VLAN specified in this VID parameter.

802.1Q Tagging

IEEE 802.1Q uses an internal tagging mechanism that inserts a 4-byte Tag field into the original Ethernet frame between the source address and type or length fields, as shown in [Figure 5-3](#). Because 802.1Q alters the frame, the trunking device recomputes the frame check sequence (FCS) on the modified frame.

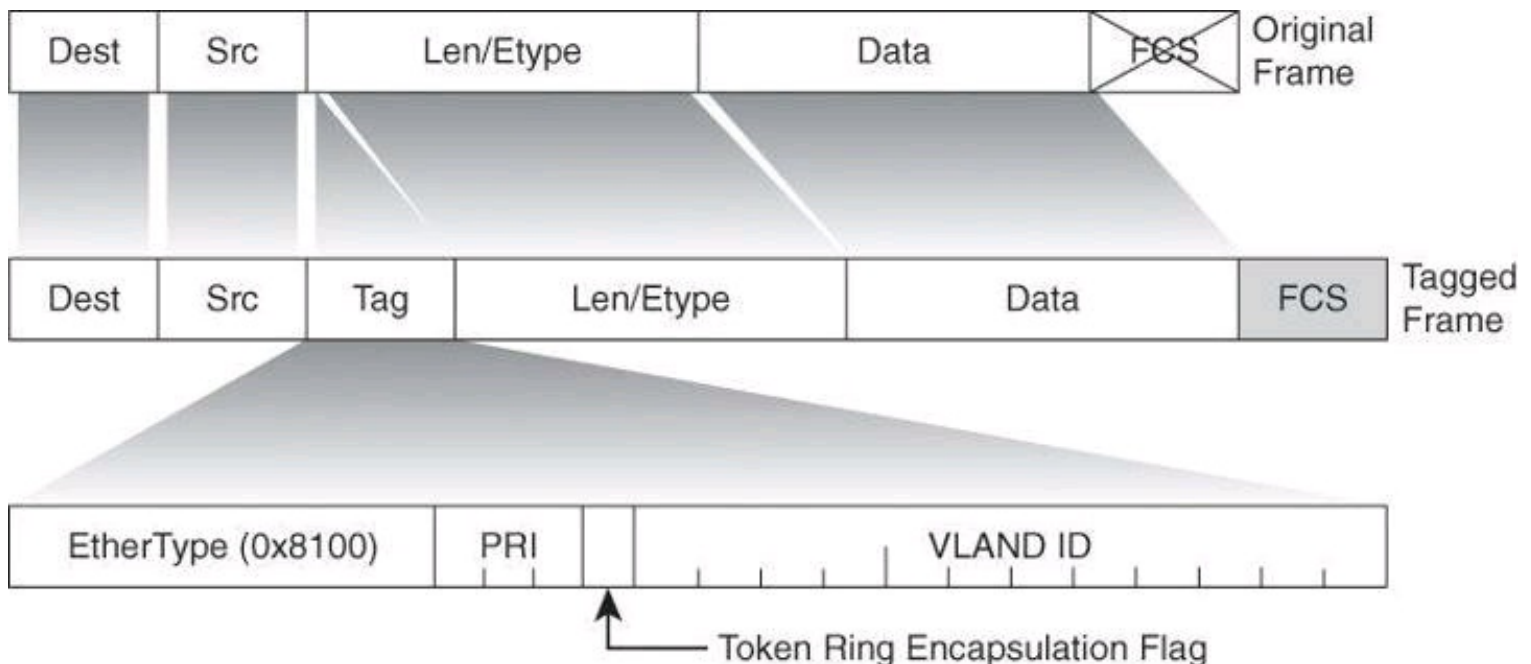


Figure 5-3. 802.1Q Frame Format

It is the responsibility of the Ethernet switch to look at the 4-byte Tag field and determine where to

deliver the frame. A tiny part of the 4-byte Tag field (3 bits, to be exact) is used to specify the priority of the frame. The details of this are specified in the IEEE 802.1p standard. The 802.1Q header contains the 802.1p field, so you must have 802.1Q to have 802.1p.

Native VLANs

An 802.1Q trunk and its associated trunk ports have a native VLAN value. 802.1Q does not tag frames for the native VLAN. Therefore, ordinary stations can read the native untagged frames but cannot read any other frame because the frames are tagged.

If a switch receives an untagged frame on a trunk port, the switch will associate the frame to the native VLAN. Similarly, outgoing frames that belong to the native VLAN will not be tagged using 802.1Q.

Configuring VLANs and Trunks

There are four general tasks required to configure VLANs and trunking and provide inter-VLAN routing. You will notice in the following steps that trunk ports are configured in a single step, while access ports require the configuration of a VLAN first, and then the assignment of the VLAN to the port.

Step 1. Configure and verify 802.1Q trunks.

Step 2. Create or modify a VLAN.

Step 3. Assign switch ports to a VLAN and verify.

Step 4. Configure inter-VLAN routing

Step 1: Configuring and Verifying 802.1Q Trunks

The interface configuration command to set a Fast Ethernet or Gigabit Ethernet port to trunk mode is

```
Switch(config-if)# switchport mode {access | dynamic {auto | desirable}  
| trunk}
```

Many Cisco Catalyst switches support the Dynamic Trunking Protocol (DTP), which manages automatic trunk negotiation.

[Table 5-1](#) describes the four options for the **switchport mode** command.

Table 5-1. *switchport mode* Command Parameters

	Description
trunk	Configures the port into permanent 802.1Q trunk mode and negotiates with the connected device to convert the link to trunk mode.
access	Disables port trunk mode and negotiates with the connected device to convert the link to nontrunk.
dynamic desirable	Triggers the port to negotiate the link from nontrunk to trunk mode. The port negotiates to a trunk port if the connected device is in either trunk state, desirable state, or auto state. Otherwise, the port becomes a nontrunk port.
dynamic auto	Enables a port to become a trunk only if the connected device has the state set to trunk or desirable. Otherwise, the port becomes a nontrunk port.

[Example 5-1](#) shows how to configure a port as a trunk.

Example 5-1. *switchport mode* Command

[Click here to view code image](#)

```
Switch(config)# interface fa0/1
Switch(config-if)# switchport mode trunk
```

You have the following options to control trunking for ports:

- For links that you do not intend to trunk across, use the **switchport mode access** interface configuration command to disable trunking.
- For links that you do intend to trunk across, take the following actions:
 - Use the **switchport mode trunk** interface configuration command to cause the interface to become a trunk link.
 - Use the **switchport nonegotiate** interface configuration command to prevent the generation of Dynamic Trunking Protocol frames. This command is valid only when the interface switchport mode is access or trunk (configured by using the **switchport mode access** or the **switchport mode trunk** interface configuration command). This command returns an error if you attempt to execute it in dynamic (auto or desirable) mode. Use the **no** form of this command to return to the default setting. When you configure a port with the **switchport nonegotiate** command, the port trunks only if the other end of the link is specifically set to trunk. The **switchport nonegotiate** command does not form a trunk link with ports in either dynamic desirable or dynamic auto mode.
 - Use the **switchport trunk native vlan *vlan_number*** interface configuration command to set the native VLAN on the trunk to an unused VLAN. The default native VLAN is VLAN 1.
 - Use the **switchport trunk allowed vlan *vlan_number*** interface configuration command to set the list of allowed VLANs that transmit traffic from this interface

in tagged format when trunking mode is on.

Verifying a Trunk

To verify a trunk configuration on many Cisco Catalyst switches, use the **show interfaces interface switchport** command or the **show interfaces interface trunk** command to display the trunk parameters and VLAN information of the port:

```
Switch# show interfaces interface [switchport | trunk]
```

In the output of the **show interfaces interface switchport** command in [Example 5-2](#), notice the distinction between the operational mode and administrative mode for each port. This is extremely helpful in troubleshooting trunks. The administrative mode is the configured mode, while the operational mode will depend on the results of trunk negotiations. If trunk negotiations fail, the port will be administratively configured as a trunk (as shown in [Example 5-2](#)), but the operational mode will be different. It could be down, as in [Example 5-2](#), or even fall back to an active port on the default VLAN 1.

Example 5-2. *show interfaces interface switchport* Command

[Click here to view code image](#)

```
SwitchX# show interfaces fa0/11 switchport
Name: Fa0/11
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
```

In the output of the **show interfaces interface trunk** command in [Example 5-3](#), you will notice in the allowed VLANs section that traffic for a particular VLAN will not traverse the trunk if the VLAN does not show in this section.

Example 5-3. *show interfaces fa0/11 trunk* Command

[Click here to view code image](#)

```
SwitchX# show interfaces fa0/11 trunk
Port      Mode           Encapsulation  Status      Native vlan
Fa0/11    desirable      802.1Q         trunking    1
Port      Vlans allowed on trunk
Fa0/11    1-4094
Port      Vlans allowed and active in management domain
Fa0/11    1-13Trunking Native Mode VLAN: 1 (default)
```

Step 2: Creating a VLAN

The maximum number of VLANs that can be created is switch-dependent. However, most Cisco Catalyst desktop switches support 128 separate spanning-tree instances, one per VLAN. VLAN 1 is the factory default Ethernet VLAN. The Cisco Discovery Protocol (CDP) advertisements and other protocols are sent on VLAN 1. The Cisco Catalyst switch IP address is in the management VLAN (VLAN 1 by default) that is the IP address used to telnet, or SSH, in the switch. This management IP address is also used by the switch if it has been configured for Simple Network Management Protocol (SNMP) or syslog.

The configuration commands to create a VLAN are as follows:

```
Switch(config)# vlan vlan-id
Switch(config-vlan)# name vlan-name
```

[Table 5-2](#) lists the commands for creating VLANs.

Table 5-2. Adding a VLAN Command

Parameter	Description
<code>vlan <i>vlan-id</i></code>	ID of the VLAN to be added and configured. For <i>vlan-id</i> , the range is 1 to 4094 when the enhanced software image is installed and 1 to 1005 when the standard software image is installed. Do not enter leading zeros. You can enter a single VID, a series of VIDs separated by commas, or a range of VIDs separated by hyphens.
<code>name <i>vlan-name</i></code>	(Optional) Specify the VLAN name, an ASCII string from 1 to 32 characters that must be unique within the administrative domain.

[Example 5-4](#) shows how to create a VLAN.

Example 5-4. Example of VLAN Creation

[Click here to view code image](#)

```
Switch(config)# vlan 2
Switch(config-vlan)# name Marketing
```

Before you create VLANs, consider these guidelines:

- Most Cisco Catalyst desktop switches support a maximum of 128 spanning-tree instances. If the number of VLANs on the switch exceeds the number of supported spanning-tree instances, it is recommended that you configure Multiple Spanning Tree Protocol (MSTP) on your switch to map multiple VLANs to a single spanning-tree instance.
- The maximum number of VLANs is switch-dependent. Many access layer Cisco Catalyst switches can support up to 250 user-defined VLANs.
- Cisco Catalyst switches have a factory default configuration in which various default VLANs are preconfigured to support various media and protocol types. The default Ethernet VLAN is VLAN 1. Cisco Discovery Protocol advertisements and other protocols

are sent on VLAN 1.

- For you to be able to communicate with the Cisco Catalyst switch remotely for management purposes, the switch must have an IP address. This IP address must be in a management VLAN; the default management VLAN is VLAN 1.

After you configure the VLAN, you should validate the parameters for that VLAN. Use the **show vlan** command to display information about a particular VLAN:

```
Switch# show vlan [brief | id vlan-id | name vlan-name]
```

[Example 5-5](#) shows the output of the **show vlan id** command.

Example 5-5. Output of the *show vlan id* Command

[Click here to view code image](#)

```
SwitchX# show vlan id 2
VLAN Name                               Status      Ports
-----
2      switchlab99                          active      Fa0/2, Fa0/12

VLAN Type  SAID    MTU   Parent RingNo BridgeNo  Stp  BrdgMode  Trans1  Trans2
-----
2      enet   100002  1500  -      -      -      -      -      0      0

SwitchX#
```

Use the **show vlan brief** command to display one line for each VLAN that displays the VLAN name, the status, and the switch ports.

Use the **show vlan** command to display information on all configured VLANs. The **show vlan** command displays the switch ports that are assigned to each VLAN. Other VLAN parameters that are displayed include the type (the default is Ethernet); the security association ID (SAID), used for the FDDI trunk; the maximum transmission unit (MTU) (the default is 1500 for Ethernet VLAN); the STP; and other parameters that are used for Token Ring or FDDI VLANs. (Chances are you have never heard of Token Ring or FDDI network. These were competing technologies prior to the obvious emergence of Ethernet as leader.)

Step 3: Assigning Switch Ports to a VLAN

After creating a VLAN, you can manually assign a port or a number of ports to that VLAN with the following command:

```
Switch(config-if)# switchport access [vlan vlan-id | dynamic]
```

A port can belong to only one VLAN at a time. When you assign a switch port to a VLAN using this method, as shown in [Example 5-6](#), it is known as a static-access port.

Example 5-6. Assigning a VLAN to a Port with the *switchport access* Command

[Click here to view code image](#)

```
SwitchX(config)# interface range fastethernet 0/2 - 4
SwitchX(config-if)# switchport access vlan 2

SwitchX# show vlan
```

VLAN	Name	Status	Ports
1	default	active	Fa0/1
2	switchlab99	active	Fa0/2, Fa0/3, Fa0/4

On most Cisco Catalyst switches, you configure the VLAN port assignment from interface configuration mode using the **switchport access** command. Use the **vlan *vlan_number*** option to set static-access membership. The **dynamic** option uses the VLAN Membership Policy Server (VMPS) method and is outside of the scope of this book.

Use the **show vlan brief** privileged EXEC command to display the VLAN assignment and membership type for all switch ports, as shown in [Example 5-7](#). It is worth mentioning that if a given port does not show under the Ports column of the table, it could be a trunk port carrying multiple VLANs.

Example 5-7. Verifying a VLAN Assignment with the *show vlan* Command

[Click here to view code image](#)

```
SwitchX# show vlan brief
VLAN Name                Status      Ports
-----
1    default                active     Fa0/1
2    switchlab99            active     Fa0/2, Fa0/3, Fa0/4
3    vlan3                  active
4    vlan4                  active
1002 fddi-default           act/unsup
1003 token-ring-default   act/unsup

VLAN Name                Status      Ports
-----
1004 fddinet-default       act/unsup
1005 trnet-default        act/unsup/4
```

Step 4: Configuring Inter-VLAN Routing

Inter-VLAN communication occurs between broadcast domains via a Layer 3 device. In a VLAN environment, frames are switched only between ports within the same broadcast domain. VLANs perform network partitioning and traffic separation at Layer 2. Inter-VLAN communication cannot occur without a Layer 3 device, such as a router.

One way to accomplish inter-VLAN routing is illustrated in [Figure 5-4](#). You can use IEEE 802.1Q to enable trunking on a router interface, and enable the router to route between VLANs. This is

sometimes known as a “router on a stick” scenario, where a router is attached to a core switch. The router can receive packets on one VLAN and forward them to another VLAN. To perform inter-VLAN routing functions, the router must know how to reach all VLANs being interconnected. There must be a separate connection on the router for each VLAN, and you must enable 802.1Q trunking on those connections. The router already knows about directly connected networks. The router must learn routes to networks to which it is not directly connected.

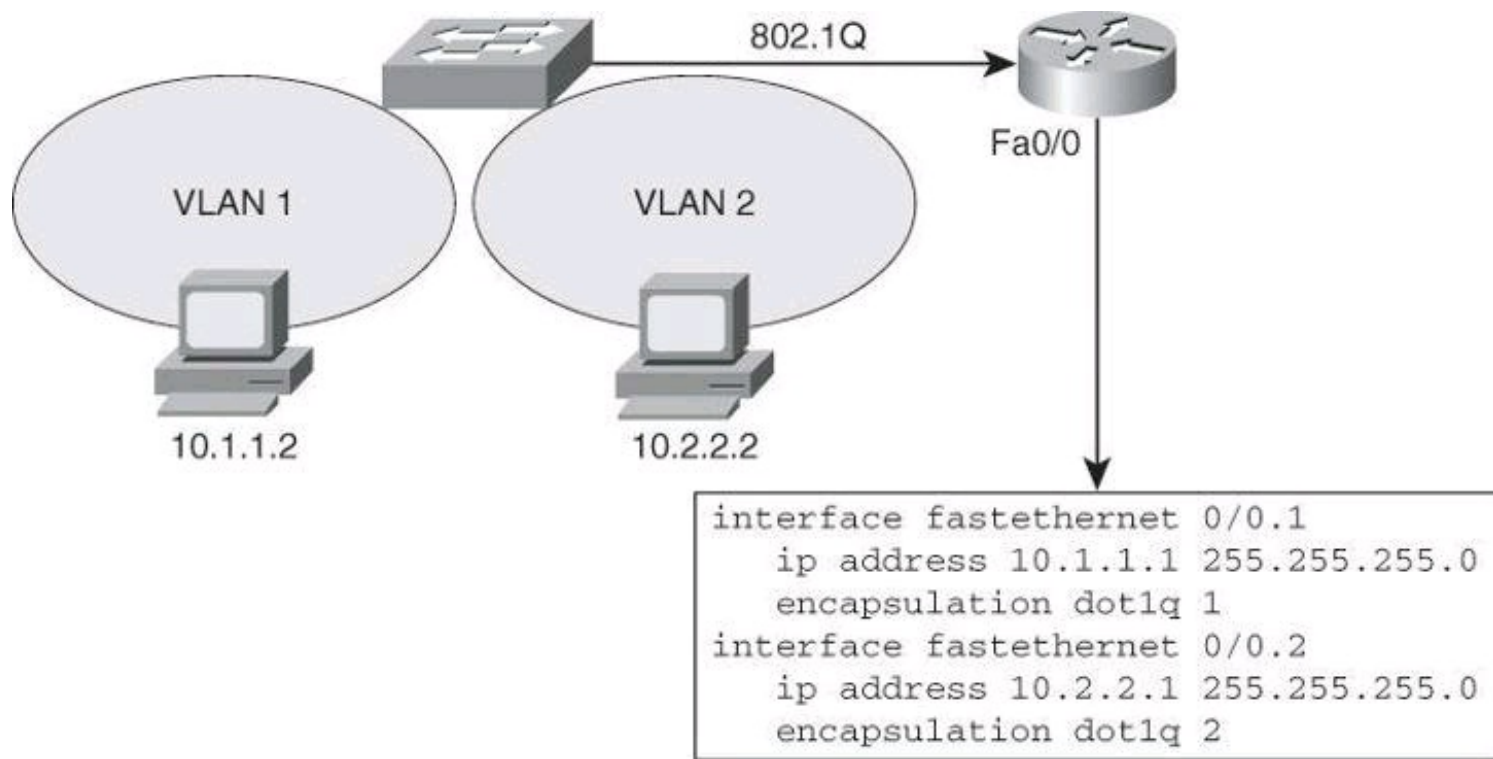


Figure 5-4. Routing Between VLANs with 802.1Q Trunks

To support 802.1Q trunking, you must subdivide the physical Fast Ethernet interface of the router into multiple, logical, addressable interfaces, one per VLAN. The resulting logical interfaces are called subinterfaces. Without this subdivision, you would have to dedicate a separate physical interface to each VLAN.

In [Figure 5-4](#), the FastEthernet 0/0 interface is divided into two subinterfaces: FastEthernet 0/0.1 and FastEthernet0/0.2. Each subinterface represents the router in each of the VLANs for which it routes. The sample configuration in [Figure 5-4](#) uses the **encapsulation dot1q *vlan-id*** command on each subinterface to enable 802.1Q encapsulation trunking. The subinterface number does not have to be the same as the dot1q VLAN number; however, management is easier when both numbers are the same.

Note

Layer 3 switches, multilayer switches, are meant to replace the router on a stick shown in [Figure 5-4](#) by incorporating the Layer 3 routing functionality and VLAN interfaces within the switch itself

It is also worth noting that the resulting configuration of FastEthernet0/0.1 on the router in [Figure 5-4](#), from issuing the command **encapsulation dot1q 1**, would be **encapsulation dot1q 1 native**. The topic of native VLANs was covered earlier in the chapter.

Spanning Tree Overview

Redundant designs, such as the one shown in [Figure 5-5](#), can mitigate the possibility of a *single point of failure*, which can cause a loss of function for the entire switched or bridged network.




Figure 5-5. Redundant Topology

However, you must consider problems that redundant designs can cause. Some of the problems that can occur with redundant links and devices in switched or bridged networks are as follows:

- **Broadcast storms:** Without some loop-avoidance process in operation, each switch or bridge floods broadcasts endlessly. This situation is commonly called a broadcast storm.
- **Multiple frame transmission:** Multiple copies of unicast frames may be delivered to destination stations. Many protocols expect to receive only a single copy of each transmission. Multiple copies of the same frame can cause unrecoverable errors.
- **MAC database instability:** Instability in the content of the MAC address table results from copies of the same frame being received on different ports of the switch. Data forwarding can be impaired when the switch consumes the resources that are coping with instability in the MAC address table.

Layer 2 LAN protocols, such as Ethernet, lack a mechanism to recognize and eliminate endlessly looping frames. Some Layer 3 protocols implement a Time to Live (TTL) mechanism that limits the number of times a Layer 3 networking device can retransmit a packet. Lacking such a mechanism, Layer 2 devices continue to retransmit looping traffic indefinitely.

A loop-avoidance mechanism is required to solve each of these problems. This mechanism is the Spanning Tree Protocol (STP).

STP Fundamentals

Spanning Tree Protocol (STP) provides loop resolution by managing the physical paths to given network segments. STP allows physical path redundancy while preventing the undesirable effects of active loops in the network. STP is an IEEE committee standard defined as 802.1D.

STP behaves as follows:

- STP forces certain ports into a standby state so that they do not listen to, forward, or flood data frames, as seen in [Figure 5-6](#), where one switch has a port in blocking mode. The overall effect is that only one active path exists to the other network segment at any time.




Figure 5-6. Loop Resolution with STP

- If there is a problem with connectivity to any of the segments within the network, STP reestablishes connectivity by automatically activating a previously inactive path, if one exists.

RSTP, the Rapid Spanning Tree Protocol, is a version of STP enhanced for fast convergence. RSTP is defined as the IEEE standard 802.1w.

PVST+ is a Cisco implementation of RSTP that provides enhancements that are aimed at scalability of the protocol, as well as provisions for traffic load sharing across STP-enabled paths.

STP performs three steps to provide a loop-free logical network topology, as shown in [Figure 5-7](#):

Step 1. Selects one root bridge.

STP has a process to elect a root bridge, which will be discussed later. Only one bridge can act as the root bridge in a given network. On the root bridge, all ports are designated ports. Designated ports are normally in the forwarding state. When in the forwarding state, a port can send and receive traffic. In [Figure 5-7](#), Switch X is the root bridge.

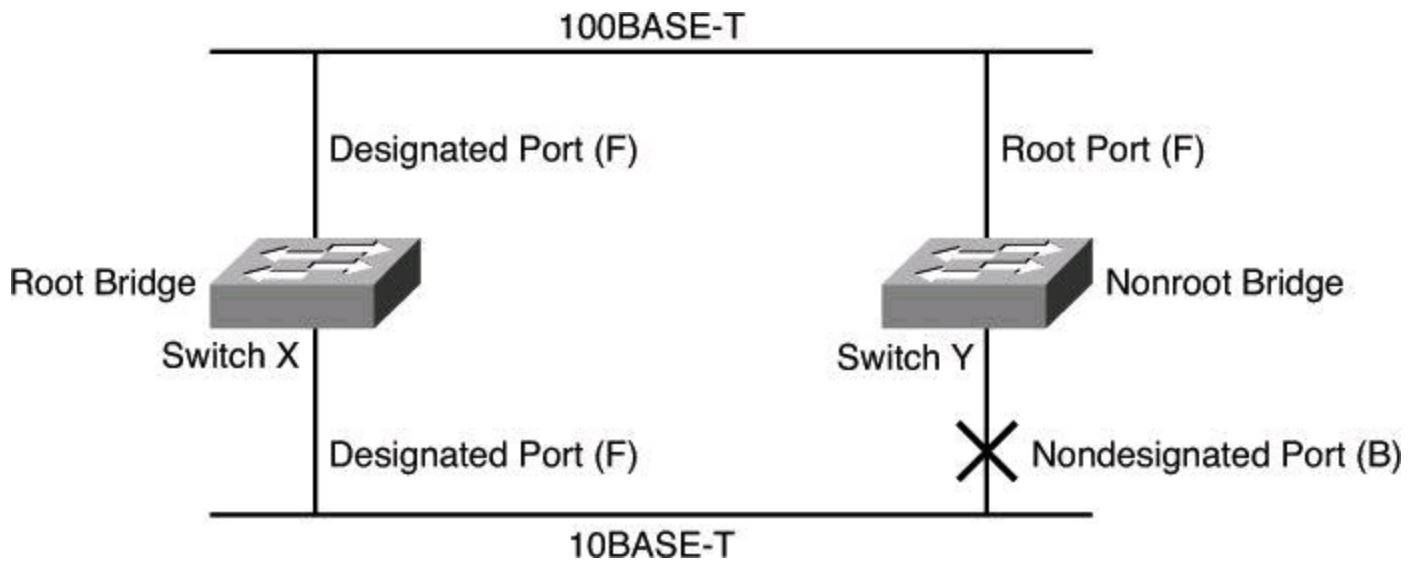


Figure 5-7. STP Operation and Resulting Topology

Step 2. Selects the root port on the nonroot bridge.

STP establishes one root port on each nonroot bridge. The root port is the lowest-cost path from the nonroot bridge to the root bridge. Root ports are normally in the forwarding state. Spanning-tree path cost is an accumulated cost that is calculated on the bandwidth. In [Figure 5-7](#), the lowest-cost path to the root bridge from Switch Y is through the 100BASE-T Fast Ethernet link.

Step 3. Selects the designated port on each segment.

On each segment, STP establishes one designated port. The designated port is selected on the bridge that has the lowest-cost path to the root bridge. Designated ports are normally in the forwarding state, forwarding traffic for the segment. In [Figure 5-7](#), the designated port for both segments is on the root bridge because the root bridge is directly connected to both segments. The 10BASE-T Ethernet port on Switch Y is a nondesignated port because there is only one designated port per segment. Nondesignated ports are normally in the blocking state to logically break the loop topology. When a port is in the blocking state, it is not forwarding traffic but can still receive traffic.

Switches and bridges running the spanning-tree algorithm exchange configuration messages with other switches and bridges at regular intervals (every 2 seconds by default). Switches and bridges exchange these messages using a multicast frame called the bridge protocol data unit (BPDU). One of the pieces of information included in the BPDU is the *bridge ID* (BID).

STP calls for each switch or bridge to be assigned a unique BID. Typically, the BID is composed of a priority value (2 bytes) and the bridge MAC address (6 bytes). The default priority, in accordance with IEEE 802.1D, is 32,768 (1000 0000 0000 0000 in binary, or 0x8000 in hex format), which is the midrange value. The root bridge is the bridge with the lowest BID.

Note

A Cisco Catalyst switch uses one of its MAC addresses from a pool of MAC addresses that

are assigned either to the backplane or to the supervisor module, depending on the switch model.

The bridge ID (BID) is made of the bridge priority plus the bridge MAC address.

Key
Topic

In [Figure 5-8](#), both switches are using the same default priority. The switch with the lowest MAC address is the root bridge. In [Figure 5-8](#), Switch X is the root bridge with the default priority of 0x8000 (hex), or 32,768 in decimal, and a MAC address of 0c00.1111.1111.




Figure 5-8. STP: Root Bridge Selection

Verifying RSTP and PVRST+

PVST is enabled by default in some Cisco Catalyst switch models. You can change the mode to Rapid PVST with the **spanning-tree mode rapid-pvst** command:

```
Switch(config)# spanning-tree mode rapid-pvst
```

This spanning-tree mode is the same as PVST+ except that it uses a rapid convergence that is based on the IEEE 802.1w standard.

For verification, you can use the **show spanning-tree vlan *vlan-range*** command. [Example 5-8](#) shows the result of this command.

Example 5-8. Verifying Spanning-Tree Configuration for vlan 21

[Click here to view code image](#)

```
Switch# show spanning-tree vlan 21
VLAN0021
Spanning tree enabled protocol rstp
Root ID    Priority      32789
           Address      88f0.77c5.0f80
           Cost        19
           Port        1 (FastEthernet0/1)
           Hello Time  2 sec  Max Age 20 sec  Forward Delay
15 sec
```

```

Bridge ID   Priority   32789   (priority 32768   sys-id-ext 21)
Address
Hello Time  d0c2.82c5.6b00
                2 sec   Max Age 20 sec   Forward Delay
15 sec

                Aging Time 300 sec

```

```

Interface          Role Sts Cost          Prio.Nbr Type
-----
Fa0/1              Root FWD 19            128.1   P2p
Fa0/2              Desg FWD 19            128.2   P2p
Fa0/8              Desg FWD 19            128.8   P2p

```

Mitigating Layer 2 Attacks

As stated at the beginning of the chapter, like routers, both Layer 2 and Layer 3 switches have their own set of network security requirements. 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 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.

Basic Switch Operation

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

Layer 2 is the data link layer in the OSI model and is one of seven layers designed to work together but with autonomy. Layer 2 operates above the physical layer, but below the network and transport layers, as shown in [Figure 5-9](#).

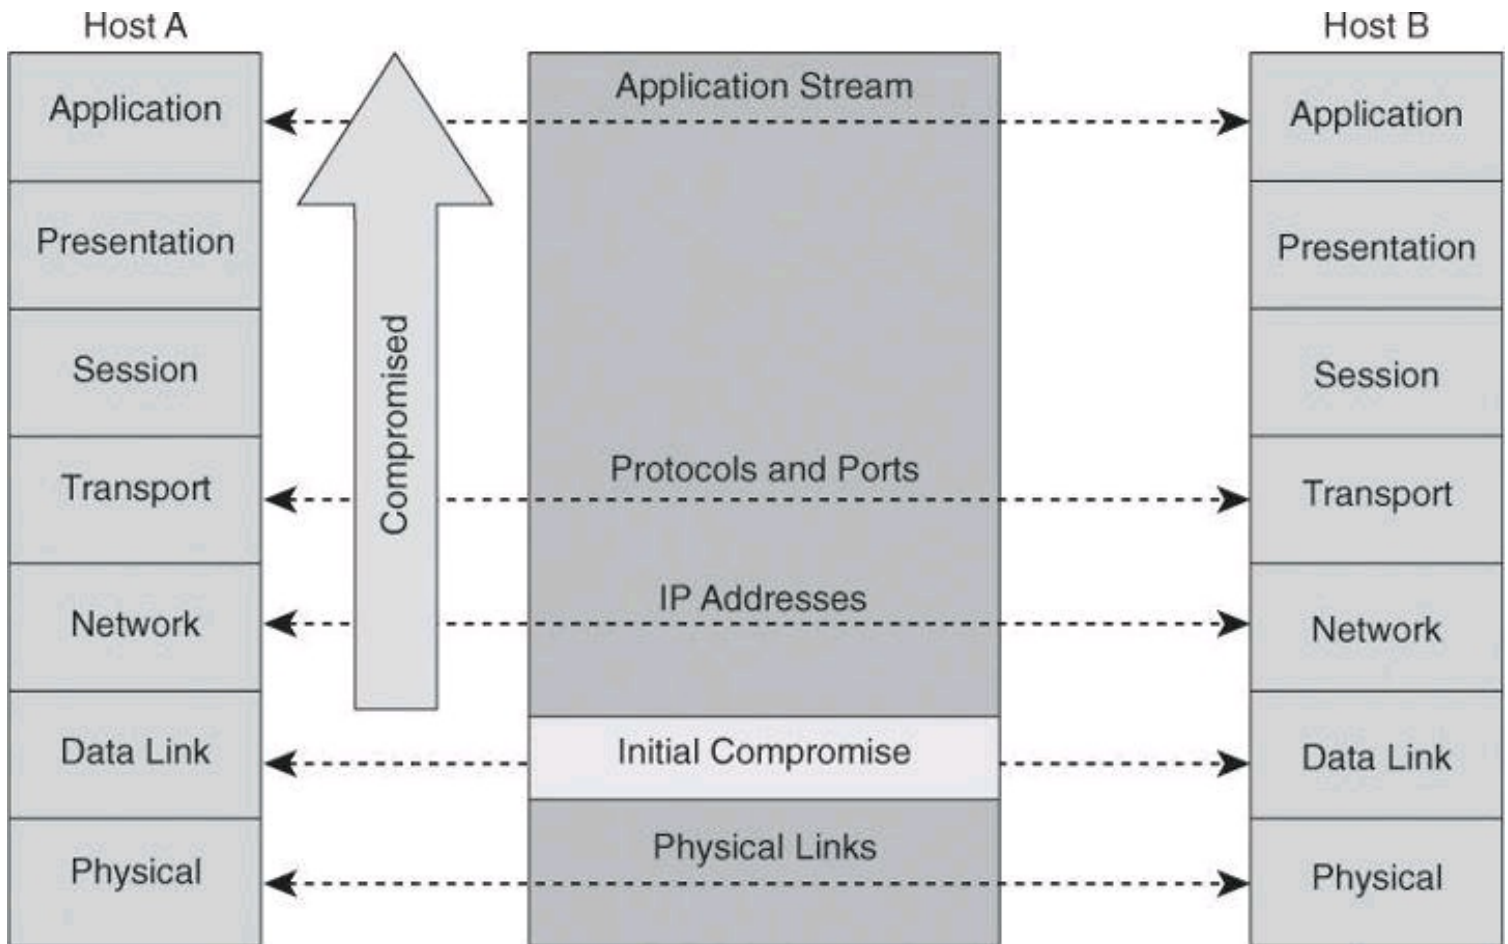


Figure 5-9. Domino Effect If Layer 2 Is Compromised

Layer 2 independence enables interoperability and interconnectivity. However, from a security perspective, Layer 2 independence creates a challenge because a compromise at one layer is not always known by the other layers. If the initial attack comes in at Layer 2, the rest of the network can be compromised in an instant. Network security is only as strong as the weakest link, and that link might be the data link layer.

Layer 2 Best Practices

The following list suggests Layer 2 security best practices. All of these suggestions are dependent upon your security policy.

- Manage switches in as secure a manner as possible (SSH, OOB, permit lists, and so on).
- Whenever practical, declare the VLAN ID used on trunk ports with the **switchport trunk allowed vlan** command
- Do not use VLAN 1 for anything.
- Set all user ports to nontrunking (unless you are using Cisco VoIP).
- Use port security where possible for access ports.
- Selectively use SNMP and treat community strings like root passwords.
- Enable STP attack mitigation (BPDU guard, root guard).
- Use Cisco Discovery Protocol only where necessary (with phones it is useful).
- Disable all unused ports and put them in an unused VLAN.

It is important to manage switches like routers, using secure protocols or out-of-band methods if

policy permits it. Because VLAN 1 is a known management VLAN, it is recommended that you avoid using it. Turn off services that are not necessary and ports that are not being used. Implement the various security services that have been covered in this chapter as necessary and as supported by your hardware. Turn CDP off on ports that do not connect to network devices, with the exception of ports that connect to Cisco IP phones.

Layer 2 Protection Toolkit

Multiple security features and technologies are available to implement recommended practices, in a manner that streamlines configuration and management while strengthening the overall security posture. [Figure 5-10](#) lists some of them, categorized based on type of attack (spoofing and denial of service [DoS]), type of security control (identity services and device hardening), and protocol (STP, DHCP, and Address Resolution Protocol [ARP]).

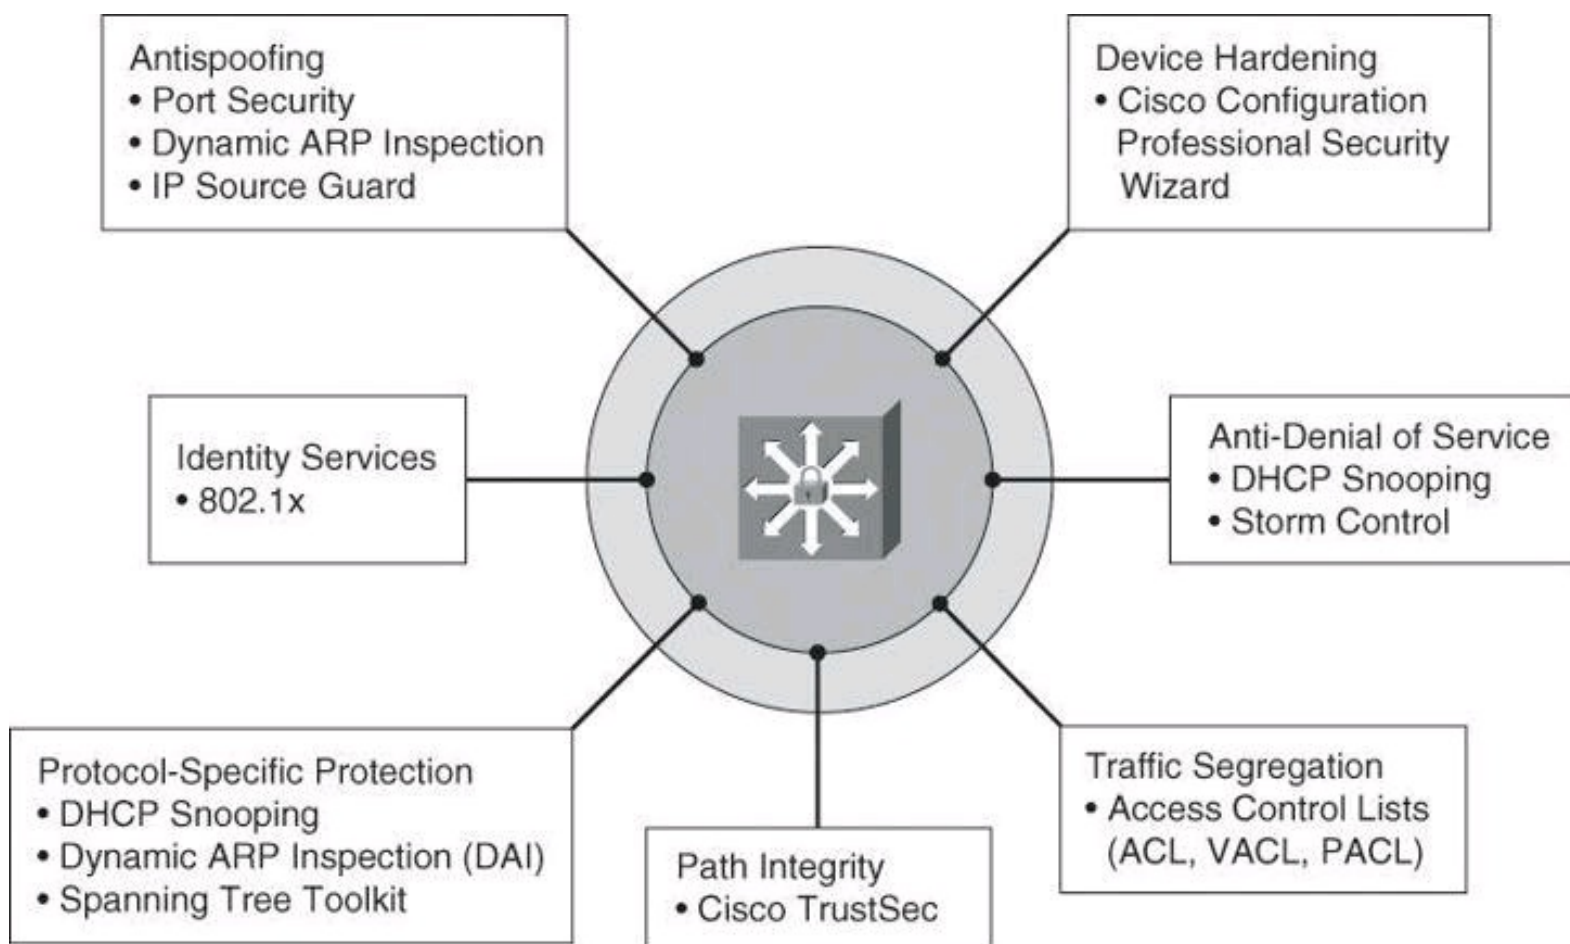


Figure 5-10. Components of Layer 2 Protection Toolkit

Some of these controls will be covered in more detail in this chapter, while some others, such as DHCP snooping and IP Source Guard, are covered in other Cisco Press books, such as *CCNP Security SECURE 642-637 Official Cert Guide*.

Mitigating VLAN Attacks

As mentioned at the beginning of this chapter, a VLAN is a logical broadcast domain that can span multiple physical LAN segments. Ports in a VLAN share broadcasts; ports in different VLANs do not share broadcasts. Containing broadcasts within a VLAN improves the overall performance of the network.

VLAN Hopping

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

VLAN Hopping by Rogue Trunk

In a basic VLAN hopping attack, the attacker takes advantage of the default automatic trunking configuration on most switches. The network attacker configures a system to spoof itself as a switch. This spoofing requires that the network attacker be capable of emulating either ISL or 802.1Q signaling along with Dynamic Trunking Protocol (DTP) signaling, as shown in [Figure 5-11](#). By tricking a switch into thinking it is another switch that needs to trunk, an attacker can gain access to all the VLANs allowed on the trunk port. To succeed, this attack requires a configuration on the port that supports trunking, such as auto. As a result, the attacker is a member of all the VLANs that are trunked on the switch and can “hop” (that is, send and receive traffic) on all of those VLANs.




Figure 5-11. VLAN Hopping by Rogue Trunk

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

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

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

VLAN Hopping by Double Tagging

The double-tagging (or double-encapsulated) VLAN hopping attack takes advantage of the way that

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

A double-tagging VLAN hopping attack follows four steps, as shown in [Figure 5-12](#):

Step 1. The attacker sends a double-tagged 802.1Q frame to the switch. The outer header has the VLAN tag of the attacker, which is the same as the native VLAN of the trunk port. For the purposes of this example, assume that this is VLAN 10. The inner tag is the victim VLAN, VLAN 20.

Step 2. The frame arrives on the switch, which looks at the first 4-byte 802.1Q tag. The switch sees that the frame is destined for VLAN 10 and sends it out all VLAN 10 ports (including the trunk), because there is no CAM table entry. The switch does not add a VLAN 10 tag to the frames because VLAN 10 is the native VLAN, and as specified by the 802.1Q specification, native VLAN traffic is not tagged. At this point, the second VLAN tag is still intact and has not been inspected by the first switch.

Step 3. The frame arrives at the second switch but has no knowledge that it was supposed to be for VLAN 10.

Step 4. The second switch looks only at the 802.1Q tag (the former inner tag that the attacker sent) and sees that the frame is destined for VLAN 20 (the victim VLAN). The second switch sends the packet on to the victim port, or floods it, depending on whether there is an existing CAM table entry for the victim host.

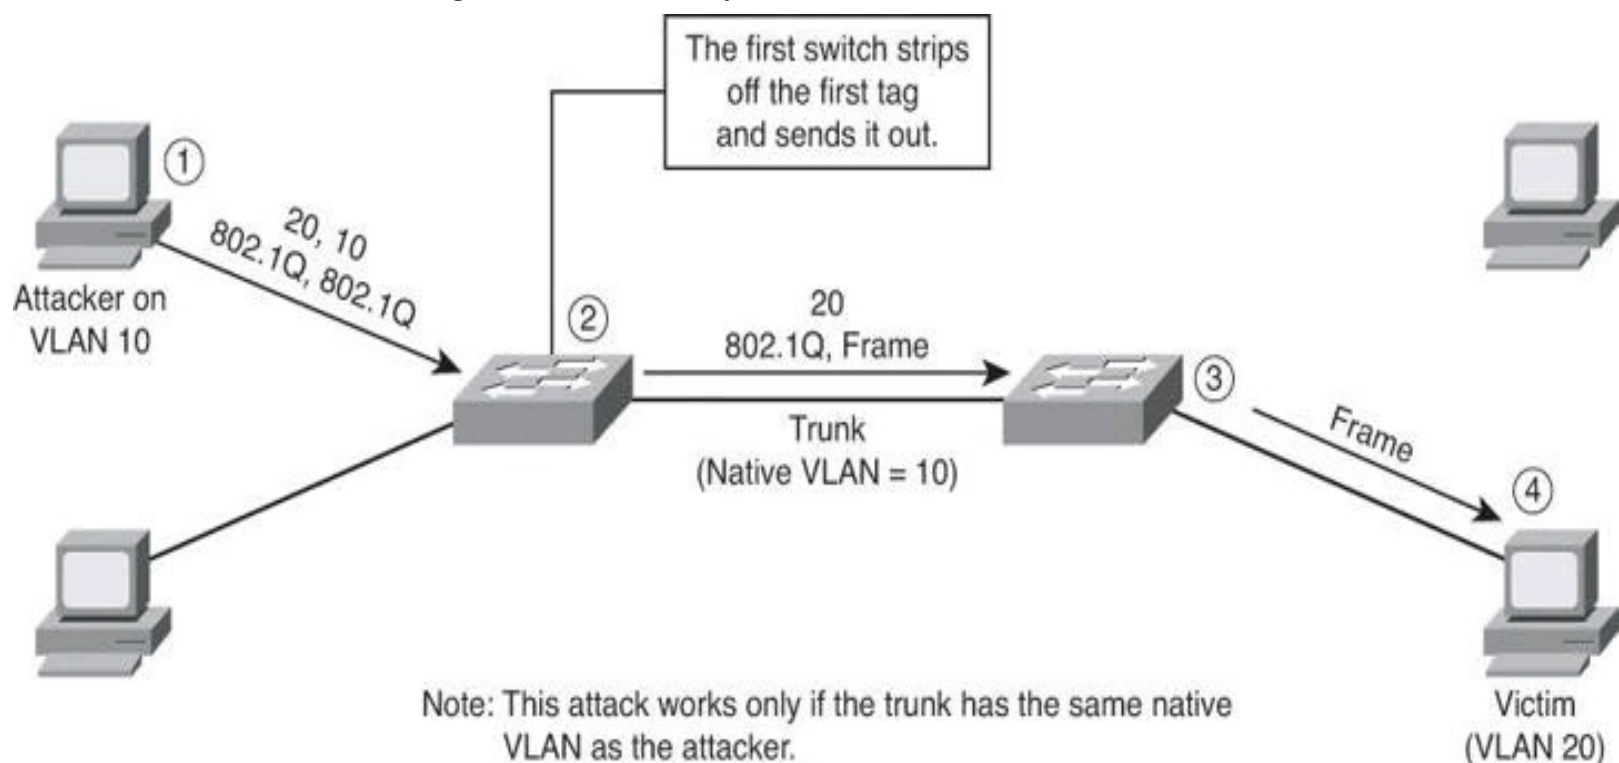


Figure 5-12. VLAN Hopping by Double Tagging

It is important to note that this attack, as shown in [Figure 5-12](#), is unidirectional and works only when the attacker and trunk port have the same native VLAN. Thwarting this type of attack is not as

easy as stopping basic VLAN hopping attacks. The best approach is to ensure that the native VLAN of the trunk ports is different from the native VLAN of the user ports.

To prevent a VLAN hopping attack that uses double 802.1Q encapsulation, the switch must look further into the packet to determine whether more than one VLAN tag is attached to a given frame. Unfortunately, the application-specific integrated circuits (ASIC) that most switches use are only hardware optimized to look for one tag and then switch the frame. The issue of performance versus security requires administrators to balance their requirements carefully.

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

Mitigating Spanning Tree Attacks

[Figure 5-13](#) shows 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. The attacker sends out BPDUs with a better bridge ID and, as a result, becomes the root bridge. Now all the traffic for this switch domain passes through the new root bridge, which is actually the attacker system.

Root Bridge Priority = 8192
MAC Address = 0000.00C0.1234

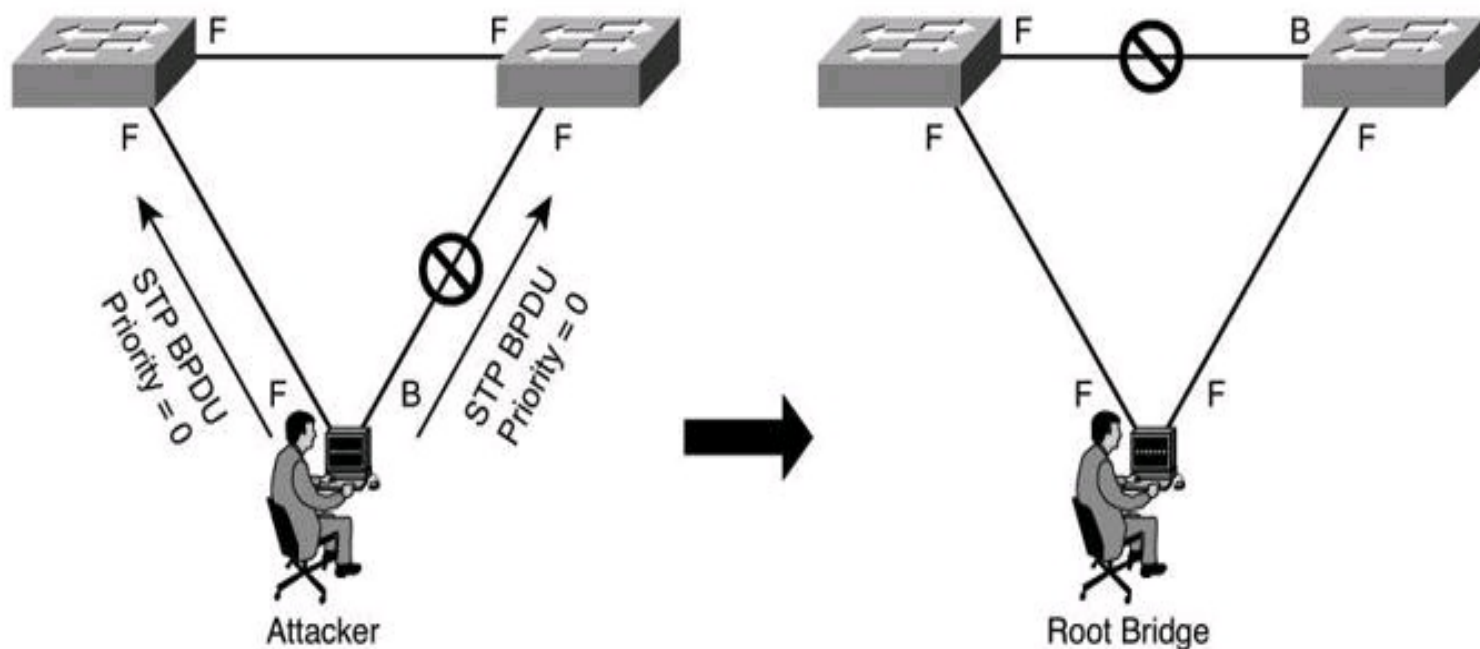


Figure 5-13. STP Manipulation

By manipulating the STP root bridge parameters, network attackers hope to spoof their system, or a rogue switch that they add to the network, as the root bridge in the topology. To do this, the network attacker broadcasts STP configuration and topology change BPDUs in an attempt to force spanning-tree recalculations. The BPDUs sent out by the system or switch of the network attacker announce that the attacking system has a lower bridge priority. If successful, the network attacker becomes the root bridge and sees a variety of frames that otherwise would not be seen.

Note

This attack can be used against all three security objectives of confidentiality, integrity, and availability.

PortFast

The spanning-tree PortFast feature causes an interface configured as a Layer 2 access port to transition from the blocking state to the forwarding state immediately, bypassing the listening and learning states. You can use PortFast on Layer 2 access ports that connect to a single workstation or server, as shown in [Figure 5-14](#), to allow those devices to connect to the network immediately, instead of waiting for spanning tree to converge.

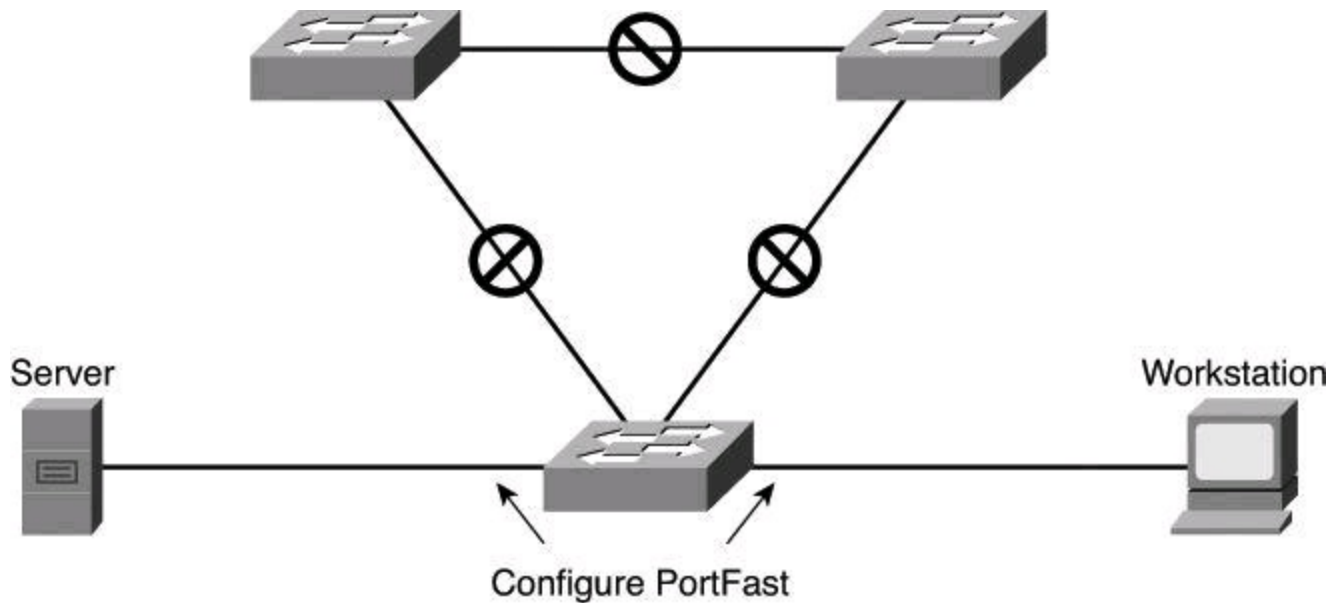


Figure 5-14. Using PortFast

If a port that is configured with PortFast receives a BPDU, STP can put the port into the blocking state by using a feature called BPDU guard.

Caution

Because the purpose of PortFast is to minimize the time that access ports must wait for spanning tree to converge, it should be used only on access ports. If you enable PortFast on a port connecting to another switch, you risk creating a spanning-tree loop.

[Table 5-3](#) lists and describes the commands that you use to implement and verify PortFast on an interface.

Table 5-3. PortFast Commands

Command	Description
Switch(config-if)# spanning-tree portfast [trunk]	Enables PortFast on a Layer 2 access port and forces it to enter the forwarding state immediately. Use the trunk keyword for trunk interfaces not connected to switches (e.g., routers or servers with trunks).
Switch(config-if)# no spanning-tree portfast	Disables PortFast on a Layer 2 access port. PortFast is disabled by default.
Switch(config)# spanning-tree portfast default	Globally enables the PortFast feature on all nontrunking ports.
Switch# show running-config interface type slot/port	Indicates whether PortFast has been configured on a port.

BPDU Guard

To mitigate STP manipulation, use the BPDU guard and root guard enhancement commands available on Cisco switches to enforce the placement of the root bridge in the network and enforce the STP domain borders.

The STP BPDU guard feature is designed to enable network designers to keep the active network topology predictable. BPDU guard is used to protect the switched network from the problems that may be caused by the receipt of BPDUs on ports that should not be receiving them. The receipt of unexpected BPDUs might be accidental or might be part of an unauthorized attempt to add a switch to the network.

BPDU guard is best deployed toward user-facing ports to prevent rogue switch network extensions by an attacker.

The global command to activate BPDU guard on all ports with PortFast enabled is as follows:

```
Switch(config)# spanning-tree portfast bpduguard default
```

In [Figure 5-15](#), the attacker starts sending out spoofed BPDUs in an effort to become the root bridge. Upon receipt of a BPDU, the BPDU guard feature disables the port.

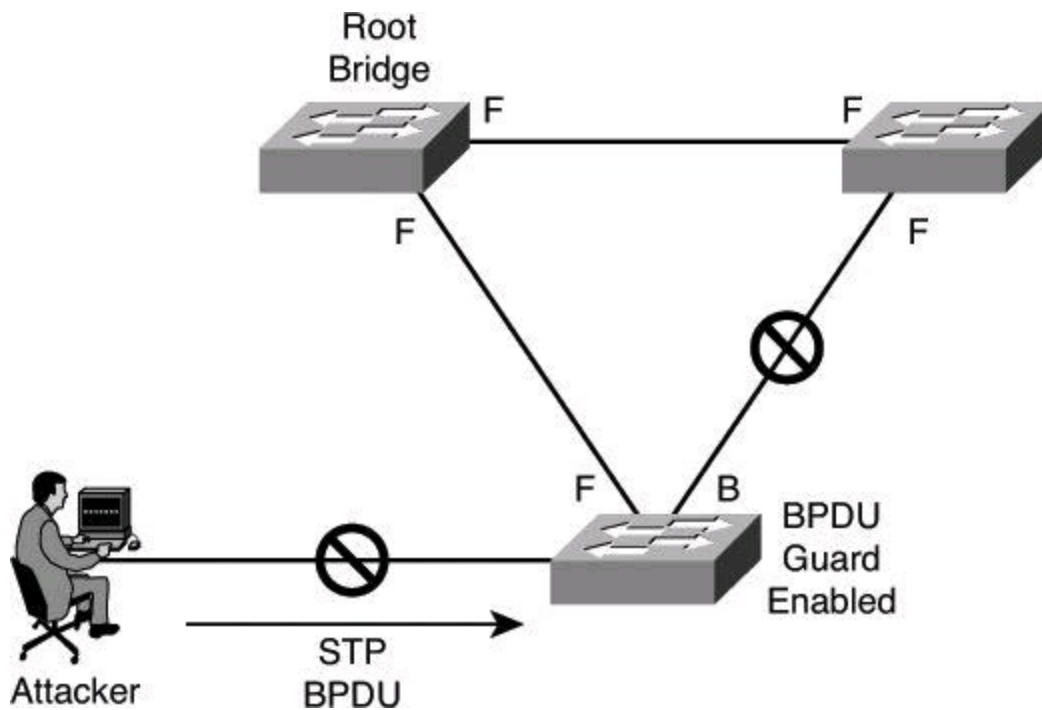


Figure 5-15. BPDU Guard

BPDU Filters

Another command used to prevent BPDU filtering, which prevents a port from sending and receiving BPDUs, is the following the interface command:

```
Switch(config-if) # spanning-tree bpduguard enable
```

Be careful when you enter this command because it overrides the PortFast configuration, explained previously.

This command has three states:

```
Switch(config-if) # spanning-tree bpduguard enable
```

This command state unconditionally enables BPDU filtering on the interface.

```
Switch(config-if) # spanning-tree bpduguard disable
```

This command state unconditionally disables BPDU filtering on the interface.

```
Switch(config-if) # no spanning-tree bpduguard
```

This command state enables BPDU filtering on the interface if the interface is in operational PortFast state and if you configure the **spanning-tree portfast bpduguard default** command.

Applying BPDU Guard Globally Versus Per Port

At the global level, you can enable BPDU guard on PortFast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. In a valid configuration, PortFast-enabled ports do not receive BPDUs. Receiving a BPDU on a PortFast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature

puts the port into the error-disabled state.

At the interface level, you can enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the PortFast feature. When the port receives a BPDU, it is put into the error-disabled state.

Root Guard

The root guard feature of Cisco switches is designed to provide a way to enforce the placement of root bridges in the network. Root guard limits the switch ports out of which the root bridge can be negotiated. Root guard is configured on a per-port basis. If a root guard-enabled port receives BPDUs that are superior to those that the current root bridge is sending, that port is moved to a root-inconsistent state, which is effectively equal to an STP listening state, and no data traffic is forwarded across that port. When the port stops receiving superior BPDUs, it will be unblocked again and will transition through STP states like any other port.

Because an administrator can manually set the bridge priority of a switch to zero, root guard might seem unnecessary. However, setting the priority of a switch to zero does not guarantee that switch will be elected as the root bridge, because another switch could have a priority of zero and a lower MAC address, and therefore a lower BID.

Root guard is best deployed toward ports that connect to switches that should not be the root bridge.

Recovery requires no intervention. A root guard port is in an STP-designated port state. When root guard is enabled on a port, the switch does not allow that port to become an STP root port. The port remains an STP-designated port.

The command to enable root guard on a per-interface basis is as follows:

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

In [Figure 5-16](#), the attacker starts sending out spoofed BPDUs in an effort to become the root bridge. Upon receipt of a BPDU, the switch with the root guard feature configured on that port ignores the BPDU and puts the port in a root-inconsistent state. The port will recover as soon as the offending BPDUs cease.

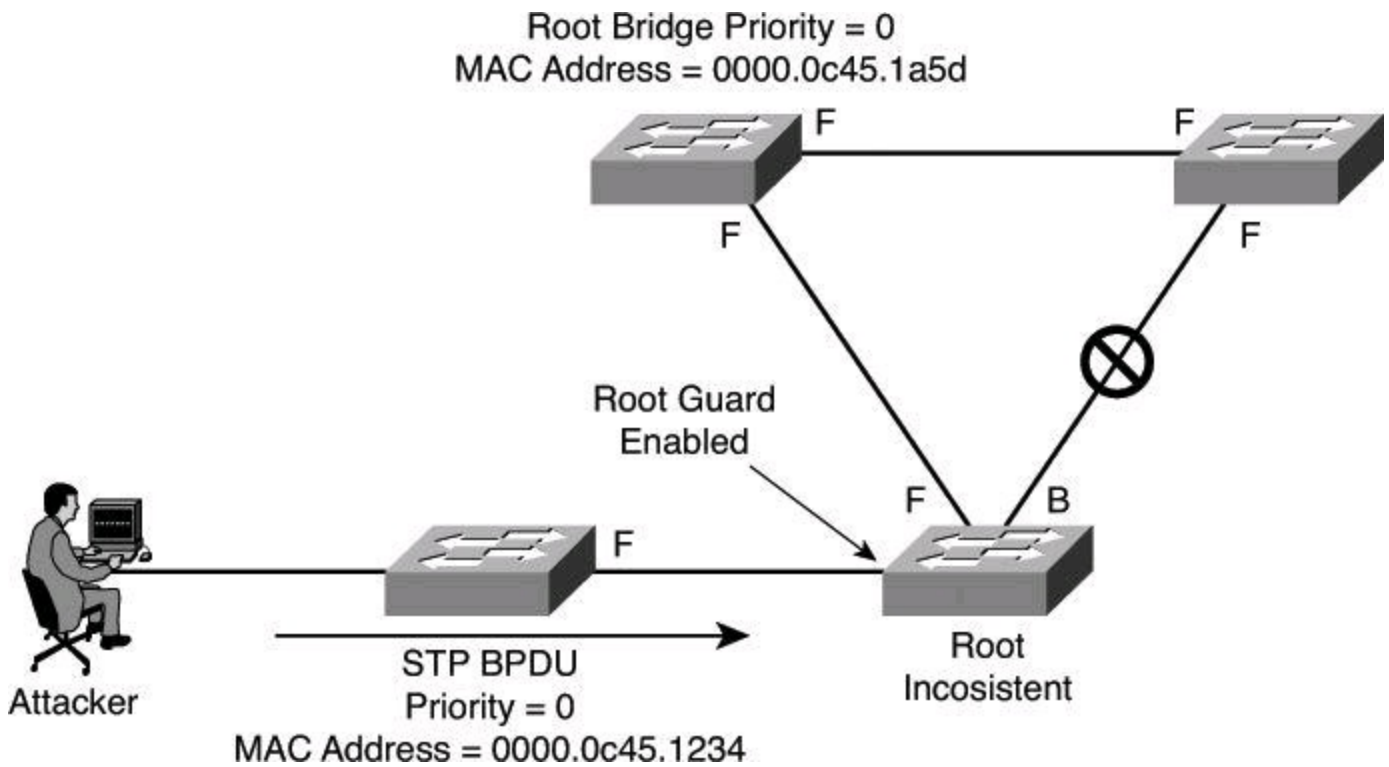


Figure 5-16. Root Guard

Confirming Spanning-Tree State and BPDU Guard

To display information about the state of spanning tree and BPDU guard, use the **show spanning-tree summary** command, as shown in [Example 5-9](#).

Example 5-9. Status of BPDU Guard with the *show spanning-tree summary* Command

[Click here to view code image](#)

```
Switch# show spanning-tree summary
Root bridge for: Bridge group 1, VLAN0001, VLAN0004-VLAN1005
VLAN1013-VLAN1499, VLAN2001-VLAN4094
EtherChannel misconfiguration guard is enabled
Extended system ID is enabled
Portfast is enabled by default
PortFast BPDU Guard is enabled
Portfast BPDU Filter is disabled by default
Loopguard is disabled by default
UplinkFast is disabled
BackboneFast is disabled
Pathcost method used is long
<output omitted>
Switch#
```

Mitigating CAM Table Overflow Attacks

The CAM table in a switch contains the MAC addresses that can be reached off a given physical port of a switch and the associated VLAN parameters for each. When a Layer 2 switch receives a frame, the switch first populates its CAM table by creating an entry listing the source MAC address of the frame it just received and the port it was received on. This entry has an idle timeout of 5

minutes; that is, if after 5 minutes no frames have been received from that MAC address on that port, the entry will be flushed out of the CAM table.

Once the switch has populated the CAM table with the source address information, it looks in the CAM table for the destination MAC address. If an entry exists for the MAC address in the CAM table, the switch forwards the frame to the MAC address port designated in the CAM table. If the MAC address does not exist in the CAM table, the switch acts like a hub and forwards the frame out every port on the switch.

The key to understanding how CAM table overflow attacks work is to know that CAM tables are limited in size. MAC flooding takes advantage of this limitation by bombarding the switch with fake source MAC addresses until the switch CAM table is full. 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.

In a CAM table overflow attack, a network intruder floods the switch with a large number of invalid source MAC addresses until the CAM table fills up. When that occurs, the switch begins to flood all incoming traffic to all ports because there is no room in the CAM table to learn any legitimate MAC addresses. The switch, in essence, acts like a hub. As a result, the attacker can see all the frames sent from a victim host to another host without a CAM table entry. CAM table overflow floods traffic only within the local VLAN so that the intruder will see only traffic within the local VLAN to which the intruder is connected. If the intruder does not maintain the flood of invalid source MAC addresses, the switch eventually ages out older MAC address entries from the CAM table and begins to act like a switch again.

In [Figure 5-17](#), the **macof** program is running on Host C. This tool floods a switch with packets that contain randomly generated source and destination MAC and IP addresses. Over a short period, the CAM table in the switch fills up until it cannot accept new entries. When the CAM table fills up, the switch begins to flood all frames that it receives.

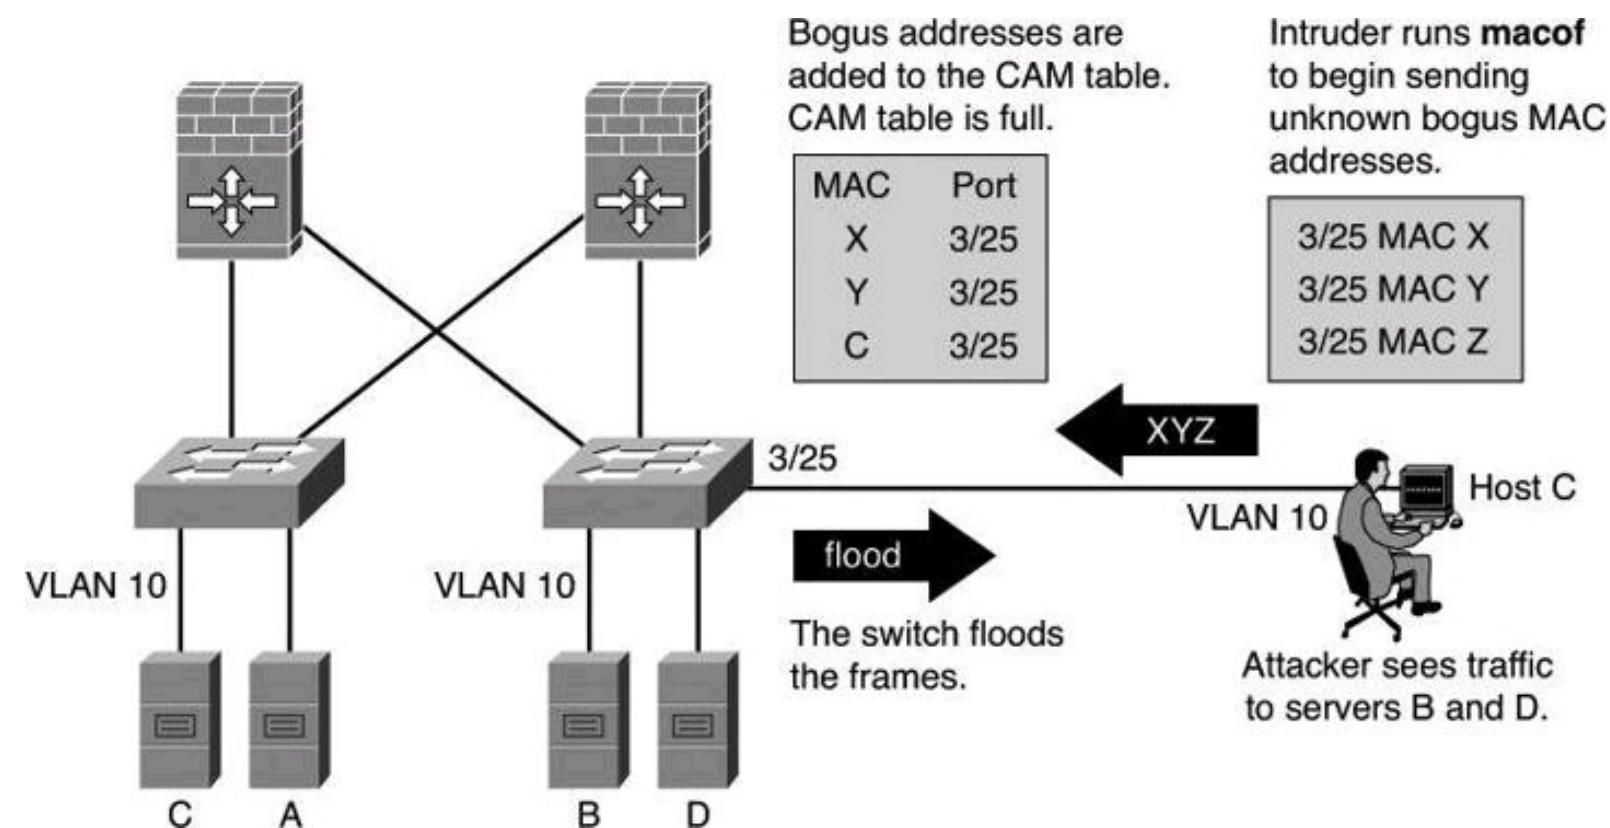


Figure 5-17. CAM Table Overflow Attack

As long as **macof** is left running, the CAM table on the switch remains full. When this happens, the switch begins to flood all received frames out every port so that frames sent from any host are also flooded out of port 3/25 on the switch.

The CAM table overflow attack can be mitigated by configuring port security on the switch. With port security (discussed later in this chapter), you can either statically specify the MAC addresses on a particular switch port or allow the switch to dynamically learn a fixed number of MAC addresses for a switch port. Statically specifying the MAC addresses on switch ports is far too unmanageable a solution for a production environment; allowing the switch to dynamically learn a fixed number of MAC addresses for a port is a more administratively scalable solution.

Mitigating MAC Address Spoofing 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 instead. Until the host sends traffic, it does 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.

[Figure 5-18](#) 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 (attacker) sends out a packet identifying itself with the source MAC address of Host A. 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 and not to Host A.

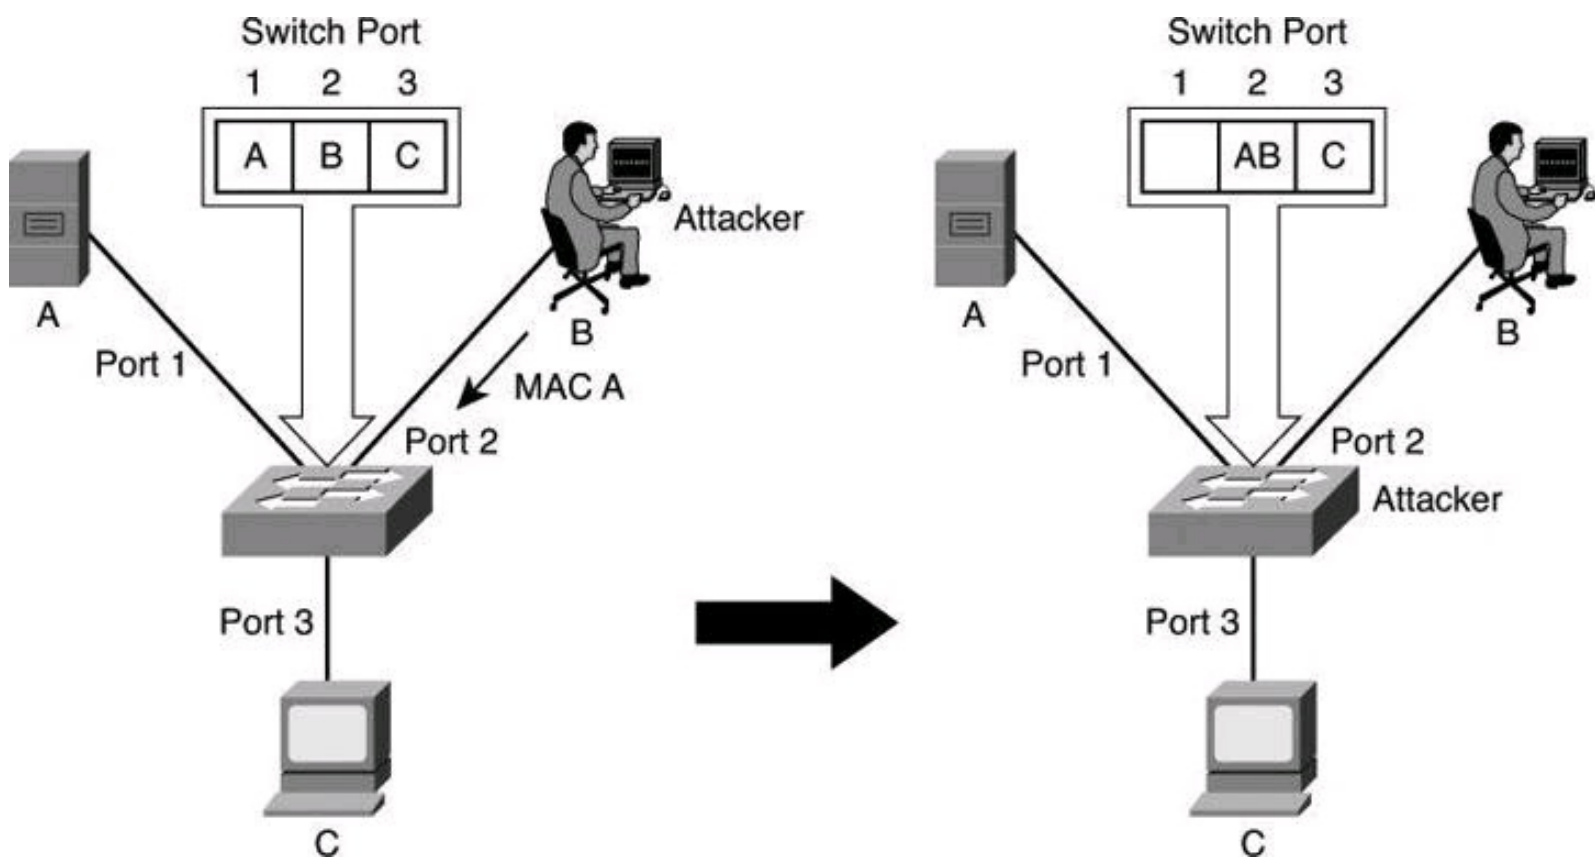


Figure 5-18. MAC Address Spoofing Attack

This attack can also be mitigated by using port security.

Using Port Security

You can use the port security feature to restrict input to an interface by limiting and identifying the MAC addresses of the stations that are allowed to access the port. When you assign 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 statically specify MAC addresses for a port or permit the switch to dynamically learn a limited number of MAC addresses. By limiting the number of permitted MAC addresses on a port to one, you can use port security to control unauthorized expansion of the network.

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, either the port shuts down until it is administratively enabled (default mode) or the port drops incoming packets from the unsecure host. The behavior of the port depends on how you configure it to respond to a security violation. In [Figure 5-19](#), traffic from Attacker 1 and Attacker 2 will be dropped at the switch because the source MAC addresses of these frames do not match MAC addresses in the list of secured (allowed) addresses.

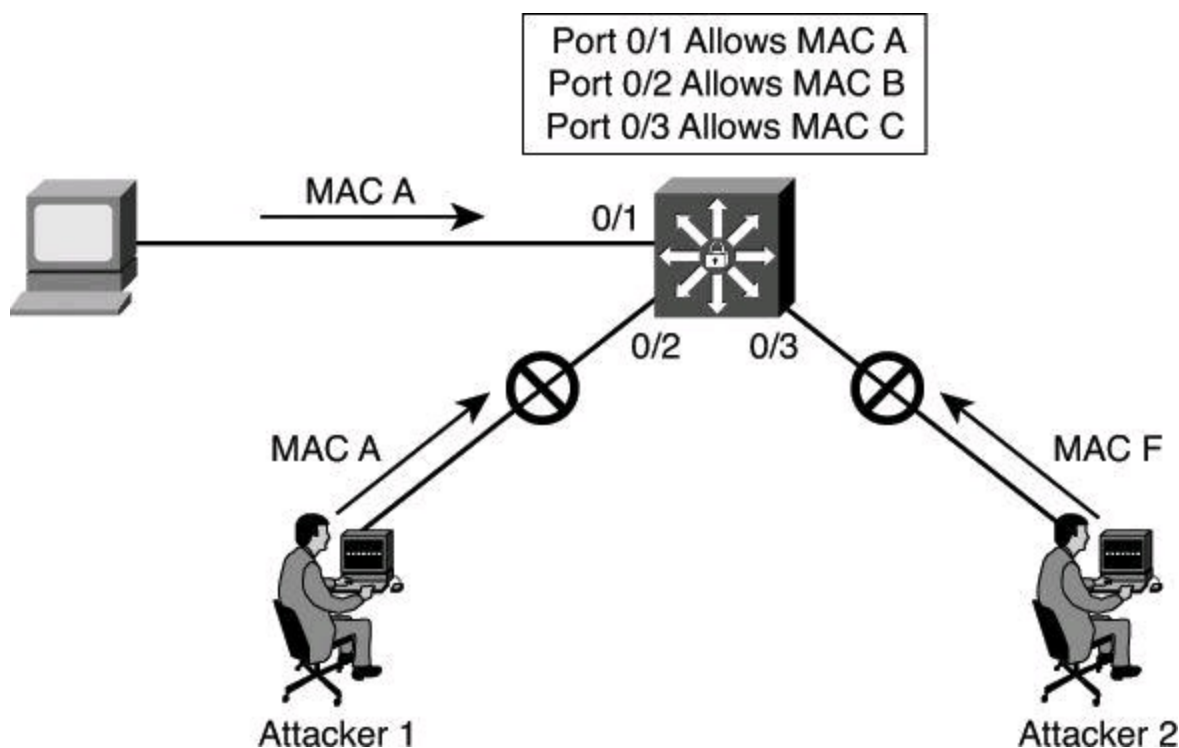


Figure 5-19. Port Security

It is recommended that you configure the port security feature to shut down a port instead of just dropping packets from insecure hosts. If port security does not shut down a port, it is possible that there will be too much load from an attack, and the port will be disabled anyway.

[Table 5-4](#) summarizes the effect of each violation mode. The parameters used to configure these violation modes will be presented later in this chapter.

Table 5-4. Configurable Port Security Violation Modes

Violation Mode	Traffic Is Forwarded	Sends SNMP Trap	Sends Syslog Message	Displays Error Message	Violation Counter Increment	Shuts Down Port
Protect	No	No	No	No	No	No
Restrict	No	Yes	Yes	No	Yes	No
Shutdown	No	Yes	Yes	No	Yes	Yes
Shutdown VLAN	No	Yes	Yes	No	Yes	No

Tip

Port security protects against too many MAC addresses per port and can dictate which MAC address is allowed to connect against which port. However, if the hacker spoofs the MAC address permitted on that port, he will gain access to the network. If you are concerned by spoofed MAC addresses, then consider implementing an 802.1X authentication solution.

Errdisable Recovery

The errdisable recovery feature also allows you to monitor spanning tree violations. If enabled with the **errdisable recovery** command, this feature monitors ports in configurable intervals to determine their stance in terms of these violations. The feature actually tries to recover the operational status of the ports when it finds them to be in violation of the policy, making the process automatic and the recovery automated. If you do not enable the recovery for the cause, the port stays in the error-disabled state until you enter the **shutdown** and **no shutdown** interface configuration commands. If you enable the recovery for a cause, the port is brought out of the error-disabled state and allowed to retry the operation when all the causes have timed out.

[Example 5-10](#) illustrates the syslog message that is generated upon a security violation. In this example, port Gi4/1 has been disabled due to a violation of the BPDU guard feature. The **show interfaces status** command displays the err-disabled status for the port.

Example 5-10. Verifying the Port Status with the *show interfaces interface status* Command

[Click here to view code image](#)

```
switch# show interfaces gigabitethernet 4/1 status
Port    Name              Status  Vlan  Duplex  Speed  Type
Gi4/1   err-disabled      100    full  1000    1000  BaseSX
```

If errdisable recovery monitoring is enabled, you can see more detailed information as to the monitored features using the **show errdisable recovery** command, as shown in [Example 5-11](#). There,

the BPDU guard feature is being monitored every 300 seconds, and one port, Gi4/1, is error-disabled for another 290 seconds (unless you use the **shutdown/no shutdown** commands to enable it and no further violations occur). The 300 seconds is a configurable option.

Example 5-11. *show errdisable recovery* Command Output

[Click here to view code image](#)

```
switch# show errdisable recovery
ErrDisable Reason      Timer Status
-----
Ulld                   Disabled
Bpduguard              Enabled
security-violatio     Disabled
channel-misconfig     Disabled
<output omitted>
Timer interval: 300 seconds
Interfaces that will be enabled at the next timeout:
Interface      Errdisable reason      Time left(sec)
-----
Gi4/1         bpduguard              290
```

To configure port security on an access port, follow these steps (see [Table 5-5](#) for command details).

Table 5-5. *switchport port-security* Command Parameters

Parameter	Description
<code>mac-address mac-address</code>	(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.
<code>mac-address sticky [mac-address]</code>	(Optional) Enable the interface for sticky learning by entering only the <code>mac-address sticky</code> 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 <code>mac-address sticky mac-address</code> keywords. Note: Although you can specify a sticky secure MAC address by entering the <code>mac-address sticky mac-address</code> keywords, it is recommended that you use the <code>mac-address mac-address</code> interface configuration command to enter static secure MAC addresses.
<code>maximum value</code>	(Optional) Set the maximum number of secure MAC addresses for the interface. The maximum number of secure MAC addresses that you can configure on a switch is set by the maximum number of available MAC addresses allowed in the system. The active Switch Database Management (SDM) template determines this number. This number represents the total available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces. The default setting is 1.
<code>vlan [vlan-list]</code>	(Optional) For trunk ports, you can set the maximum number of secure MAC addresses on a specific VLAN. If the <code>vlan</code> keyword is not entered, the default value is used. vlan: set a per-VLAN maximum value. vlan <i>vlan-list</i>: set a per-VLAN maximum value on a range of VLANs separated by a hyphen or a series of VLANs separated by commas. For nonspecified VLANs, the per-VLAN maximum value is used.

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

An interface in the default mode (dynamic desirable) cannot be configured as a secure port.

Step 3. Enable port security on the interface:

[Click here to view code image](#)

```
Switch(config-if)# switchport port-security [mac-address  
mac-address  
[vlan {vlan-id | {access | voice} } ] ] | [mac-address  
sticky  
[mac-address| vlan {vlan-id | {access | voice} } ]]  
[maximum value  
[vlan {vlan-list | {access | voice} } ]]
```

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.

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

[Click here to view code image](#)

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

[Table 5-6](#) provides the details of the **switchport port-security violation** command parameters.

Table 5-6. *switchport port-security violation* Parameters

Parameter	Description
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 reenale it by entering the shutdown and no shutdown interface configuration commands.
shutdown vlan	Set the security violation mode to per-VLAN shutdown. In this mode, only the VLAN on which the violation occurred is error-disabled.

Tip

When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause secure-violation** global configuration command, or you can manually reenale 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
```

In addition to port security, consider the commands shown in [Table 5-7](#) for making the port more secure.

Table 5-7. *switchport* Command Parameters

Parameter	Description
<code>vlan <i>vlan-id</i></code>	(Optional) On a trunk port only, specify the VLAN ID and the MAC address. If no VLAN ID is specified, the native VLAN is used.
<code>vlan access</code>	(Optional) On an access port only, specify the VLAN as an access VLAN.
<code>vlan voice</code>	(Optional) On an access port only, specify the VLAN as a voice VLAN. Note: The <code>voice</code> keyword is available only if voice VLAN is configured on a port and if that port is not the access VLAN.

Use the **no switchport port-security** interface configuration command to return the interface to the default condition of not being a secure port. The sticky secure addresses remain part of the running configuration. To remove the sticky secure addresses from the running configuration, use the **no mac-address *mac-address*** command.

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).

You can use port security aging to set the aging time for static and dynamic secure addresses on a port. Each port supports two types of aging:

- **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 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.

Use the **switchport port-security aging {static | time *time* | type {absolute | inactivity} }** command to enable or disable static aging for the secure port, or set the aging time or type. [Table 5-8](#) provides the details of the **switchport port-security aging** parameters.

Table 5-8. *switchport port-security aging* Parameters

Parameter	Description
<code>static</code>	This command option enables aging for statically configured secure addresses on this port.
<code>time <i>time</i></code>	This command option specifies the aging time for this port. The range is 0 to 1440 minutes. If the time is 0, aging is disabled for this port.
<code>type absolute</code>	This command option sets the aging type to absolute. All the secure addresses on this port age out exactly after the time (minutes) specified and are removed from the secure address list.
<code>type inactivity</code>	This command option sets the aging type to inactivity. The secure addresses on this port age out only if there is no data traffic from the secure source address for the specified time period.

[Example 5-12](#) shows a typical port security configuration for a voice port. Two MAC addresses are allowed, and they are to be learned dynamically. One MAC address is for the IP phone, and the other IP address is for the PC connected to the IP phone. Violations of this policy result in the port being shut down, and the aging timeout for the learned MAC addresses is set to two hours.

Example 5-12. Port Security Configuration

[Click here to view code image](#)

```
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 2
Switch(config-if)# switchport port-security violation shutdown
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security aging time 120
```

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

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.

[Example 5-13](#) shows that port security is enabled on port Fa0/12 with a maximum MAC address count of 2. Currently, there are no MAC addresses learned on that port, and the violation action has been set to shut down the port.

Example 5-13. *show port-security* Command Output

[Click here to view code image](#)

```
sw-class# show port-security
```

Secure Port Action	MaxSecureAddr (Count)	CurrentAddr (Count)	SecurityViolation (Count)	Security
Fa0/12	2	0	0	Shutdo

Total Addresses in System (excluding one mac per port) : 0
Max Addresses limit in System (excluding one mac per port) : 1024

[Example 5-14](#) demonstrates output from the **show port-security interface fa0/12** command, revealing that a violation has occurred, which means that more than one MAC address has been seen on the port. The port has been shut down because of this policy violation, as confirmed by the secure-down port status.

Example 5-14. *show port-security interface fa0/12* Command Output

[Click here to view code image](#)

```
sw-class# show port-security interface fa0/12
Port Security          : Enabled
Port status           : Secure-down
Violation mode        : Shutdown
Maximum MAC Addresses : 1
Total MAC Addresses   : 2
Configured MAC Addresses : 0
Aging time            : 120 mins
Aging type            : Absolute
SecureStatic address aging : Disabled
Security Violation Count : 1
```

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

[Example 5-15](#) shows that port Fa0/12 is in VLAN 1 and has a secured MAC address of 0000.ffff.aaaa, which means that the host with the 0000.ffff aaaa MAC address can connect to port Fa0/12.

Example 5-15. *show port-security address* Command Output

[Click here to view code image](#)

```
sw-class# show port-security address
Secure Mac Address Table
-----
Vlan  Mac Address      Type                Ports Remaining Age
-----
1     0000.ffff.aaaa    SecureConfigured    Fa0/12             -
```

Using SNMP to Monitor Access to Switch Port

Network managers need a way to monitor who is using the network and where they are. In [Figure 5-20](#), if port Fa2/1 is secure, an SNMP trap will be generated when MAC D disappears from the CAM table of the switch.

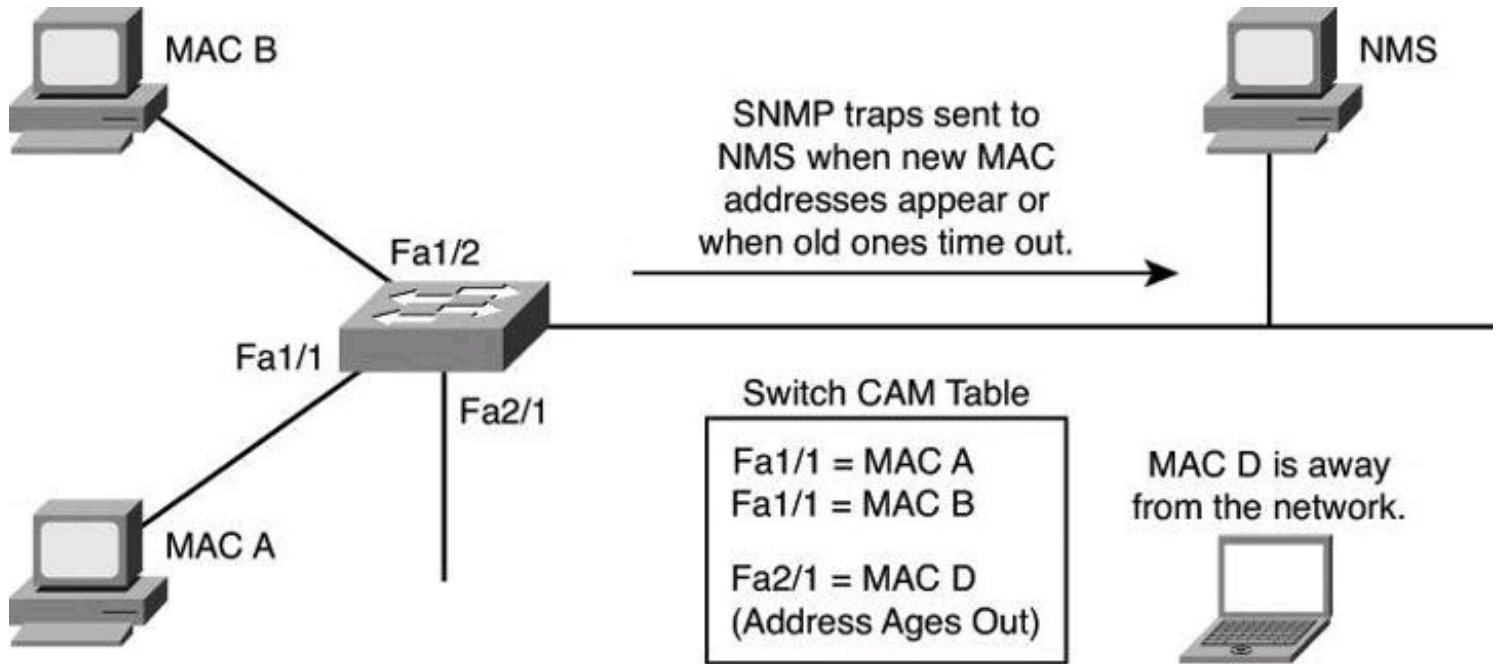


Figure 5-20. Notification of Intrusions

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

MAC address notification allows the network administrator to monitor MAC addresses that are learned and MAC addresses that age out and are removed from the switch.

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

Summary

Layer 2 security is often an overlooked aspect of network security. Buffer overflows can be the worst of these problems. The goals of endpoint security include protection from viruses, worms, and Trojan horses. SAN and voice security are also increasingly important because these technologies are growing in popularity in the modern enterprise.

The major points that were covered in this chapter are as follows:

- VLAN and trunks are susceptible to attacks such as VLAN hopping.
- A switched network can be attacked by propagating erroneous spanning tree information

between participants.

- Through proper planning and implementation, a network security strategy can effectively protect the switched data plane.
- VLAN hopping and MAC spoofing attacks can be defeated by adopting effective protection of the switch data plane, such as
 - Using a dedicated VLAN ID for trunk ports while not using VLAN 1 for anything
 - Setting user ports to nontrunking (unless you are using Cisco VoIP)
 - Using port security whenever possible for access ports and enabling STP attack mitigation (BPDU guard, root guard)
- VLAN hopping and MAC spoofing attacks, which are possible on switched networks and measures that should be put in place to protect against those attacks

References

For additional information, refer to these Cisco.com resources:

Private VLAN Catalyst Switch Support Matrix (Cisco Catalyst 6500 Series Switches), <http://tinyurl.com/2w22d6>

“Identity Based Networking Services,” http://www.cisco.com/en/US/products/ps6638/products_ios_protocol_group_home.html

“LAN Security: Introduction,” <http://tinyurl.com/594lpb>

Securing Networks with Private VLANs and VLAN Access Control Lists (Cisco Catalyst 6000 Series Switches), http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a0080

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

1. Which of the following is a valid statement regarding assigning voice traffic to specific VLANs?
 - a. Separation of voice from data traffic increases total throughput.
 - b. Separation of voice from data traffic ensures communication privacy and nonrepudiation.
 - c. Separation of voice from data traffic is seldom used in the industry.
 - d. Separation of voice from data traffic makes it easier to apply VLAN access list.
2. Which of the following commands should be used on a trunk port when attempting to protect against VLAN hopping? (Choose all that apply.)
 - a. **switchport mode access**
 - b. **switchport mode trunk**
 - c. **switchport nonegotiate**
 - d. **switchport trunk native vlan**

- 3.** Which two commands best protect a switched network from a hacker who is trying to preempt an election of the Spanning Tree Protocol?
- a. spanning-tree portfast bpduguard**
 - b. spanning-tree portfast default**
 - c. spanning-tree guard root**
 - d. switchport port-security violation**
- 4.** Which command limits the number of MAC addresses communicating through the same switch port?
- a. switchport mode access**
 - b. switchport port-security maximum**
 - c. switchport-security mac-address sticky**
 - d. switchport -security violation**
- 5.** Why should you worry about Layer 2 security?
- a.** Switches cannot regulate the flow of data between their ports.
 - b.** You don't have to worry about Layer 2 security because it is lower than the IP layer and most attacks happen at the network layer.
 - c.** With the domino effect, compromising Layer 2 means compromising the higher layers.
 - d.** VLANs are a Layer 3 function in a switch, and therefore, as with any other Layer 3 processes, it can be easily hacked.
- 6.** Put the following steps of a VLAN hopping attack in the proper order.
- a.** The frame arrives at the second switch but has no knowledge that it was supposed to be for VLAN 10.
 - b.** The second switch looks only at the 802.1Q tag (the former inner tag that the attacker sent) and sees that the frame is destined for VLAN 20 (the victim VLAN). The second switch sends the packet on to the victim port, or floods it, depending on whether there is an existing CAM table entry for the victim host.
 - c.** The attacker sends a double-tagged 802.1Q frame to the switch. The outer header has the VLAN tag of the attacker, which is the same as the native VLAN of the trunk port. For the purposes of this example, assume that this is VLAN 10. The inner tag is the victim VLAN, VLAN 20.
 - d.** The frame arrives on the first switch, which looks at the first 4-byte 802.1Q tag. The switch sees that the frame is destined for VLAN 10 and sends it out all VLAN 10 ports (including the trunk), because there is no CAM table entry. The switch does not add a VLAN 10 tag to the frames because VLAN 10 is the native VLAN, and as specified by the 802.1Q specification, native VLAN traffic is not tagged. At this point, the second VLAN tag is still intact and has not been inspected by the first switch.
- 7.** The VLAN hopping attack is the result of which condition?
- a.** Layer 2 loops with STP disabled

- b.** CAM table overflow
- c.** Poor VLAN planning
- d.** Trunking protocol vulnerabilities

8. What is one of the exceptions to the recommended practice to disable trunking on all user ports on a switch?

- a.** IP phone ports
- b.** Trusted VLANs
- c.** Unreliable ports
- d.** 802.1X ports

9. Which spanning tree protection feature disables ports when a violation occurs?

- a.** Source guard
- b.** BPDU guard
- c.** PortFast
- d.** Root guard

10. Which port security mode has as its only action to drop frames arriving on that port when a violation occurs?

- a.** Shutdown
- b.** Protect
- c.** Dynamic
- d.** Restrict

Chapter 6. Securing the Data Plane in IPv6 Environments

IPv6 shares some of the same security concerns and considerations as IPv4. Some IPv6-specific vulnerabilities and threats make it unique as it relates to your considerations and strategy to protect IPv6 infrastructures and services. In this chapter, you learn how to do the following:

- Explain the need for IPv6 from the general perspective of the transition to IPv6 from IPv4
- List and describe the fundamental features of IPv6, as well as enhancements when compared to IPv4
- Analyze the IPv6 addressing scheme, components, and design principles and configure IPv6 addressing
- Describe the IPv6 routing function
- Evaluate how common and specific threats affect IPv6
- Develop and implement a strategy for IPv6 security

The Need for IPv6

The IPv4 address space provides approximately 4.3 billion addresses. Of that address space, approximately 3.7 billion addresses are actually assignable. The other addresses are reserved for special purposes such as multicasting, private address space, loopback testing, and research.

An IPv6 address is a 128-bit binary value, which can be displayed as 32 hexadecimal digits, as shown in [Figure 6-1](#). It provides 3.4×10^{38} IP addresses. This version of IP addressing should provide sufficient addresses for future Internet growth needs—enough to allocate the equivalent of the entire IPv4 address space to every person on the planet. Another analogy to show the enormity of the IPv6 address pool is to think that there are 667,712,614,478,140,039 addresses available per square meter of the Earth's surface, including ocean surface.

IPv4
32 Bits
= 4,294,967,296 Possible Addressable Nodes

IPv6
128 Bits: Four Times Larger in Bits
= $\sim 3.4 \times 10^{38}$ Possible Addressable Nodes
= 340,282,366,920,938,463,374,607,431,768,211,456
= 51,557,934,381,960,373,252,026,455,671 Addresses
per Person on the Planet !!!



Figure 6-1. Comparing IPv4 and IPv6 Addresses

So, What Ever Happened to IPv5?

Well, IPv5 never existed. In the late 1970s, a set of protocols were defined to create the

Internet Stream Protocol (ST), which was to be used, instead of IP, for streaming. The protocol was never introduced. However, the second iteration of ST (ST-II) used Internet Protocol version 5 to distinguish itself from regular IP. Though ST was never known as IPv5, IETF decided to stay away from this nomenclature to avoid confusion.

IPv6 Penetration

According to an unofficial Cisco TAC statistic, in fall 2011, about 5 percent of networks running on Cisco equipment had IPv6 in production. More formal statistics can be found at Réseaux IP Européens Network Coordination Centre, or RIPE NCC (<http://www.ripe.net/>). As one of four regional Internet registries that supply and administer IP addresses, it has some interesting statistics. According to RIPE NCC, in 2012 about 7.5 percent of networks worldwide were running IPv6. The largest penetration was in the Asia-Pacific region with close to 11 percent. The lowest percentage of networks running IPv6 was in North America at 5 percent. Interestingly, it's in North America that the highest number of IPv4 Class A addresses can be found. (Source: RIPE NCC, <https://labs.ripe.net/Members/emileaben/interesting-graph-networks-with-ipv6-over-time>.)

Google is also tracking the IPv6 adoption rate on the Internet. As of July 2012, Google was reporting that 0.78 percent of all traffic it was seeing was IPv6. (Source: <http://www.google.com/intl/en/ipv6/statistics/>.)

If you are interested in finding the IPv6 penetration rate in your country, visit <http://v6asns.ripe.net/v/6> or <http://www.google.com/ipv6/statistics.html#tab=per-country-ipv6-adoption>.

In 2012, there were approximately 2.3 billion Internet users around the world. Between 2000 and 2011, there was a 480 percent growth of the Internet. In February 2011, IANA announced the allocation of the last few blocks of /8 address spaces.

The Internet will be transformed after IPv6 fully replaces IPv4. IPv4 address exhaustion is an imminent fact. Nevertheless, IPv4 will not disappear overnight. Rather, it will coexist with and then gradually be replaced by IPv6.

These facts are augmented by the tremendous increase of mobile and consumer devices connecting to the public Internet. The *consumerization* of IT services is an ongoing trend, where business workers use their consumer and mobility devices to access corporate networks. It is expected that by 2013, there will be 1 billion internet-connected devices. These trends, along with application trends, have driven peer-to-peer communications and more demand for a larger IP address space. The next wave is machine-to-machine (M2M) communications, smart grids, networked security cameras and motion sensors, connected home appliance, and so on.

The change from IPv4 to IPv6 has already begun, particularly in Europe and the Asia-Pacific region. These areas are exhausting their allotted IPv4 addresses, which makes IPv6 all the more attractive and necessary. Some countries, such as Japan, are aggressively adopting IPv6. The European Union is moving toward IPv6, and China is considering building new networks that are dedicated for IPv6. In 2008, the U.S. government mandated all federal agencies to demonstrate IPv6

connectivity over their backbone networks and that their public-facing web sites be IPv6 ready by Sept. 30, 2012.

Other IPv4 weaknesses affect the need for IPv6. IPv4 was designed without a number of modern-day network requirements for security, device roaming, quality of service, address depletion, and others. Incorporating additional features in IPv4 has been costly in terms of complexity and flexibility.

The supporting cast of IPv4 protocols and solutions that manage address scarcity (Network Address Translation [NAT], Dynamic Host Configuration Protocol [DHCP], variable-length subnet mask [VLSM], and classless interdomain routing [CIDR]) are good examples of added complexity and lower performance due to an unplanned event such as IP address space depletion.

Security is one of those requirements. The push for peer-to-peer communications, where each peer has a public address, demands security controls that work end to end and are initiated by the IP hosts themselves. IP Security (IPsec) was introduced to solve the problem of the inherent lack of security in IPv4 transmission.

Quality of service (QoS) was also an area neglected by IPv4. Resource Reservation Protocol (RSVP) and other protocols were introduced to provide QoS.

Mobility is not built in to IPv4. Mobile IP is required to deploy a roaming approach to IP addressing and service continuity.

IPv6 Features and Enhancements

IPv6 is a powerful enhancement to IPv4. Several features in IPv6 offer functional improvements. What IP developers learned from using IPv4 suggested changes to better suit current and probable network demands:

- **Larger address space:** A larger address space includes several enhancements:
 - Improved global reachability and flexibility
 - Aggregation of prefixes that are announced in routing tables
 - Multihoming to several ISPs
 - Autoconfiguration that can include data link layer addresses in the address space
 - Plug-and-play options
 - Public-to-private readdressing end to end without address translation
 - Simplified mechanisms for address renumbering and modification
- **Simpler header:** A simpler header offers several advantages over IPv4:
 - Better routing efficiency for performance and forwarding-rate scalability
 - No broadcasts and thus no potential threat of broadcast storms
 - No requirement for processing checksums
 - Simpler and more efficient extension header mechanisms
 - Flow labels or per-flow processing with no need to open the transport inner packet to identify the various traffic flows

• **Mobility and security:** Mobility and security help ensure compliance with Mobile IP and IPsec standards functionality. Mobility enables people with mobile network devices—many with wireless connectivity—to move around in networks. Mobile IP is an Internet Engineering Task Force (IETF) standard that is available for both IPv4 and IPv6, enabling mobile devices to move without breaks in established network connections. Because IPv4 does not automatically provide this kind of mobility, you must add it with additional configurations.

- In IPv6, mobility is built in, which means that any IPv6 node can use mobility when necessary. The routing headers of IPv6 make mobile IPv6 much more efficient than Mobile IPv4 for end nodes.
 - IPsec is the IETF standard for IP network security, available for both IPv4 and IPv6. Although the functionalities are essentially identical in both environments, IPsec is mandatory in IPv6. IPsec can be used transparently on every IPv6 host without additional software, making the IPv6 Internet potentially more secure. IPsec also requires keys for each party, which implies a global key deployment and distribution.
- **Transition richness:** There are several ways to incorporate existing IPv4 capabilities with the added features of IPv6:
- One approach is to implement a dual-stack method, with both IPv4 and IPv6 configured on the interface of a network device.
 - Tunneling is another technique that will become more prominent as the adoption of IPv6 grows. There are various IPv6-over-IPv4 tunneling methods. Some methods require manual configuration, while others are more automatic.

Cisco IOS Release 12.3(2)T and later also include Network Address Translation Protocol Translation (NAT-PT) between IPv6 and IPv4. This translation allows direct communication between hosts that use different versions of the IP protocol.

IPv6 Headers

The new IPv6 header is simpler than the IPv4 header, in the following ways:

- Half of the previous IPv4 header fields are removed. This enables simpler processing of the packets, enhancing the performance and routing efficiency.
- All fields are aligned to 64 bits, which enables direct storage and access in memory by fast lookups.
- No checksum occurs at the IP layer, and no recalculation is performed by the routers. Error detection is done by the link layer and transport layer.

IPv6 header enhancements improve hardware-based processing, which provides scalability of the forwarding rate for the next generation of high-speed networks.

IPv6 uses a different approach from IPv4 to manage optional information in the header. It defines extension headers that form a chain of headers linked by the Next Header field that is contained in each extension header. This approach provides efficiency gains over IPv4 in the way that options and special functions are packaged. It enables a faster forwarding rate and leaves the router with less work to do for each packet.

IPv4 Versus IPv6 Header Fields

With IPv4, the Fragmentation field is always present in the header regardless of whether the packet is fragmented or not. With IPv6, such field would appear only if the functionality is needed.

Key
Topic

All extension headers are daisy-chained, each header pointing to the next header until they reach the transport layer data, as shown in [Figure 6-2](#). This arrangement allows an IPv6 packet to be customized with features and functionality.

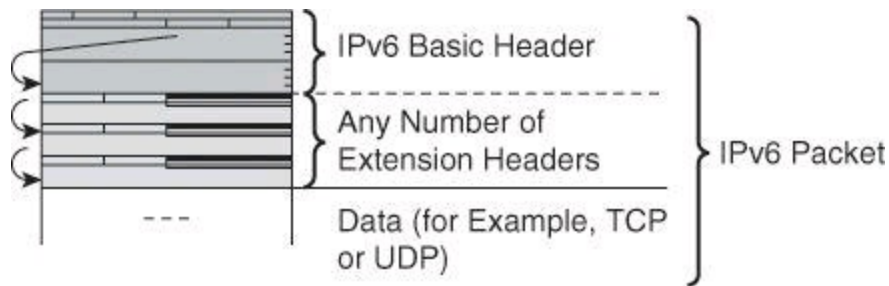


Figure 6-2. Daisy-chaining IPv6 Extension Headers

A good example of extension headers is the routing header (RH) in IPv6. It can be used to implement the source routing function that is widely known in IPv4, where the packet lists the intermediary hosts that it will visit before arriving at its final destination. If the Next Header field is equal to 0, the next header is a Hop-by-Hop field. This header contains information that must be examined by each node on the path.

A second routing header, with a value of 2, has been defined for use with IPv6 Mobility. It is formatted similarly to the type 0 routing header, except that it only carries one intermediate hop.

All IPv6 hosts, including end stations, process the RH.

Stateless Address Autoconfiguration

Stateless address autoconfiguration is an important feature of IPv6. It allows serverless basic configuration of the nodes and easy renumbering. It uses the Neighbor Discovery Protocol (NDP), which is based on ICMP version 6 (ICMPv6) protocol messaging, a topic we will discuss next. NDP replaces Address Resolution Protocol (ARP), and multicast replaces broadcast.

Stateless address autoconfiguration uses the information in the router advertisement messages to configure the node. The prefix included in the router advertisement is used as the /64 prefix for the node address. The other 64 bits are obtained by the dynamically created interface identifier, which, in the case of Ethernet, is the modified extended universal identifier 64-bit (EUI-64) format.

Router advertisements are sent periodically. When a host boots, it needs to have its address in the early stage of the boot process, as illustrated in [Figure 6-3](#). Instead of waiting for the next router advertisement to get the information to configure its interfaces, a node sends a router solicitation message asking the routers on the network to reply immediately with a router advertisement so that the

node can immediately autoconfigure. All the routers respond with a normal router advertisement message with the all-nodes multicast address as the destination address.

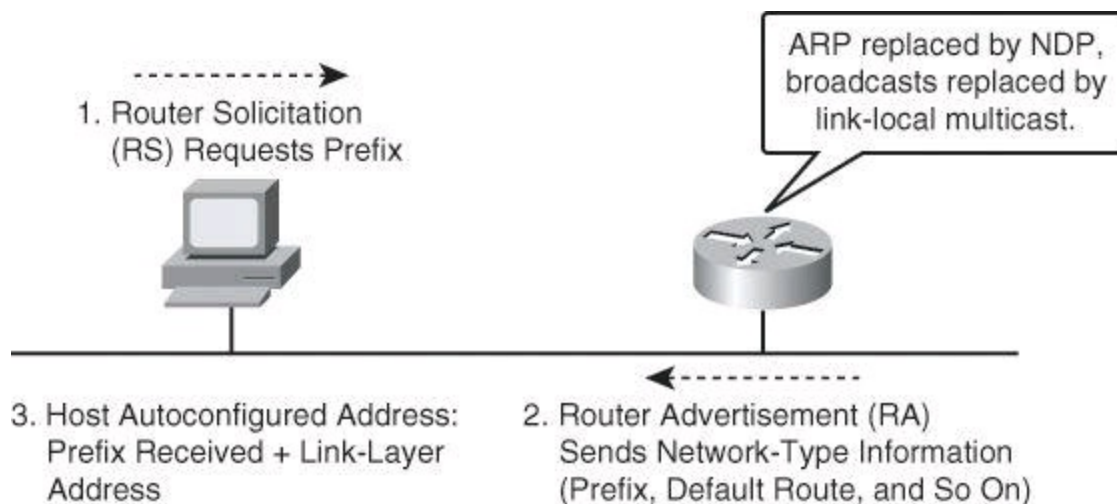


Figure 6-3. IPv6 - Stateless Autoconfiguration

Autoconfiguration enables plug-and-play configuration of an IPv6 device, which allows devices to connect themselves to the network without any configuration from an administrator and without any servers, such as DHCP servers. This key feature enables deployment of new devices on the Internet, such as smartphones, wireless devices, home appliances, and other consumer devices.

The NDP functions therefore include:

- Router, prefix, and parameter discovery
- Address autoconfiguration and resolution
- Duplicate address detection

IPv6 Network Discovery Protocol (NDP) replaces IPv4 Address Resolution Protocol (ARP). NDP is a messaging protocol that relies on ICMPv6 and facilitates the discovery of neighboring devices over a network.



Internet Control Message Protocol Version 6

ICMPv6 is similar to ICMP version 4 (ICMPv4) in that it enables nodes to make diagnostic tests and report problems. [Table 6-1](#) lists the differences between ICMPv4 and ICMPv6.

Table 6-1. ICMPv4 and ICMPv6 Message Types

ICMP Message Type	ICMPv4	ICMPv6
Connectivity Checks	X	X
Informational/Error Messaging	X	X
Fragmentation Needed Notification	X	X
Address Assignment		X
Address Resolution		X
Router Discovery		X
Multicast Group Management		X
Mobile IPv6 Support		X

Like ICMPv4, ICMPv6 implements two kinds of messages: error messages, such as “destination unreachable,” “packet too big,” or “time exceeded,” and informational messages, such as “echo request” and “echo reply.”

ICMPv4 is often blocked by security policies in corporate firewalls because of known attacks that are based on ICMP. ICMPv6 can be subject to similar attacks; however, there are substantial changes to its scope. The responsibilities of ICMPv6 go beyond mere messaging, and into address resolution and assignment, router discovery, and Mobile IP.

IPv6 General Features

IPv6 has many advantages over IPv4, such as

- A large address space makes global reachability possible from every IPv6 node.
- Autoconfiguration is essential for deploying many appliances. It would not be possible, practically, to manually configure IP addresses. You need some autoconfiguration mechanism that scales. DHCP may not be the right way to manage millions of clients.
- IPsec is mandated in the architecture.
- NAT includes end-to-end security in networks by requiring that you trust the end devices.
- Mobile IPv6 improves routing efficiency over IPv4.
- IPv6 is the same as IPv4 in QoS and header compression features. Both areas benefited from the work on IPv6. The IPv6 header compresses better than the IPv4 header because there are fewer fields.
- Other features are equivalent, except for a few details, such as scoped addresses (defined below) in multicast, or the concept of stateless DHCP where only static parameters are provided by the DHCP server.

RFC 4007 defines scoped addresses as follows: “Internet Protocol version 6 includes support for addresses of different ‘scope’; that is, both global and non-global (e.g., link-local) addresses.”

[Table 6-2](#) summarizes the general features and benefits of IPv6 and compares them to IPv4.

Table 6-2. IPv4 and IPv6 Compared

IP Service	IPv4	IPv6
Addressing range	32-bit, NAT	128-bit, multiple scopes
Autoconfiguration	DHCP	Stateless, stateful (DHCPv6)
Security	IPsec	IPsec mandated, works end to end
Mobility	Mobile IP	Mobile IP with optimized routing
Quality of service	Differentiated service, integrated service	Differentiated service, integrated service
IP Multicast	Heavy application use	Heavy application and protocol stack use
ICMP	Messaging mostly	Messaging and protocol functions

Transition to IPv6

The transition from IPv4 does not require upgrades on all nodes at the same time. Many transition mechanisms enable smooth integration of IPv4 and IPv6. Other mechanisms that allow IPv4 nodes to communicate with IPv6 nodes are available. All of these mechanisms are applied to different situations. Tunnels are established manually, semiautomatically, or automatically:

- Manually
 - IPv6-in-IPv4
 - GRE (not discussed)
 - VPN (not discussed)
- Semiautomatically
 - Tunnel broker (proxying)
- Automatically
 - 6to4
 - ISATAP
 - Teredo

The three most common techniques to transition from IPv4 to IPv6 are as follows:

- **Dual stack:** Dual stack is an integration method in which a node has implementation and connectivity to both an IPv4 network and an IPv6 network. As a result, the node and its corresponding routers have two protocol stacks, as shown in [Figure 6-4](#).
- **Tunneling:** There are several tunneling techniques available:
 - **Manual IPv6-over-IPv4 tunneling:** An integration method in which an IPv6 packet is encapsulated within the IPv4 protocol. This method requires dual-stack routers.
 - **Dynamic 6to4 tunneling:** A method that automatically establishes the connection of IPv6 islands through an IPv4 network, typically the Internet. The

6to4 tunneling method dynamically applies a valid, unique IPv6 prefix to each IPv6 island, which enables the fast deployment of IPv6 in a corporate network without address retrieval from the ISPs or registries.

- **Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunneling:** An automatic overlay tunneling mechanism that uses the underlying IPv4 network as a data link layer for IPv6. ISATAP tunnels allow individual IPv4 or IPv6 dual-stack hosts within a site to communicate with other such hosts on a virtual link, creating an IPv6 network using the IPv4 infrastructure.

- **Teredo tunneling:** An IPv6 transition technology that provides host-to-host automatic tunneling instead of gateway tunneling. It is used to pass unicast IPv6 traffic when dual-stacked hosts (hosts that are miming both IPv6 and IPv4) are located behind one or multiple IPv4 NATs.

- **Proxying and translation (NAT-PT):** A translation mechanism that sits between an IPv6 network and an IPv4 network. The job of the translator is to translate IPv6 packets into IPv4 packets and vice versa.




Figure 6-4. IPv6-in-IPv4 Tunnel

IPv6 Addressing

IPv6 addresses are, at first, the most noticeable change compared to IPv4. Not only are IP addresses going from 32 bits to 128 bits, but how addresses are represented and how they are classified are also new. There are different types of IPv6 addresses and different representations.

IPv6 Address Representation

IPv6 addresses are 128 bits long. Addresses are represented as a series of eight 16-bit hexadecimal fields that are separated by colons. The A, B, C, D, E, and F in hexadecimal fields are not case sensitive. A typical IPv6 address would therefore look like 2001:cb7:46d1:0:0:8a2e:370:7334.

Here are some ways to shorten the writing of IPv6 addresses:

- The leading zeros in a field are optional, so 010F = 10F and 0000 = 0.
- Successive fields of zeros can be represented as a double colon (::), but only once in an address. An address parser is able to identify the number of missing zeros by separating the two parts and filling in zeros until the 128 bits are completed. But if two double colons are

placed in the address, there is no way to identify the size of each block of zeros. Therefore, only one double colon is possible in a valid IPv6 address.

The use of the double-colon technique makes many addresses very small. For example, FF01:0:0:0:0:0:0:1 becomes FF01::1, as shown in [Figure 6-5](#). The “unspecified” address is written as a double colon (::), because it contains only zeros. The loopback address, 0:0:0:0:0:0:0:1, is represented by ::1.

```
Successive Zeros: FF01:0:0:0:0:0:0:1 → FF01::1
Loopback Address: 0:0:0:0:0:0:0:1 → ::1
Unspecified Address: 0:0:0:0:0:0:0:0 → ::

Example:
2031:0000:130F:0000:0000:09C0:876A:130B
- Can be represented as 2031:0:130f::9c0:876a:130b.
- Cannot be represented as 2031::130f::9c0:876a:130b.
```

Figure 6-5. Examples of IPv6 Address Representation in Their Long Form and Shortened Form

[Figure 6-5](#) shows the use of the double colon to represent multiple contiguous 16-bit chunks of zeros in an IPv6 address. The second representation that is shown in [Figure 6-5](#) is incorrect—the double colon (::) notation can appear only once in an address, because multiple uses may make the address ambiguous. In the example, the parser cannot tell whether the missing bits (four 16-bit sections) are apportioned with 16 at the first double colon and 48 at the last double colon or some other combination.

IPv6 addresses are presented with a prefix. The high-order bits of an IPv6 address represent the network. All the host addresses of one network would have the same first few bits. These first few network bits, n bits, are called the “prefix.” We use / n to denote a prefix n bits long. If an address is presented with 2001:cb7::/32, it means that the first 32 bits belong to the network and that the other 96 bits belong to subnet and host addresses.

IPv6 Address Types

IPv6 supports three types of addresses:

- Unicast
 - Address is for a single interface
 - IPv6 has several types (for example, global, reserved, link-local, and site-local)
- Multicast
 - One-to-many
 - Enables more efficient use of the network
 - Uses a larger address range
- Anycast
 - One-to-nearest (allocated from unicast address space)
 - Multiple devices share the same address

- All anycast nodes should provide uniform service
- Source devices send packets to anycast address
- Routers decide on closest device to reach that destination
- Suitable for load balancing and content delivery services

Unicast and multicast work the same as with IPv4. As for anycast, think of it as *a shared secondary address*. As an example, think of a network topology with two routers as potential default gateways deserving the same subnet. In IPv4, a host would be configured to point to one of the two gateways. If the selected gateway is no longer available, a reconfiguration of the host would be necessary for it to start using the other gateway on the network. In IPv6, however, the two routers would be configured with the same anycast address, in addition to their own unicast address. A host would send its packets to the anycast address, and either router could provide the service.

Because there are no broadcasts with IPv6, there is no need for a broadcast address.



Key
Topic

Each address type has specific rules regarding its construction and use, as discussed next.

IPv6 Unicast Addressing

IPv6 unicast addresses can be aggregated with prefixes of arbitrary bit length, similar to IPv4 addresses under classless interdomain routing (CIDR).

There are several types of unicast addresses in IPv6, including global addresses, site-local addresses (deprecated), unique local addresses, and link-local addresses. There are also some special-purpose subtypes of global unicast, such as the unspecified address, loopback address, and IPv6 addresses with embedded IPv4 addresses. Additional address types or subtypes could be defined in the future.

IPv6 address types have the following patterns:

- **Global:** Starts with 2000::- **Reserved:** Used by the IETF
- **Private:** Link local (starts with FE80::- **Loopback:** (::1)
- **Unspecified:** (::)

A single interface may be assigned multiple IPv6 addresses of any type: unicast, anycast, or multicast. IPv6 addressing rules are covered by multiple RFCs, including RFC 4291.

IPv6 Global Unicast and Anycast Addresses

Global unicast addresses correspond to the principal use of IPv6 addresses for generic global IPv6 traffic and consume the most important part of the address space.

The structure of a global unicast address, shown in [Figure 6-6](#), is as follows:

- A global routing prefix, typically a /48 assigned to a site, is a structure that enables aggregation upward, eventually to the ISP
- A subnet ID used to identify links within a site, typically 16 bits
- A 64-bit interface ID to identify the interface of the node




Figure 6-6. IPv6-in-IPv4 Tunnel

IPv6 has the same address format for global unicast and for anycast addresses. Every IPv6-enabled interface contains at least one loopback address (::1/128) and one link-local address. Optionally, every interface can have multiple unique local and global addresses.

Examples of global addresses can be found in RFC 3587, “IPv6 Global Unicast Address Format.” The structure that is proposed in that document provides for aggregation of routing prefixes to limit the number of entries in the global routing table. An example of aggregation is shown in [Figure 6-7](#).

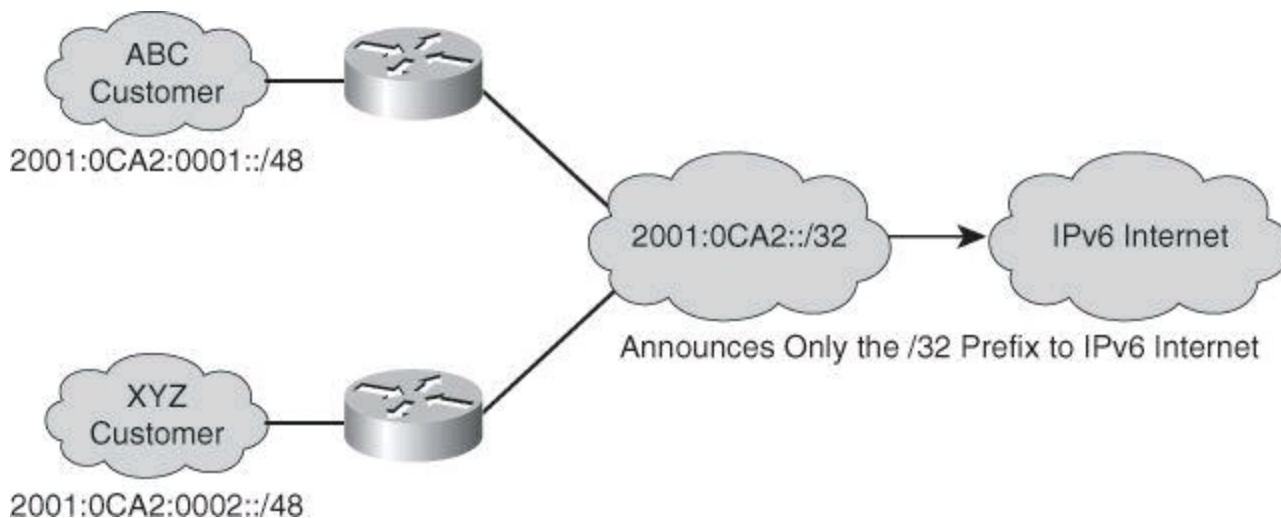


Figure 6-7. Example of IPv6 Address Aggregation

Link-Local Addresses

All IPv6-enabled interfaces must have a link-local address.

Link-local addresses are used for addressing on a single link, so they have a scope that is limited to the link. Link-local addresses are dynamically created on all IPv6 interfaces by using a specific link-local prefix, FE80::/10, and a 64-bit interface identifier, as shown in [Figure 6-8](#).




Figure 6-8. Link-Local Address

Link-local addresses are used for automatic address configuration, neighbor discovery, and router discovery. They are also used by many routing protocols.

Link-local addresses can serve as a way to connect devices on the same local network without requiring global or unique local addresses. When communicating with a link-local address, you must specify the outgoing interface because every interface is connected to FE80::/10.

IPv6 has a 128-bit address space, but 64 bits are used for the host number on the subnet. A better way to look at the address space is to say that IPv6 supports 2^{64} subnets, and each subnet can have a practically unlimited number of hosts. In any case, there are more than enough networks and hosts for the future.

Multicast Addresses

A multicast scope is new in IPv6. Multicast is used in the context of one-to-many. A multicast address identifies a group of interfaces. Traffic that is sent to a multicast address is sent to multiple destinations at the same time. An interface may belong to any number of multicast groups. Multicast is used in the core of many functions in IPv6.

Multicast addresses are defined by the prefix FF00::/8, as shown in [Figure 6-9](#).

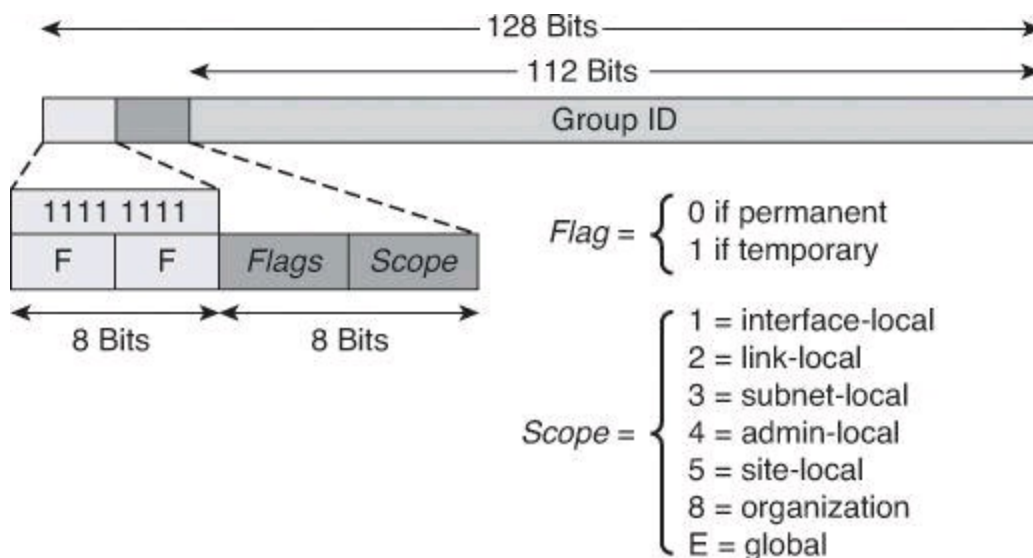


Figure 6-9. Multicast Address

The second octet defines the flags and the scope of the multicast address. Flags are defined as “ORPT,” and these conditions apply:

- 0 is reserved and must equal 0.
- R stands for “rendezvous point” and is almost always set to 0.
- P indicates “prefix dependency” and is almost always set to 0.
- T is the “temporary” bit. For a temporary multicast address, T equals 1. For a permanent multicast address, T equals 0.

Note

If R equals 1, P and T must also equal 1.

The scope parameters are used to specify in which part of the network address is valid and unique. Some addresses are unique only on the local network, while others are globally unique. [Table 6-3](#) provides definitions of scope.

Table 6-3. IPv6 Scope Parameters and Definitions

Scope	Definition
1	Interface-local scope (loopback transmission)
2	Link-local scope (similar to unicast link-local scope)
3	Subnet-local scope, where subnets may span multiple links
4	Administrative-local scope (administratively configured)
5	Site-local scope
8	Organizational scope (multiple sites)
E	Global scope

For example, a multicast address starting with FF02::/16 is a permanent multicast address with a link-local scope.

The lower 112 bits of the multicast address constitute the multicast group ID.

Multicast is frequently used in IPv6 and replaces broadcast. There is no broadcast in IPv6. There is no Time to Live (TTL) in IPv6 multicast. The scoping is defined inside the address.

The multicast addresses FF00:: to FF0F:: are reserved. Inside that range, the following addresses are assigned:

- **FF02::1**: All nodes on the link-local scope
- **FF02::2**: All routers on the link-local scope
- **FF02::9**: All Routing Information Protocol (RIP) routers on the link
- **FF02::1**: FFX:XXXX: solicited-node
- **FF05::101**: All Network Time Protocol (NTP) servers on the site-local scope
- **FF0X::103**: Rwhod (rwho plus ruptime daemon)

- **FF0X::102**: Silicon Graphics Dogfight (Internet game)
- **FF0X::127**: Cisco RP announce (multicast rendezvous point)
- **FF0X::128**: Cisco RP discovery
- **FF05::1:3**: All DHCP servers in site

Assigning IPv6 Global Unicast Addresses

Interface identifiers in IPv6 addresses are used to identify interfaces on a link. They can also be thought of as the “host portion” of an IPv6 address. Interface identifiers are required to be unique on a specific link. Interface identifiers are always 64 bits and can be dynamically derived from a Layer 2 media and encapsulation.

There are several ways to assign an IPv6 address to a device:

- Static assignment using a manual interface ID
- Static assignment using an EUI-64 interface ID
- Stateless autoconfiguration
- DHCP for IPv6 (DHCPv6)

Manual Interface Assignment

One way to statically assign an IPv6 address to a device is to manually assign both the prefix (network) and interface ID (host) portion of the IPv6 address. To configure an IPv6 address on a Cisco router interface and enable IPv6 processing on that interface, use the **ipv6 address *ipv6-address prefix-length*** command in interface configuration mode.

[Example 6-1](#) shows how to enable IPv6 processing on the interface and configure an address based on the directly specified bits.

Example 6-1. Manually Assigning an IPv6 Address to a Router Interface

[Click here to view code image](#)

```
R1(config)# interface fa 0/0
R1(config-if)# ipv6 address 2001:0DB8:2222:7272::72/64
```

Note

In [Example 6-1](#), the IPv6 address could have been configured without the leading 0 in the second most significant hexadecimal field, as shown here:

```
R1(config-if)# ipv6 address 2001:DB8:2222:7272::72/64
```

EUI-64 Interface ID Assignment

Another way to statically assign an IPv6 address is to configure the prefix (network) portion of the IPv6 address and derive the interface ID (host) portion from the Layer 2 MAC address of the device, which is known as the EUI-64 interface ID. You will see later in this chapter how the 48-bit MAC

address is expanded to provide a 64-bit interface ID.

To configure an IPv6 address for an interface and enable IPv6 processing on the interface using an EUI-64 interface ID in the low-order 64 bits of the address (host), use the **ipv6 address *ipv6-prefix/prefix-length eui-64*** command in interface configuration mode.

[Example 6-2](#) assigns the IPv6 address 2001:0DB8:0:1::/64 to Ethernet interface 0 and uses an EUI-64 interface ID in the low-order 64 bits of the address.

Example 6-2. Configuring the Prefix Portion for an EUI-64 Interface ID Assignment

[Click here to view code image](#)

```
R1(config)# interface fa 0/0
R1(config-if)# ipv6 address 2001:0DB8:0:1::/64 eui-64
```

Stateless Autoconfiguration

Autoconfiguration, as the name implies, is a mechanism that automatically configures the IPv6 address of a node. In IPv6, it is assumed that non-PC devices, as well as computer terminals, will be connected to the network. The autoconfiguration mechanism was introduced to enable plug-and-play networking of these devices, to help reduce administration overhead.

DHCPv6 (Stateful)

DHCP for IPv6 enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. It offers the capability of automatic allocation of reusable network addresses and additional configuration flexibility. This protocol is a stateful counterpart to IPv6 stateless address autoconfiguration, and can be used separately or concurrently with IPv6 stateless address autoconfiguration to obtain configuration parameters.

IPv6 EUI-64 Interface Identifier

The 64-bit interface identifier in an IPv6 address identifies a unique interface on a link. A link is a network medium over which network nodes communicate using the data link layer. The interface identifier can also be unique over a broader scope. In many cases, an interface identifier is the same as, or is based on, the data link layer (MAC) address of an interface. As in IPv4, a subnet prefix in IPv6 is associated with one link.

Interface identifiers in global unicast and other IPv6 address types must be 64 bits long and can be constructed in the 64-bit EUI-64 format. As shown in [Figure 6-10](#), this format expands the 48-bit MAC address to 64 bits by inserting “FFFE” into the middle 16 bits. The EUI-64 format interface ID is derived from the 48-bit link layer (MAC) address by inserting the hexadecimal number FFFE between the upper 3 bytes (Organizationally Unique Identifier [OUI] field) and the lower 3 bytes (serial number) of the link layer address. To ensure that the chosen address is from a unique Ethernet MAC address, the seventh bit in the high-order byte is inverted (equivalent to the IEEE G/L bit) to indicate the uniqueness of the 48-bit address. To make sure that the chosen address is from a unique Ethernet MAC address, the U/L bit is set to 0 for global scope (1 for local scope).

U/L Bit: MAC Versus EUI-64

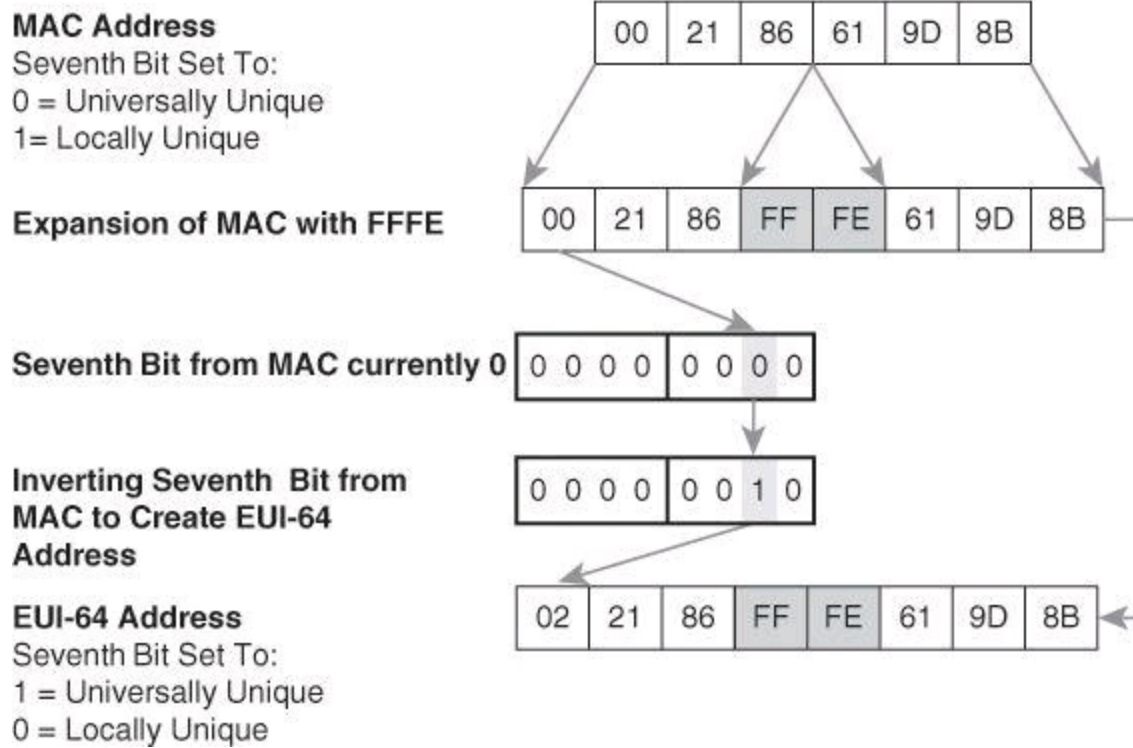


Figure 6-10. Creating an EUI-64 Format Interface ID

IPv6 and Cisco Routers

Cisco IOS Software Release 12.2(2)T and later are IPv6-ready. As soon as you configure basic IPv4 and IPv6 on the interface, the interface is dual-stacked and the router forwards IPv4 and IPv6 traffic on that interface.

There are two basic steps to activate IPv6 on a router. First, you must activate IPv6 traffic forwarding on the router, and then you must configure each interface that requires IPv6. By default, IPv6 traffic forwarding is disabled on a Cisco router.

To activate IPv6 traffic forwarding between interfaces, you must configure the global command **ipv6 unicast-routing**:

Key
Topic

```
R1(config)# ipv6 unicast-routing
```

This command enables the forwarding of IPv6 datagrams.

The **ipv6 address** command can configure a global IPv6 address:

```
R1(config-if)# ipv6 address ipv6prefix/prefix-length eui-64
```

Key
Topic

The link-local address is automatically configured when an address is assigned to the interface. You must specify the entire 128-bit IPv6 address or specify to use the 64-bit prefix by using the EUI-

64 option, as shown in [Example 6-2](#).

IPv6 Address Configuration Example

You can completely specify the IPv6 address or compute the host identifier (rightmost 64 bits) from the EUI-64 identifier of the interface. In [Example 6-3](#), the IPv6 address of the interface is configured using the EUI-64 format, based on the topology shown in [Figure 6-11](#).




Figure 6-11. Topology of IPv6 Networks for [Example 6-3](#)

Example 6-3. Configuring and Verifying a Router for EUI-64 Interface ID Assignment

[Click here to view code image](#)

```
R1(config)# ipv6 unicast-routing
R1(config)# interface fa0/0
R1(config-if)# ipv6 address 2001:db8:c18:1::/64 eui-64
R1# show ipv6 interface fa0/0
FastEthernet0/0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::218:B9FF:FE21:9278
Global unicast address(es):
2001:DB8:c18:1:218:B9FF:FE21:9278, subnet is 2000:1:2:3::/64
Joined group address(es):
FF02::1:FF21:9278
FF02::1
FF02::2
MTU is 1500 bytes
<output omitted>
```

Alternatively, you can completely specify the entire IPv6 address to assign a router interface an address using the `ipv6 address ipv6-address/prefix-length` command in interface configuration mode.

Routing Considerations for IPv6

IPv6 uses longest-prefix match routing just like CIDR does for IPv4. CIDR will be covered in [Chapter 8](#), “[Access Control Lists for Threat Mitigation](#).”

Many of the common routing protocols have been modified to manage longer IPv6 addresses and different header structures. The following updated routing protocols are currently available:

- Static
- RIPng (RFC 2080)
- OSPFv3 (RFC 2740)
- IS-IS for IPv6
- MP-BGP4 (RFC 2545/2858)

- EIGRP for IPv6

This book discusses static routing and RIPng. IPv6 routing is discussed in greater detail in *Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide: Foundation Learning for the ROUTE 642-902 Exam* by Diane Teare.

You can use and configure IPv6 static routing in the same way you would with IPv4. There is an IPv6-specific requirement per RFC 2461 that a router must be able to determine the link-local address of each of its neighboring routers to ensure that the target address of a redirect message identifies the neighbor router by its link-local address. This requirement means that using a global unicast address as a next-hop address with IPv6 routing is not recommended.

Routing protocols use the link-local address as a source for exchanging routing updates, and this link-local address is actually used as next-hop address.



The Cisco IOS global command to enable IPv6 is **ipv6 unicast-routing**. You must enable IPv6 unicast routing before an IPv6-capable routing protocol or an IPv6 static route will work.

Routing Information Protocol next generation (RIPng) (RFC 2080) is a distance-vector routing protocol with a limit of 15 hops that uses split horizon and poison reverse to prevent routing loops. RIPng includes the following features:

- Based on IPv4 RIP version 2 (RIPv2) and is similar to RIPv2
- Uses IPv6 for transport
- Includes the IPv6 prefix and next-hop IPv6 address
- Uses the multicast group FF02::9, the all-RIP-routers multicast group, as the destination address for RIP updates
- Sends updates on UDP port 521
- Supported by Cisco IOS Release 12.2(2)T and later

Revisiting Threats: Considerations for IPv6

The good news, if this can be considered good, is that when considering the TCP/IP protocol stack, the Internet layer (Open Systems Interconnection [OSI] network layer) is the only difference between IPv4 and IPv6. Therefore, when the transition to IPv6 occurs, the layers above and below IPv6 will remain the same. If your web application is vulnerable to attacks in an IPv4 environment, it will also be vulnerable to attacks when IPv6 is used. This means that the threat and vulnerability landscape is similar between the two protocol stacks.

IPv4 and IPv6 are both datagram protocols, and there are many similarities between the two headers. Both headers still have a version, a quality of service (QoS) marking field, a payload length field, a counter to detect how far the packet has traveled, the value of the next upper-layer protocol, and of course a pair of addresses. Therefore, in general, many types of attacks are similar between IPv4 and IPv6, as listed below. For some attack types, additional information is provided.

- Reconnaissance

- Not so easy in IPv6 due to large address space
- Scanners will make router trigger NDP, wasting CPU and resources
- Attack tools exist today (Parasit6, Fakerouter6, Scapy6, others)
- Viruses and worms
 - Scanning will probably use alternative techniques
- Application layer attacks
 - Same implications
 - Peer-to-peer nature of IPv6 augments the problem
- Unauthorized access
- Man-in-the-middle attacks
 - Still a possibility
 - Myth: mandatory IPsec resolves the issue
 - Reality: IPsec is a mandatory part of the stack, but you still have to configure it
- Sniffing or eavesdropping
- Denial of service (DoS) attacks
- Spoofed packets: forged addresses and other fields
 - Still a possibility
 - Bogons (bogus IP addresses) a reality today
- Attacks against routers and other networking devices
- Attacks against the physical or data link layers

However, there is also some bad news. IPv6 is a bit different and, as such, there are threats that have been slightly changed by the fact that IPv6 does things slightly differently than IPv4. The following is a list of threats that are only slightly modified by IPv6:

- LAN-based attacks (NDP)
- Attacks against DHCP or DHCPv6
- DoS against routers (hop-by-hop extension headers rather than router alerts)
- Fragmentation (IPv4 routers performing fragmentation versus IPv6 hosts using a fragment extension header)
- Packet amplification attacks (IPv4 uses broadcast; IPv6 uses multicast)

As far as the protocol is concerned, IPv6 is no more or less secure than IPv4, but the IPv6 protocol is unique and has its own security considerations. The fields within the IPv6 header that are unique to IPv6 include the flow label and extension headers.

Even though IPv6 does not significantly transform the IP header, there will be attacks unique to IPv6.

Following is a list of threats that are unique to IPv6 networks:

- **Reconnaissance and scanning worms:** Brute-force discovery is more difficult.
- **Attacks against ICMPv6:** ICMPv6 is a required component of IPv6.

- **Extension header (EH) attacks:** EHs need to be accurately parsed.
- **Autoconfiguration:** NDP attacks are simple to perform.
- **Attacks on transition mechanisms:** Migration techniques are required by IPv6.
- **Mobile IPv6 attacks:** Devices that roam are susceptible to multiple vulnerabilities.
- **IPv6 protocol stack attacks:** Because of the code freshness of IPv6, bugs in the protocol stack exist.

IPv6 introduces the following difficulties or vulnerabilities:

- Training and planning
 - Lack of knowledge, poor planning even for basic security controls (example: weak ingress filtering, or no filtering at all)
- End nodes are exposed to many threats:
 - **Address configuration parameters:** Rogue configuration parameters
 - **Address initialization:** Denial of address insertion
 - **Address resolution:** Address stealing
 - **Default gateway discovery:** Rogue routers
 - **Neighbor reachability tracking:** Rogue neighbor status
- Header extensions
 - Hosts process routing headers (RH)
 - Header extensions can be exploited (example: routing header for source routing and reconnaissance)
 - Amplification attacks based on routing header

The reliance of IPv6 on multicasting and ICMPv6 makes those protocols subject to multiple exploits to implement DoS, man-in-the-middle, and spoofing attacks, such as:

- Multicasting facilitates reconnaissance (example: FF02::1 is all hosts, FF02::2 is all routers).
- ICMPv6 is a vehicle for autoconfiguration, subject to spoofing and multiple exploits.

The built-in tunneling capabilities of IPv6 become a vulnerability when not properly controlled. Tunnels are inherently covert channels that will very likely be passed through by firewalls that do not support IPv6 or do not have a strong IPv6 filtering mechanism. Something similar happens with other network security controls such as intrusion prevention systems (IPS). Tunneling is pervasive and sometimes automatic in IPv6, so it does not take much to initiate an unwanted tunnel that could be invisible to your security controls. As an example, Teredo runs over IPv4 UDP port 41 and could also run over any UDP port. Also, most IPv4/IPv6 transition mechanisms have no authentication built in.

Dual stacking represents another issue. You might think IPv6 is not present in your network, but most operating systems support it and enable it by default. A fully protected IPv4 host can be exploited if IPv6 is a weak link. Specific dual-stacking vulnerabilities include the following:

- IPv6 is on by default in most operating systems. You may have it on and not know that it's

running.

- Applications can be subject to attack on both IPv6 and IPv4.
- The network is only as secure as the least secure stack.

Examples of Possible IPv6 Attacks

In [Figure 6-12](#), the attacker manipulates the routing header to create a traffic loop. DoS attacks can be performed using this feedback loop to consume resources or amplify the packets that are sent to a victim. RH0 packets could be created with a list of embedded IPv6 addresses. The packet would be forwarded to every system in the list before finally being sent to the destination address. If the embedded IPv6 addresses in an RH0 packet were two systems on the Internet listed numerous times, it could cause a type of feedback loop.




Figure 6-12. Traffic Loop from Exploiting Routing Header

In [Figure 6-13](#), the attacker abuses NDP by using a router to amplify a network scan. The router sends Neighbor Solicitation (NS) messages to all the hosts in the LAN segment, using the all-nodes multicast address.

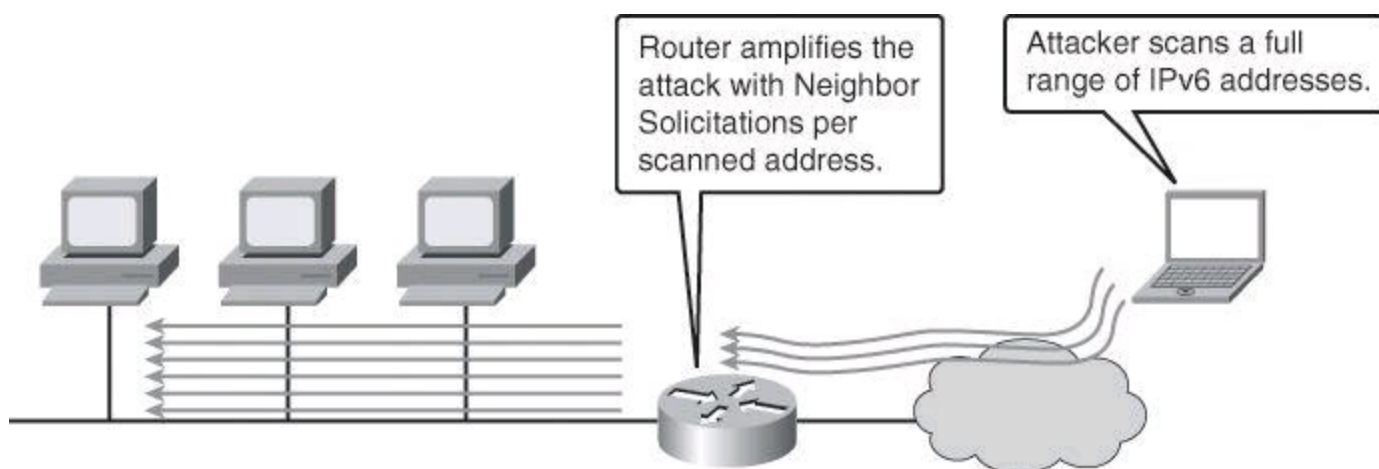


Figure 6-13. Network Scan from Exploiting NDP

By combining multiple techniques, attackers can accomplish stealth attacks that result in trust exploitation and information theft. [Figure 6-14](#) illustrates an attack that combines dual-stacked hosts, which are subject to rogue router advertisements. This type of attack could exploit the routing header (RH) to pivot using multiple hops; and by using automatic tunnels, it could stealthily go through firewalls and IPS sensors.

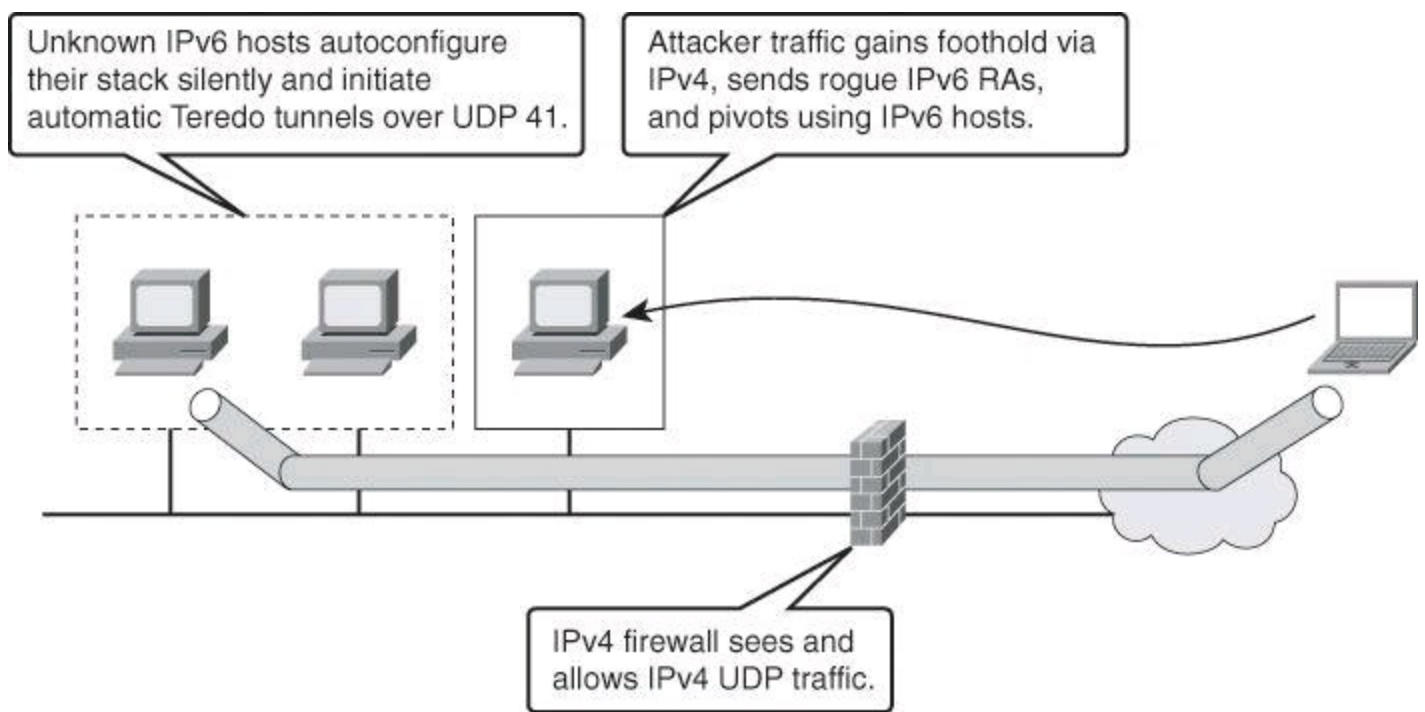


Figure 6-14. Combo Attack on IPv6

The attacker gains a foothold in the IPv4 network. The compromised host sends rogue router advertisements, triggering unwilling dual-stacked hosts to obtain an IPv6 address. These compromised hosts could trigger an automatic Teredo tunnel, which would go unnoticed by the firewall. The attacker can also use the routing header to pivot around multiple hosts in the internal network before sending traffic out.

Recommended Practices

The same best current practices for protecting IPv4 networks are still appropriate for IPv6. Standard perimeter security architecture still applies to IPv6 as it does to IPv4. Filtering at the edge and trying to protect the interior are still the order of the day. The network architecture model of core, distribution, and access will still be the way that IPv6 networks are designed. Many of the same protection mechanisms that are used in IPv4 networks will be adapted to work on IPv6. The same operational guidelines and forensic search also apply to IPv6.

The following list summarizes additional guidelines for IPv6 security. Some of these guidelines will be discussed in more detail in other chapters of this book, such as IPv6 access lists in [Chapter 8](#) and IPv6 aware intrusion prevention in [Chapter 11](#), "[Intrusion Prevention Systems](#)."

- Ingress filtering is key:
 - Deny Bogon addresses.
 - Filter multicast packets at your perimeter based on their scope.
 - Permit only packets that have as a destination address your allocated block of addresses or multicast group address or your link-local address for NDP.
 - Granularly filter ICMPv6 messages at the perimeter (remember, ICMPv6 is needed for protocol operations such as NDP).
 - Drop RH0 packets and unknown extension headers at the perimeter and throughout the interior of the network.

- Favor dual stack as the transition mechanism, but secure each protocol equally.
- Control the use of tunneling:
 - Configure manual tunnels if possible.
 - Do not allow tunnels through the perimeter unless required.
- Consider current and future security enhancements:
 - Secure NDS (SeND) from RFC 3971 provides a cryptographic method to Neighbor Discovery.
 - RA Guard, from RFC 6105, is an alternative and complement to SeND, filtering at Layer 2.

Summary

In this chapter, you learned about the need for IPv6. You saw that IPv6 offers more benefits than IPv4, including a larger address space, easier address aggregation, and integrated security.

A review of the IPv6 addressing scheme revealed that IPv6 addresses are 128 bits long, with a 48-bit global prefix, a 16-bit subnet ID, and a 64-bit interface identifier. You saw how EUI-64 addresses are derived from the MAC address. You also saw that network addresses could be assigned statically, through a stateless configuration, or through DHCPv6.

You learned that IPv6 is not automatically enabled on Cisco IOS 12.2(2)T and later; to turn it on, you have to use the **ipv6 unicast-routing** command, and to assign an IPv6 address to an interface, you have to use the **ipv6 address** interface command. You also saw that Cisco supports all major IPv6 routing protocols: RIPng, OSPFv3, and EIGRP.

The chapter also discussed the transition from IPv4 to IPv6 using dual stacks, tunneling, and possibly NAT-PT. After the overview of IPv6's general characteristics, you learned how common and specific threats could affect IPv6, and how to develop and implement a strategy for IPv6 security.

References

For additional information, refer to these resources:

Cisco Systems, Inc. *Cisco IOS IPv6 Configuration Guide, Release 12.4, Implementing IPv6 Addressing and Basic Connectivity*, http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-addrg_bsc_con.html

Cisco Systems, Inc. *IPv6 and IPv4 Threat Comparison and Best-Practice Evaluation (v1.0)*, http://www.cisco.com/web/about/security/security_services/ciag/documents/v6-v4-threats.pdf

RFC 2464, "Transmission of IPv6 Packets over Ethernet Networks," <http://www.ietf.org/rfc/rfc2464.txt>

RFC 3146, "Transmission of IPv6 Packets over IEEE 1394 Networks," <http://www.ietf.org/rfc/rfc3146.txt>

RFC 3587, "IPv6 Global Unicast Address Format," <http://www.ietf.org/rfc/rfc3587.txt>

RFC 4007, "IPv6 Scoped Address Architecture," <http://www.ietf.org/rfc/rfc4007.txt>

RFC 4291, "IP Version 6 Addressing Architecture," <http://tools.ietf.org/html/rfc4291>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions.](#)”

1. Which global command enables IPv6 on a Cisco router?
 - a. **ipv6 transition**
 - b. **ipv6 routing**
 - c. **ipv6 unicast-routing**
 - d. **ipv6 anycast**
2. Which are advantages of IPv6 over IPv4? (Choose two.)
 - a. Larger address space
 - b. Complex header
 - c. Security
 - d. Efficient broadcast
3. Which of the following functions is unique to ICMPv6, as compared to ICMPv4?
 - a. Router discovery
 - b. Connectivity checks
 - c. Informational/error messaging
 - d. Fragmentation requiring notification
4. Which is not a valid IPv6 address?
 - a. ::
 - b. FF01::1
 - c. 2031::130f::9c0:876a:130b
 - d. 2031::130f:9c0:876a:130b
5. Which IPv6 address types can be acquired automatically by an IPv6 host?
 - a. Global unicast address
 - b. Anycast address
 - c. Multicast address
 - d. Link-local address
 - e. All of the above
6. Which IPv6 feature or component, when unprotected, is more likely to result in covert channels that go undetected by firewalls?
 - a. Routing header
 - b. 6to4 transition deployments
 - c. ICMPv6
 - d. Autoconfiguration
7. Which IPv4-to-IPv6 transition technology is implemented as a translation mechanism?

- a. Dynamic 6to4 tunnels
- b. Dual stacking
- c. ISATAP
- d. NAT-PT

8. Choose two threats that result directly from the use of the routing header alone in IPv6?

- a. Denial of service
- b. Confidentiality exploits
- c. Advanced reconnaissance
- d. Worm amplification
- e. Spoofing

9. Which of the following is a valid EUI-64 address for MAC 0021.8661.9D8B?

- a. 002186FFFF619D8B
- b. 002186FFFE619D8B
- c. 012186FFFE619D8B
- d. 022186FFFE619D8B

10. Which of the following is an IPv6 routing protocol?

- a. RIPv6
- b. OSPFv6
- c. EIGRP for IPv6
- d. MP-BGP6

Part III: Threat Control and Containment

Chapter 7. Planning a Threat Control Strategy

Current trends in security threat vectors require a carefully planned threat control strategy. Trends affecting security include persistent application layer exploits, using social engineering to exploit the trust architecture of the enterprise, the pervasiveness of mobility and consumerization, and the insidious motivations behind the behavior of the attackers. All of these trends result in the need for dynamic security intelligence gathering and distribution, early warning systems, and application layer inspection for mobile services where data and applications are hosted in the cloud. This chapter suggests design principles to plan a threat control and containment strategy using firewalls and intrusion prevention systems in Cisco IOS environments. (Note that in previous chapters you have started to see how to use switches and routers to provide protection; in upcoming chapters, we will further explore the protection provided by routers when used for filtering, firewalling, and encrypting traffic.) In this chapter, we will

- Evaluate the current state of enterprise security in the presence of evolving threats
- Describe design considerations for a threat protection strategy to mitigate threats as part of a risk management strategy
- Describe how Cisco strategizes threat control and containment

Threats Revisited

Previous chapters of this book discussed the evolution of threats in information security. Recent trends indicate that threats have become global in nature, increasingly sophisticated, and pervasive. New exploits are designed to have a global impact in minutes. Blended threats, which use multiple means of propagation, are more sophisticated than ever. The trends are becoming regional and global in nature. Attackers are also learning to be dynamic and reactive, changing their approach rapidly if needed. I remember a discussion I had with the director of network security for the federal reserve bank of his country; he remarked, “We have a small team of professionals working 8 hours per shift at protecting our network, but there are hundreds of hackers coming at us 24 hours per day, nonstop.” This network security team had to make sure that every crease of their network was secure. The hacker had to find only one open, minuscule gap to exploit.

Threats are also becoming persistent. After the initial attack starts, subsequent attacks may come in waves as infected systems join the network. Because infections are so complex and have so many end users (employees, vendors, and contractors), multiple types of endpoints (company desktop, home, and server), and multiple types of access (wired, wireless, VPN, and dial-up), infections are difficult to eradicate.

Trends in Network Security Threats

Evolution in technology and craftiness of hackers bring in new threat vectors. Threat vectors are the methods used to get to the desired target. More recent threat vectors are increasingly sophisticated, and the motivation of the attackers is reflected in their impact. Recent threat vectors include the following:

- Cognitive threats: social networks (likejacking)
- Smartphones, tablets, and consumer electronics exploits

- Widespread website compromises
- Disruption of critical infrastructure
- Virtualization exploits
- Memory scraping
- Hardware hacking
- IPv6-based attacks

The following is a list of the specific trends that can be gathered from the evolution of threats in information security:

- **Insidious motivation, high impact:** The primary motivation of the attacker is no longer just fame and notoriety. It is mostly political and financial, aimed at economic espionage and money-making activities. Malicious attackers use a businesslike, structured process not only to write code, but also to “conduct business.” It is common to see a lifecycle approach to software development and a businesslike environment where developers cooperate and make attack tools publicly available.
- **Targeted, mutating, stealth threats:** The modus operandi is increasingly stealth, using mutating malware that changes by the minute. In some instances, hundreds of variants of the same malware are created in a matter of hours.

Cybercriminals see value in fine-tuning their efforts so that their malware reaches a single high-profile target or performs a specific function. An example of a high-profile target attack effort would include attacking critical infrastructure, government operations, and organizations with massive amounts of customer identity information. Attackers use strength in numbers. Botnets with specific targets are the norm, whereas large-scale worms with no individual objective are less common or effective. In the past, botnets would use centralized command and control. Botnets today use peer-to-peer control, so there is no central controller to take down.

- **Threats consistently focusing on the application layer:** Known application layer attacks such as cross-site scripting and SQL injection are still relevant, as well as other vulnerabilities such as iFrame injection, which result from programming errors in web services environments. Increasingly, attacks have demonstrated that criminals looking for financial gain can exploit vulnerabilities resulting from web programming errors as new ways of penetrating important organizations. Attacks exploiting web browser vulnerabilities are common.
- **Social engineering front and center:** Social engineering continues to attract victims who are willing to share information with people they believe are known to them. Nowhere is this trend more visible than in social networks, a prime target of attacks. One noticeable shift in social engineering is that criminals are spending more time figuring out how to assume someone’s identity, perhaps by generating emails from an individual’s computer or social networking account.
- **Threats exploiting the borderless network:** Smartphones are general-purpose computers, so worms, viruses, and other malware will increasingly target them. Open mobile platforms are a prime target, and as more workers use their consumer devices for

business-critical activities, endpoint security becomes a critical security control.

Data is now stored everywhere, with the advent of cloud computing and cloud services. Cloud security is a major concern for organizations of all kinds. A technology is often the victim of its own success. As a technology becomes more popular, it appears more prominently on hackers' radar screens, thus motivating them to attack this new technology. As an example, it is expected that as IPv6 becomes more pervasive, IPv6 vulnerabilities and threats will become more commonplace.

Threat Mitigation and Containment: Design Fundamentals

The result of the recent trends in information security threats is the need for an updated, carefully planned threat control and mitigation strategy, and a revision of old design paradigms.

- **Policies and process definition:** The nature and frequency of attacks, and the increasingly high stakes in terms of organizational and business risk, call for a more formal approach to defining threat control policies. What probably started as a reactive process, and then moved to a more socialized process, needs to become a formal process in the wake of botnets, industrial espionage, and identity theft.
- **Mitigation technologies:** The mutating and stealth nature of attacks requires more accuracy in detection and prevention technologies. It also requires a more automated process, where networks and applications react to attacks and mitigate them with little or no human intervention. This used to be a more manual process in the past; however, organizations can no longer risk the potential consequences of this manual approach.
- **End-user awareness:** Data is stored everywhere, and users access mission-critical resources from multiple locations at multiple times, using multiple types of devices. The new paradigm must get the end user more involved with security, where in the past this was a function that was fully performed by IT.

Threat Control Design Guidelines

These new paradigms result in specific design guidelines for the threat control and containment architecture:

- **Stick to the basics:** The fundamental design guidelines are still valid. Implement a layered security approach, where defense in depth defines redundant, overlapping security controls to mitigate risk and provide fault tolerance. The threat control design should consider both internal and external networks to be untrusted and apply policy for both inbound and outbound traffic.
- **Risk management:** This should be a central consideration. Logging and monitoring are still critical functions.
- **Distributed security intelligence:** Because data and access are everywhere, there should also be security "sensors" everywhere. Security needs to be embedded in network elements, operating systems, and applications. This network of sensors needs to be integrated. Actionable, relevant, highly dynamic security information should be available to security devices and controls.
- **Security intelligence analysis:** This army of sensors and security controls will produce

vast amounts of information. Event management, incident management, and, in more general terms, security intelligence management require a central approach to analysis and correlation, including correlation with third parties and external sources. This analysis should result in relevant threat information that should be automatically distributed to security devices and controls in a timely manner.

- **Application layer visibility:** Application layer inspection is critical and should be implemented efficiently and effectively.
- **Incident response:** Incident response should be automated and dynamic, providing end users the means to respond and mitigate.

Application Layer Visibility

Application layer security controls provide visibility at higher layers, up to Layer 7 (as shown in [Figure 7-1](#)), necessary to match protocol and application anomalies that are purposefully crafted to exploit application and operating system vulnerabilities. This functionality is necessary in firewalls and intrusion prevention systems, and it should be implemented with effectiveness and efficiency in mind. Application layer controls are often implemented in hardware, and they require a good amount of capacity planning because they trigger deep packet inspection and are more CPU intensive than Layer 4 controls such as access control lists (ACL).

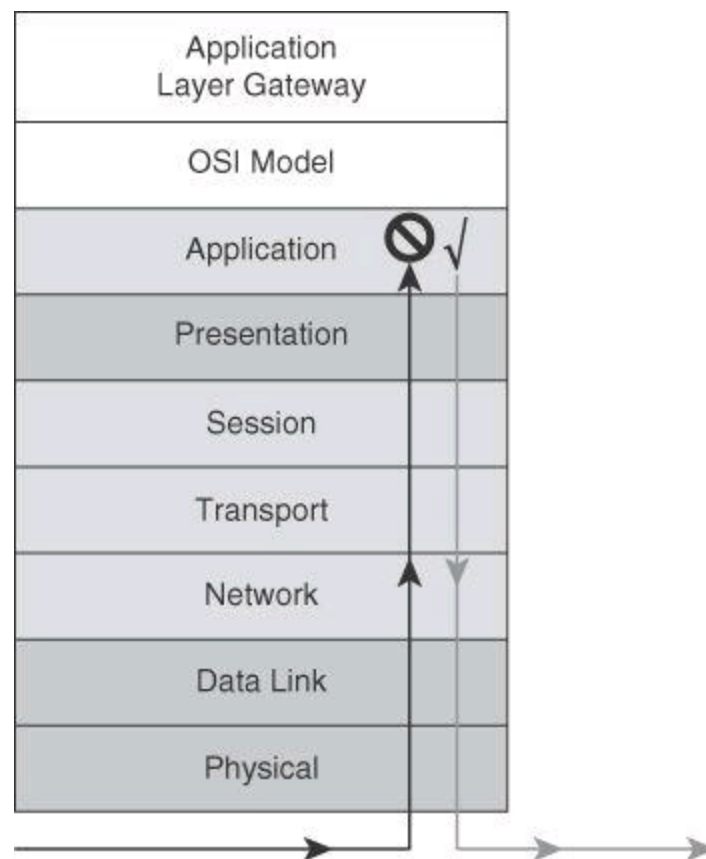


Figure 7-1. Application Layer Visibility Analyses Traffic Up to the Application Layer

Application layer visibility provides the intelligence to distinguish botnet command and control activity from otherwise innocuous web traffic, and can distinguish between valid encrypted data and data meant to obfuscate the application. Without application layer visibility, you cannot comprehensively detect, respond to, and investigate policy violations, exploits, and intrusions. Cisco IronPort Web Security Appliance (WSA) is a device that provides application layer visibility. The

application layer visibility, which is the bread and butter of the IronPort WSA, can also be accomplished with Layer 7 filtering on the Cisco ASA firewalls and on Cisco IOS routers, though it necessitates more effort.

Distributed Security Intelligence

The distributed approach to gathering and providing relevant and actionable security intelligence follows the defense-in-depth guidelines of security architecture design. For threat management, this approach creates an army of sensors and threat controls that work collaboratively and automatically to provide a more effective threat detection and response mechanism. This results in higher threat coverage and greater accuracy, as new threats detected in one location are fed back into the system, making other locations know about the threat and update their detection and response capabilities to manage it.

This system is based on telemetry, collaboration, and proactive protection, where sensors in the field feed a central system that distributes the information to other sensors, as shown in [Figure 7-2](#). This approach acts as an early warning system to prevent new attacks and automatically tune the response to current attacks. Cisco offers such systems with its AnyConnect telemetry module, which communicates with Cisco IronPort WSA. The web filtering infrastructure uses this data to strengthen its web security scanning algorithms, improve the accuracy of the URL categories and web reputation database, and ultimately provide better URL filtering rules for the organization.

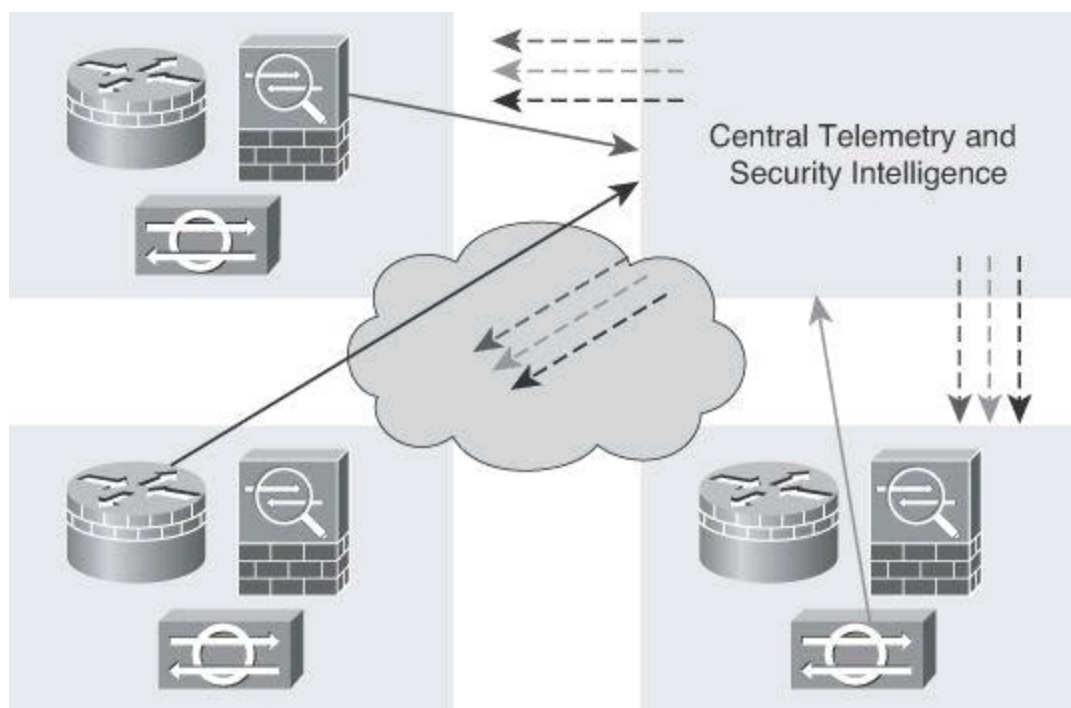


Figure 7-2. Distributed Security Intelligence Using Telemetry

Security Intelligence Analysis

The proliferation of security controls could overflow the ability of the system to gather and analyze the security intelligence information. Building an effective security intelligence analysis capability becomes a crucial component of the threat control architecture design.

Event and incident management are critical components of this strategy. Monitoring and logging are fundamental guidelines that may be already in place. Local and global correlation of the monitoring and log data becomes increasingly important, for accuracy and automation, and also for forensics and

post-mortem analysis.

In [Figure 7-3](#), multiple devices send a log entry to the syslog server, on the right, regarding the single event generated by the hacker attacking the target. All these alerts are analyzed by the central event management and correlation server, and the administrator is notified once about the issue. Without proper correlation analysis, the network administrator might have been notified by the alert from the router, by the alert from the firewall, and by the alert generated by the IPS sensor. Cisco MARS, now in end-of-life status but still supported by Cisco, was an example of such central event management and correlation system. As mentioned in [Chapter 4](#), Cisco now refers its customers to one of its Security Information and Event Management (SIEM) ecosystem partners, such as LogLogic, Splunk, and so forth. More information about SIEM partners can be found at http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns744/ns1090/landing_siem.html.

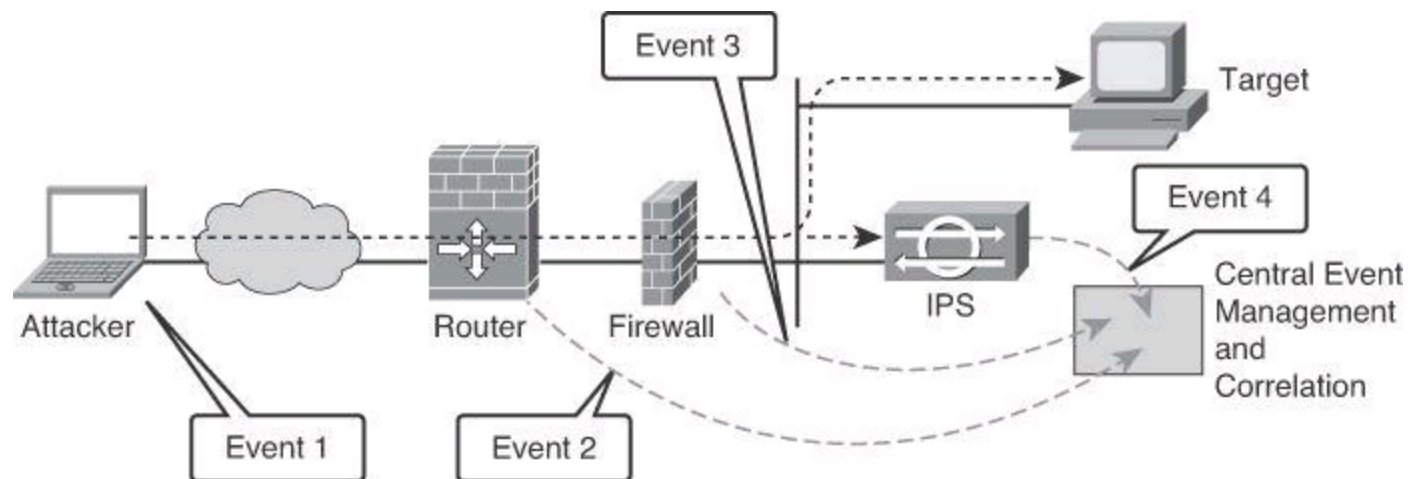


Figure 7-3. Security Information and Event Management

Integrated Threat Control Strategy

The Cisco threat control and containment solution offers comprehensive protection for your network through network-wide visibility, simplified policy control, and proactive system protection. The Cisco threat control and containment solution regulates network access, isolates infected systems, prevents intrusions, and protects critical business assets. This solution uses centralized policy, configuration, and threat event management to counteract malicious traffic, such as worms, viruses, and malware, before it affects your business.

Cisco Threat Control and Containment Categories

Threat prevention solutions have been a flagship of Cisco security for a long time. Integrated security controls found in Cisco Adaptive Security Appliance (ASA), Integrated Services Routers (ISR), and Intrusion Prevention System (IPS) sensors provide a multidimensional approach to detecting, mitigating, and responding to threats at various levels. Threat prevention mechanisms, such as application layer inspection, Modular Policy Framework (MPF), signatures, and anomaly detection and firewalling, are integrated into the fabric of the network. Flexible deployment options include standalone appliances, service modules that are embedded into routers and switches, and virtualized devices, such as Cisco ASA contexts (virtual firewalls). These mechanisms can be deployed on premises, off premises, or as part of a hybrid, managed approach using cloud services.

[Table 7-1](#) provides an overview of the different technologies incorporated into the Cisco threat

control and containment solution. These technologies are looked at more closely in the rest of this chapter.

Table 7-1. Cisco Threat Control and Containment

Category	Method	Recommendations
Integrated approach to threat control	Stick to the basics, formalize	Sensors everywhere Multiple form factors for firewall and IPS (physical and virtual appliance, hardware module, software based)
Application awareness	Distributed security intelligence	Advanced Inspection and Control (AIC) Modular Policy Framework (MPF) Flexible Packet Matching (FPM) Network-Based Application Recognition (NBAR)
Application-specific gateways	Application layer visibility	IronPort email and web security ScanSafe web security
Security management	Security intelligence analysis	Multiple options for event and log management (device managers, Cisco Configuration Professional, Cisco Security Manager)
Cisco SIO	Incident response	Global Threat Operations Centers (TOC), dynamic updates SensorBase and IntelliShield Correlation and data mining

Cisco security technologies facilitate the implementation of the design guidelines previously described by incorporating the following features and technologies:

Integrated Approach to Threat Control

The network is the system that detects and prevents threats. As such, threat control is embedded and integrated in the network using multiple deployment options: appliance-based controls, hardware modules for existing network elements, software-based controls, and virtualized security elements, are all part of this strategy to prevent intrusion and control threats.

Application Awareness

Application awareness is implemented across the board in multiple network elements and security devices, such as the Zone-Based firewall that performs application layer (Layer 7) inspection and deep packet inspection. These solutions are implemented in the form of multiple application

inspection mechanisms, such as, among others, the following:

- Any alphanumeric character
- Modular Policy Framework (MPF)
- Network Based Application Recognition (NBAR)
- Flexible Packet Matching (FPM)

Application-Specific Gateways

Application awareness is also implemented in the form of application-specific gateways, geared toward the inspection and analysis of traffic related to specific applications. IronPort web and email security appliances and ScanSafe web security technologies are examples of those gateways.

Security Management

A comprehensive approach to event management is available in all security management tools, with scalable features available for different scenarios. From device managers such as Cisco IPS Device Manager (IDM), to configuration builders such as Cisco Configuration Professional, to central management applications such as Cisco Security Manager, event management and correlation are a key function in Cisco security management.

Cisco Security Intelligence Operations Site

Cisco SIO is the back-end security ecosystem that detects threat activity, researches and analyzes the threats, and provides real-time updates and best practices to keep organizations informed and protected. Cisco SIO consists of three pillars:

- Threat intelligence, which is called Cisco SensorBase, shown in [Figure 7-4](#)
- The automatic and human development process, called the Threat Operations Center
- The automated and best practices content that is pushed to network elements in the form of dynamic updates

The screenshot shows the Cisco IronPort SenderBase Security Network web page. At the top, there is a navigation bar with links: HOME, THREAT OVERVIEW, TOP SENDERS, REPUTATION LOOK UP, HELP, and ABOUT. The main heading is "Cisco IronPort SenderBase Security Network".

Threat Activity Source: A world map showing threat activity sources across North America, South America, Europe, Africa, and Asia. The map is powered by Google and includes a legend for Email Traffic, Spam, Viruses, and More Details.

Current Threat Outbreaks: A table listing various threats with their names and detection times.

Name	Time
Troj/Agent-WVO	07/04/2012 21:25
Troj/Agent-WXF	07/04/2012 14:26
Troj/Agent-WXB	07/04/2012 10:40
Troj/Agent-WWY	07/04/2012 06:34
Troj/Agent-WWV	07/04/2012 00:42
Troj/Agent-WXG	07/04/2012 00:03
Troj/VB-FZM	07/03/2012 18:46
Troj/VB-FZJ	07/03/2012 16:59
Troj/Mdrop-EIS	07/03/2012 11:21

Today's Global Email Traffic Watch: A table showing IP addresses, volumes, and countries.

IP Address	Volume (m)	Country
195.228.75.125	13.2 ↑	—
62.225.150.103	11.5 ↓	—
195.245.231.130	10.8 ↓	—
255.255.255.255	10.5 ↓	—
200.49.102.79	10.2 ↑	—
195.245.230.34	8.6 ↓	—
95.211.161.171	8.2 ↑	NL
95.211.160.141	8.1 ↑	NL
66.220.155.156	7.5 ↓	US
66.220.155.147	7.5 ↓	US

Web and Email Reputation Look Up: A section with a magnifying glass over the number -6.5, stating "You are only as credible as your online reputation." It includes a search box for IP address, URI, or Domain.

Right Sidebar: Contains a "Look up your network:" search box, a "Look Up" button, and "Reputation Look Up" section with "QUICK LINKS" (Blocked?) and "EXTERNAL LINKS" (Threat Operations Center, Web Reputation, Email Reputation). There are also promotional banners for "FREE SPYWARE AUDIT" and "Cisco Security Center".

Figure 7-4. Cisco IronPort SenderBase Web Page

Cisco SIO is a security intelligence center that baselines the current state of threats on a worldwide basis and provides the network system with valuable information to detect, prevent, and react to threats. Cisco SIO acts as an early warning system by correlating threat information from the SensorBase, analyzed by the Threat Operations Center. Cisco SIO then feeds this information to enforcement elements for live threat prevention that is based on malware outbreaks, current vulnerabilities, and zero-day attacks.

More information on Cisco SIO can be found at <http://tools.cisco.com/security/center/home.x> and at <http://www.senderbase.org>.

Cisco Threat Control and Containment Solutions Fundamentals

Cisco offers multiple solutions in multiple different platforms to provide threat control and containment. For the purpose of the CCNA Security, we will focus on the most common technology used by small and medium businesses, which are Firewalling, including access-control lists, and Intrusion Prevention Systems. For a more exhaustive list of Cisco's offerings, visit www.cisco.com.

Cisco Security Appliances

Cisco security appliances, commonly called firewalls, are integrated devices that offer a new generation of threat control mechanisms. Solutions that include firewalling services are available in multiple footprints to implement the distributed security intelligence approach to threat prevention:

- **Cisco ASA:** These multifunction security devices include firewalls, VPNs, IPSs, and content security technologies. They are available in multiple form factors, and benefit multiple places in the network, from the small office to the data center and service provider environment.
- **Hardware modules:** Designed for data center and large enterprise networks, these security modules integrate into the Catalyst 6500 switch:
 - Cisco Catalyst 6500 ASA Services Module
 - Cisco Catalyst 6500 Firewall Services Module (FWSM)
- **Cisco IOS Firewall:** Integrated in the Cisco IOS software in Integrated Services Routers and Layer 3 switches, this Cisco IOS option implements the zone-based policy firewall and allows a distributed approach to firewall services.
- **Cisco Virtual Security Gateway (VSG):** Designed for data centers where multitenancy and service virtualization are common, this virtual appliance integrates with commercial server virtualization and switch virtualization environments.

Note

In February 2012, Cisco announced the latest firewall OS, called Cisco ASA CX, which will bring application awareness and visibility to the traffic similarly to, but not completely the same as, how a secure web gateway would do it. Customers who wish for the full array of context awareness should still invest in a Cisco IronPort Web Security Appliance.

Cisco also announced new platforms, called the Cisco ASA 5500-X Series capable of running the Cisco ASA CX operating system. One of the many notable characteristics of the ASA 5500-X Series is its capacity to run IPS in software.

For more on Cisco ASA CX and the Cisco ASA 5500-X Series platforms, visit Cisco.com.

The different firewalls listed above implement various access control mechanisms for the new landscape of information security threats that are described in this module:

- Zone-based firewall
- ACLs
- FPM
- AIC
- MPF

- URL filtering
- User-based access control (cut-through proxy)
- Stateful failover

Cisco IPSs

Cisco IPSs are another example of integrated devices offering a new generation of threat control mechanisms. The IPS functionality is available in multiple footprints to implement the distributed security intelligence approach to threat prevention:

- **Cisco IPS 4200 Series Sensors:** These IPS appliances are available in multiple form factors, and benefit multiple places in the network, from the small office to the data center and service provider environment.
- **Hardware modules:** Integrating into ASA appliances, Catalyst 6500 switches, and ISRs, these modules embed the function of IPS sensors into the fabric of the network, and at the same time integrate with global correlation solutions such as the Cisco SIO service:
 - Cisco ASA Advanced Inspection and Prevention Security Services Module (AIP-SSM) and Cisco ASA Advanced Inspection and Prevention Security Services Card (AIP-SSC)
 - Cisco Intrusion Detection Services Module (IDSM-2) for Cisco Catalyst 6500
 - Cisco IPS Network Module Enhanced (IPS NME) and Cisco IPS Advanced Integration Module (IPS AIM) for Cisco Integrated Services Router Generation 2 (ISR G2)
- **Cisco IOS IPS:** Integrated in the Cisco IOS operating system in ISRs and Layer 3 switches, this Cisco IOS option implements IPS technologies consistent with the rest of the Cisco IPS solutions.

These IPSs implement various intrusion management solutions for the new landscape of information security threats that are described in an upcoming chapter:

- Rich set of detection mechanisms
- Signatures
 - Anomaly detection
 - Normalization
 - Correlation
- Automatic signature updates
- Multiple deployment modes
 - Inline
 - Promiscuous

[Figure 7-5](#) illustrates the Cisco threat control and containment architecture for a small network scenario. Notice the distributed security intelligence approach, implemented through various threat control devices, including Cisco IOS zone-based and Cisco ASA firewalls, Cisco IOS IPS, and Cisco Security Manager. Application layer inspection, signature- and anomaly-based protection, and collaboration via the Cisco SIO Center are showcased in this scenario.

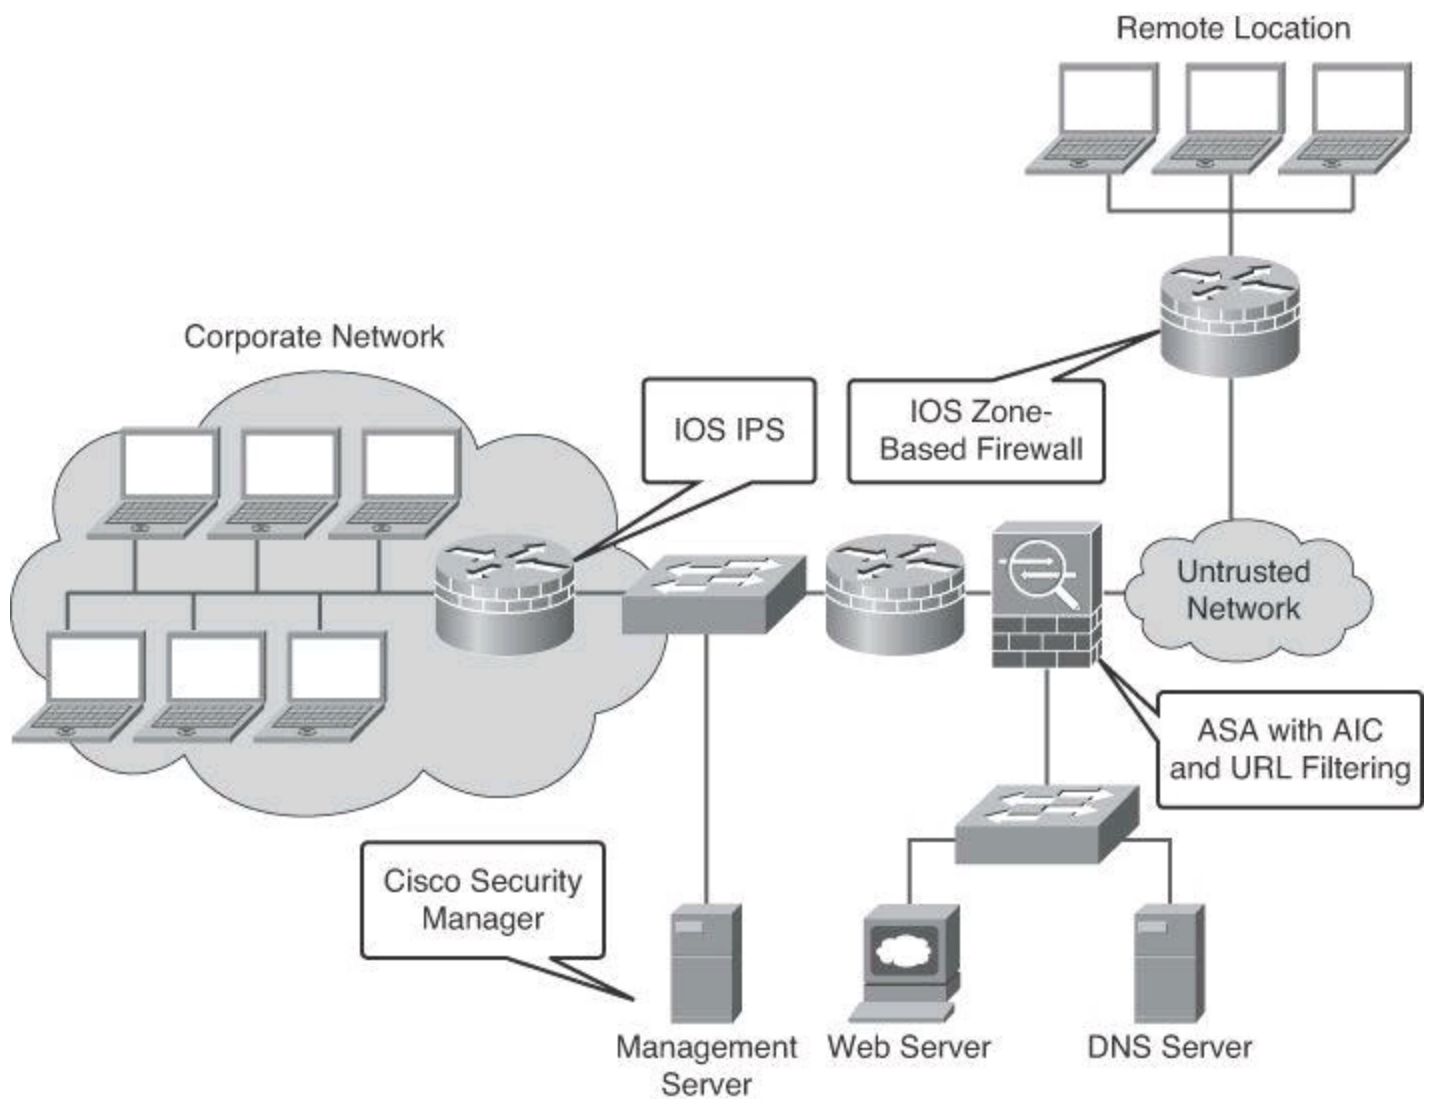


Figure 7-5. Threat Control Scenario for a Small Business

Summary

Threat vectors are stealth, targeted, aimed at the application. The motivations of attackers are increasingly based on social engineering and financial gains.

The following are the main points conveyed in this chapter:

- Threat control and containment should distribute security intelligence, improve incident analysis and correlation, and respond automatically.
- Cisco threat control and containment solutions provide multiple deployment options: appliance, hardware module, software based, and virtualized.
- Cisco threat control and containment is a solution for small, medium, and large businesses.

References

For additional information, refer to these Cisco.com resources:

“Cisco Security Intelligence Operations,” <http://tools.cisco.com/security/center/home.x>

“Cisco 5500 Series Adaptive Security Appliances,” <http://www.cisco.com/en/US/products/ps6120/index.html>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

- 1.** The term “likejacking” refers to which type of threat vector?
 - a.** Phishing
 - b.** IPv6 Teredo tunnels
 - c.** Social network exploits
 - d.** Botnets
- 2.** Which Cisco product or service provides the global correlation service?
 - a.** IPS
 - b.** SIO
 - c.** Third-party SIEM systems
 - d.** SecureX
- 3.** Which two security appliances best fit a data center scenario where server virtualization and multitenancy are among the main architectural components?
 - a.** FWSM
 - b.** ASA
 - c.** Zone-based firewall
 - d.** VSG
- 4.** What is meant by end-user awareness?
 - a.** End users are consulted and made aware of the different applications tested by the organization.
 - b.** End users are aware of the potentially harmful traffic generated by other users.
 - c.** End users are involved in security functions that were strictly reserved to IT in the past, such as the user-acceptable policies and being conscious of potential threats.
 - d.** End-user awareness is a network security training program developed by IT for future dissemination to end users.
- 5.** Which of the following is seen as the best protection against zero-day attacks?
 - a.** NBAR
 - b.** AIC
 - c.** SIO
 - d.** MPF
 - e.** FPM
 - f.** ACL

Chapter 8. Access Control Lists for Threat Mitigation

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

This chapter outlines the types of ACLs that are available and provides guidelines that help create ACLs to provide network security in IPv4 and IPv6 environments. More precisely, this chapter

- Lists the benefits of ACLs
- Describes the building blocks and operational framework of ACLs
- Describes summarizable address blocks in the context of CIDR and VLSM environments, demonstrating how ACL wildcard masks allow for threat mitigation in those environments
- Lists design considerations when deploying ACLs
- Demonstrates the use of Cisco Configuration Professional and the CLI to deploy and verify a threat containment strategy using ACLs
- Demonstrates the use of Cisco Configuration Professional and the CLI to correlate ACL log and alarm information in order to monitor their impact and effectiveness
- Demonstrates how to configure object groups to streamline the implementation of ACLs for threat control
- Demonstrates how to configure ACLs in IPv6 environments, highlighting the operational differences with IPv4 ACLs

ACL Fundamentals

ACLs provide packet filtering for routers and firewalls to protect internal networks from the outside world. ACLs filter network traffic in both directions by controlling whether to forward or block packets at the router interfaces, based on the criteria that you specify within the ACLs. ACL criteria could be the source address of the traffic, the destination address of the traffic, the upper-layer protocol, or other information. Be aware, however, that sophisticated users (hackers) can sometimes successfully evade or fool basic ACLs not only because authentication is not required, but mainly because of the inability of an ACL to track the state of a connection.

ACLs provide a basic level of security for accessing your network. If you do not configure ACLs on your router, all packets passing through the router could get to all parts of your network. You can use ACLs on a router that is positioned between two parts of your network to control traffic entering or exiting a specific part of your internal network. An ACL on the router allows one host to access a part of your network and prevents another host from accessing the same area. The ACL shown in [Figure 8-1](#) allows Host A to access the Human Resources network but prevents Host B from accessing the Human Resources network.

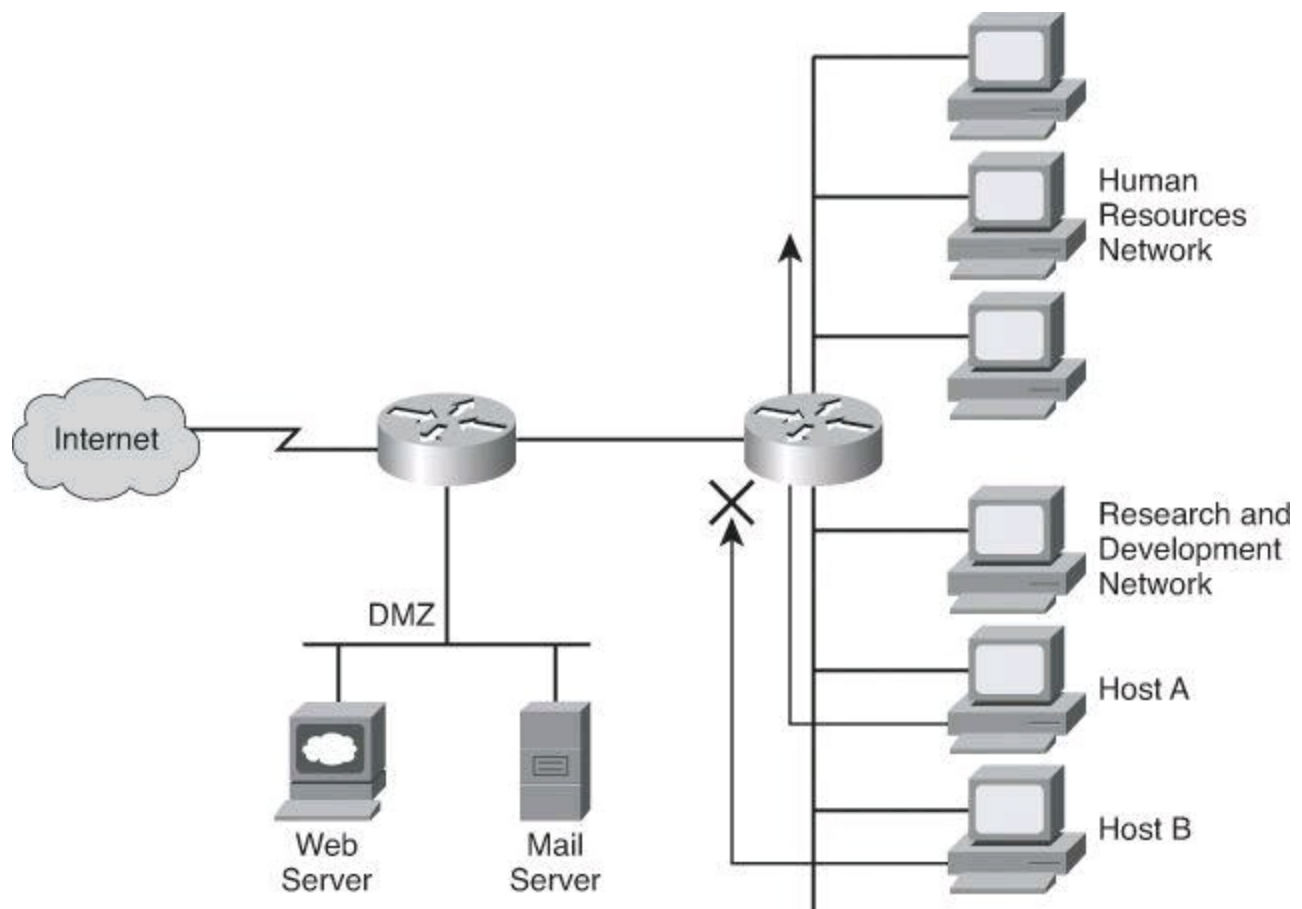


Figure 8-1. Filtering Host B Traffic Ingress Using an ACL

To provide the security benefits of ACLs, you should, at a minimum, configure ACLs at the perimeter of your networks. This configuration provides a basic buffer from the outside network, or from a less-controlled area of your own network into a more sensitive area of your network. On these network edge routers, you should configure ACLs for each network protocol that is configured on the router interfaces.

You can use ACLs to mitigate many threats:

- IP address spoofing (inbound)
- IP address spoofing (outbound)
- DoS TCP SYN attacks (blocking external attacks)
- DoS TCP SYN attacks (using TCP intercept)
- DoS Smurf attacks
- Filtering ICMP messages (inbound)
- Filtering ICMP messages (outbound)
- Filtering traceroute

Note

The System Administration, Networking, and Security (SANS) Institute offers a guide to protecting your Cisco IOS router, *Cisco Router Hardening Step-by-Step*, available at http://www.sans.org/reading_room/whitepapers/firewalls/cisco-router-hardening-step-by-step_794. The National Security Agency (NSA) also offers a guide on this topic, *Router*

Tip

Two terms you'll encounter in this chapter that are commonly confused are egress and ingress. *Egress* refers to traffic leaving the network or device. *Ingress* refers to traffic entering the network or device.

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

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

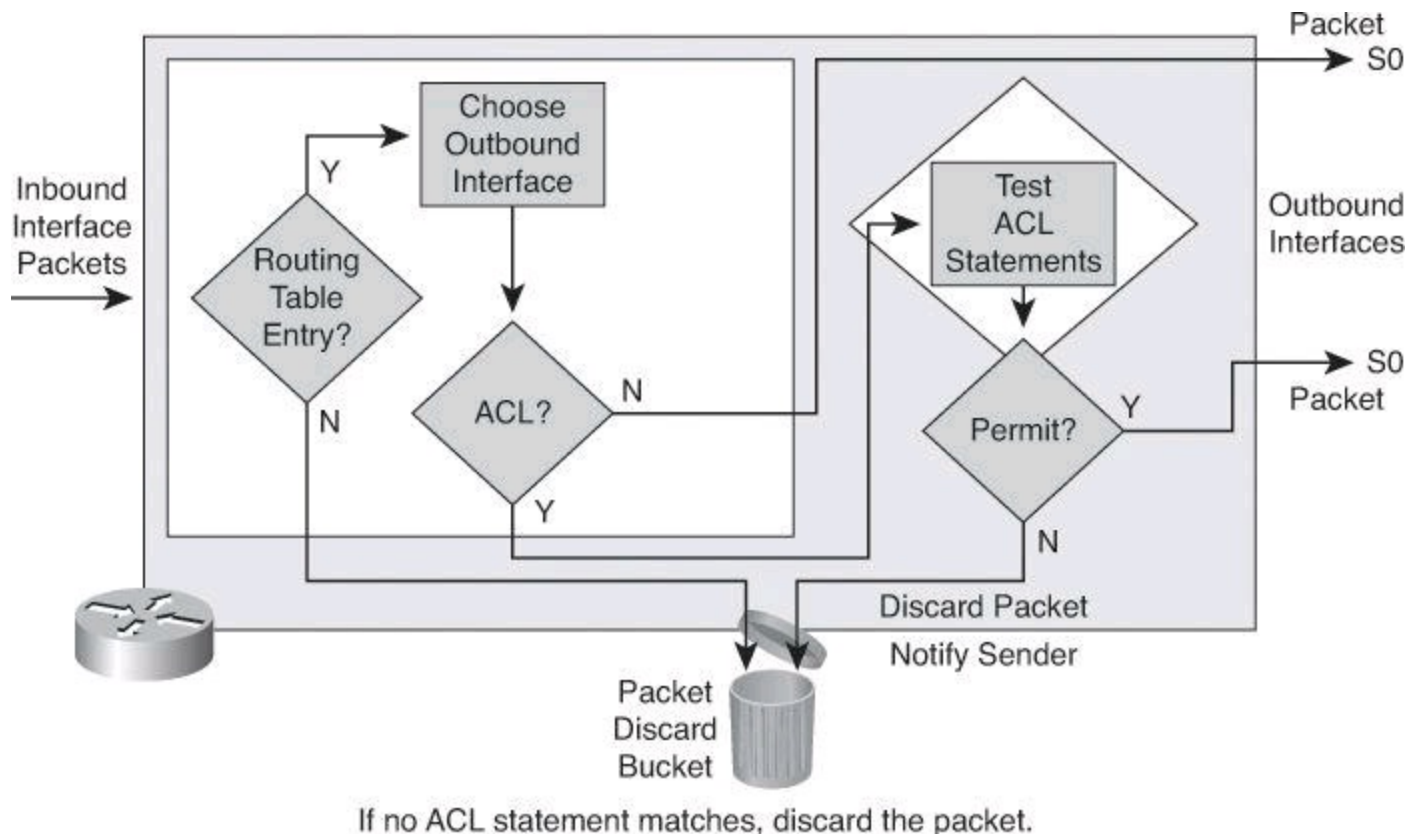
Cisco router security depends strongly on well-written ACLs to restrict access to router network services and to filter packets as they traverse the router.

ACLs express the sets of rules that give added control over packets that enter inbound interfaces, packets that relay through the router, and packets that exit outbound interfaces of the router. ACLs do not act on packets that originate from the router itself. Instead, ACLs are statements that specify conditions of how the router handles the traffic flow through specified interfaces.

ACLs operate in two ways:

- **Inbound:** Incoming packets are processed before they are routed to an outbound interface. An inbound ACL is efficient because it saves the overhead of routing lookups if the packet will be discarded after it is denied by the filtering tests. If the packet is permitted by the tests, it is then processed for routing.
- **Outbound:** Packets arriving on the inside interface are routed to the outbound interface, and then they are processed through the outbound ACL.

[Figure 8-2](#) shows an example of an outbound ACL. When a packet enters an interface, the router checks the routing table to see whether the packet is routable. If the packet is not routable, it is dropped.



If no ACL statement matches, discard the packet.

Figure 8-2. Outbound ACL Operation

Next, the router checks to see whether the destination interface is grouped to an ACL. If the destination interface is not grouped to an ACL, the packet can be sent to the output buffer. Examples of outbound ACL operation are as follows:

- If the outbound interface selected by the routing process is not grouped to an outbound ACL, the packet is sent directly to that outbound interface, which is S0 in the example in [Figure 8-2](#).
- If the outbound interface selected by the routing process is grouped to an outbound ACL, the packet is not sent out directly to that outbound interface, which is S0 in the example in [Figure 8-2](#). The packet must first be tested by the combination of ACL statements associated with that interface. Based on the ACL tests, the packet is permitted or denied.

For outbound lists, to *permit* means to send the packet to the output buffer, and to *deny* means to discard the packet.

With an inbound ACL, when a packet enters an interface, the router checks to see whether the source interface is grouped to an ACL. If the source interface is not grouped to an ACL, the router checks the routing table to see whether the packet is routable. If the packet is not routable, the router drops the packet. Examples of inbound ACL operation are as follows:

- If the inbound interface is E0, for example, which has not been grouped to an inbound ACL, the packet is processed normally and the router checks to determine whether the packet is routable.
- If the inbound interface is E1, for example, which has been grouped to an inbound ACL, the packet is not processed and the routing table is not consulted until the packet is tested by the ACL that is associated with that interface. Based on the ACL tests, the packet is

permitted or denied.

For inbound lists, *permit* means to continue to process the packet after receiving it on an inbound interface, and *deny* means to discard the packet.

ACL statements operate in a sequential, logical order, as shown in [Figure 8-3](#). They evaluate packets from the top down, one statement at a time. If a packet header and an ACL statement match, the rest of the statements in the list are skipped, and the packet is permitted or denied as determined by the matched statement. If a packet header does not match an ACL statement, the packet is tested against the next statement in the list. This matching process continues until the end of the list is reached.

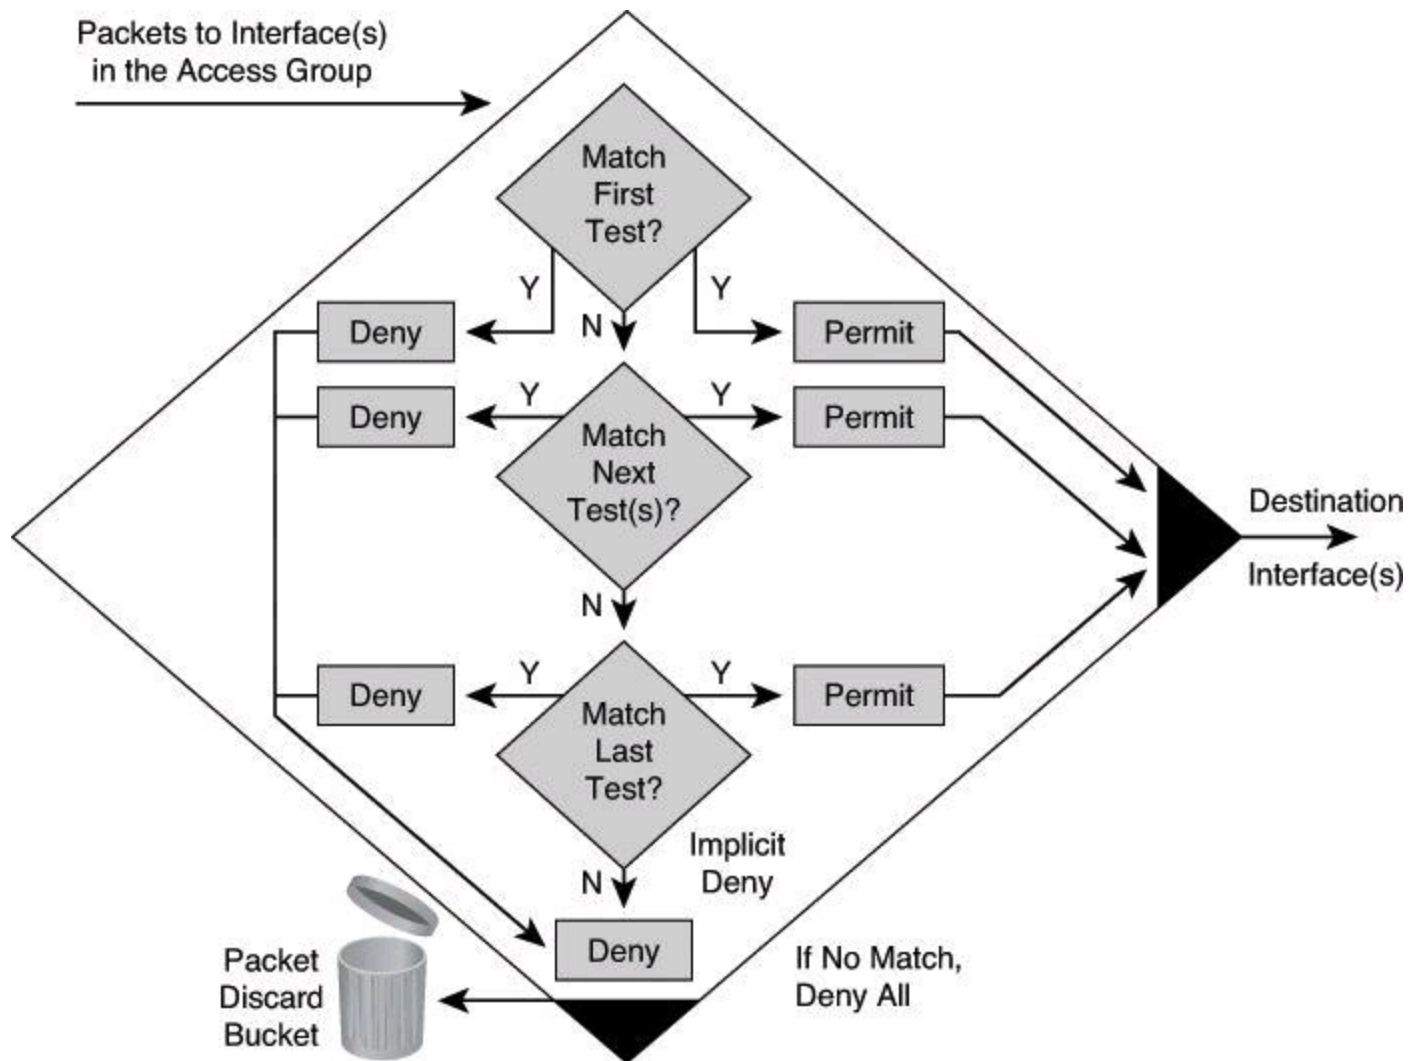


Figure 8-3. Top-Down Process of Tests: Deny or Permit

A final implied statement covers all packets for which conditions did not test true. This final test condition matches all other packets and results in a “deny” instruction. Instead of proceeding into or out of an interface, the router drops all of these remaining packets. This final statement is often referred to as the “implicit **deny any** statement.” Because of this statement, an ACL should have at least one **permit** statement in it; otherwise, the ACL blocks all traffic.

Types of IP ACLs

Cisco routers support two types of IP ACLs:

- **Standard ACLs:** Standard IP ACLs check the source addresses of packets that can be

routed. The result either permits or denies the output for an entire protocol suite, based on the source network, subnet, or host IP address.

- **Extended ACLs:** Extended IP ACLs check both the source and destination packet addresses. They can also check for specific protocols, port numbers, and other parameters, which allows administrators more flexibility and control.

The two general methods you can use to create ACLs are as follows:

- **Numbered ACLs:** Use a number for identification.
- **Named ACLs:** Use an alphanumeric string for identification.

Using numbered ACLs is an effective method on smaller networks with more homogeneously defined traffic. Because each ACL type is limited to an assigned range of numbers, it is easy to determine the type of ACL that you are using.

Specifying an ACL number from 1 to 99 or 1300 to 1999 instructs the router to accept numbered standard IP Version 4 (IPv4) ACL statements. Specifying an ACL number from 100 to 199 or 2000 to 2699 instructs the router to accept numbered extended IPv4 ACL statements.

The named ACL feature allows you to identify IP standard and extended ACLs with an alphanumeric string (name) rather than the numeric representations. Named IP ACLs provide you more flexibility in working with the ACL entries.

Tip

The number of the ACL determines which protocol it is filtering: 1 to 99 and 1300 to 1999 define standard IP ACLs. 100 to 199 and 2000 to 2699 define extended IP ACLs.

Named ACLs have been available since Cisco IOS Software Releases 11.2. Names contain alphanumeric characters. Names cannot contain spaces or punctuation and must begin with an alphabetic character. Named ACLs enable you to add or delete entries within the ACL.

There are two benefits to IP ACL entry sequence numbering:

- You can edit the order of ACL statements.
- You can remove individual statements from an ACL.

Where additions are placed in an ACL depends on whether you use sequence numbers. If you don't specify a sequence number, new ACL statements are placed at the end of the ACL.

Removing ACL statements is also streamlined with sequence numbers. If you want to remove a single entry, you can do it by referencing the entry number (or sequence number) with the **no access-list** command. It is best practice to space out the entry sequence number to leave room to add intermediate entries in the future. As an example, the first entry could be sequenced 10, the second entry 20, and so forth. Should an entry need to be inserted between 10 and 20, you could sequence it 15. Also note that Cisco lets you resequence an ACL and select the increment. A good reference on this can be found at http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsaclseq.html.

Note

ACL Wildcard Masking and VLSM Review

Address filtering occurs when you use ACL address wildcard masking to identify how to check or ignore corresponding IP address bits. However, prior to jumping into ACL wildcard masking, let's review subnetting, including variable-length subnet masking (VLSM), which was covered amply in the CCNA material that you studied for the ICND2 exam. Please consult *CCNA Preparation Library* (Cisco Press) for thorough presentations on IP addresses, subnetting, and VLSM.

Subnetting Overview

To understand ACL filtering in classless interdomain routing (CIDR) and VLSM environments, we must start with a review of subnetting.

Remember that an IP address has 32 bits and comprises two parts—a network ID and a host ID. The length of the network ID and host ID depends on the class of the IP address. The number of hosts available also depends on the class of the IP address.


The default number of bits in the network ID is referred to as the *classful prefix length*. Therefore, a Class A address has a classful prefix length of /8, a Class B address has a classful prefix length of /16, and a Class C address has a classful prefix length of /24.

The number of hosts per subnet available depends upon the number of host ID bits not borrowed.

Subnetting Example: Class C

The subnet address is created by taking address bits from the host-number portion of Class A, Class B, and Class C addresses. Usually a network administrator assigns the subnet address locally. Like IP addresses, each subnet address must be unique.

Each time one bit is borrowed from a host field, there is one less bit remaining in the host field that can be used for host numbers, and the number of host addresses that can be assigned per subnet decreases by a power of 2. [Figure 8-4](#) shows an example for Class C addresses.



Number of Bits Borrowed (s)	Number of Subnets Possible (2^s)	Number of Bits Remaining in Host ID ($8 - s = h$)	Number of Hosts Possible Per Subnet ($2^h - 2$)
1	2	7	126
2	4	6	62
3	8	5	30
4	16	4	14
5	32	3	6
6	64	2	2
7	128	1	0

Figure 8-4. Subnetting a Class C Address

When you borrow bits from the host field, it is important to note the number of additional subnets that are being created each time one more bit is borrowed. Borrowing two bits creates four possible subnets ($2 \times 2 = 4$). Each time another bit is borrowed from the host field, the number of possible subnets increases by a power of 2 and the number of individual host addresses decreases by a power of 2.

The following are examples of how many subnets are available, based on the number of host bits that you borrow:

- Using 3 bits for the subnet field results in 8 possible subnets ($2^3 = 8$).
- Using 4 bits for the subnet field results in 16 possible subnets ($2^4 = 16$).
- Using 5 bits for the subnet field results in 32 possible subnets ($2^5 = 32$).
- Using 6 bits for the subnet field results in 64 possible subnets ($2^6 = 64$).

In general, you can use the following formula to calculate the number of usable subnets, given the number of subnet bits used:

$$\text{Number of subnets} = 2^s \text{ (in which } s \text{ is the number of subnet bits)}$$

In general, you can use the following formula to calculate the number of hosts per subnet, given the number of host bits used:

$$\text{Number of hosts per subnet} = 2^h - 2, \text{ in which } h \text{ is the number of host bits not borrowed.}$$

We subtract two host addresses from the total because one address is reserved as the network

address (the address where all host bits are set to 0) and one address is reserved as the broadcast address (the address where all host bits are set to 1).

Subnetting Example

Subnet a network with a private network address of 172.16.0.0/16 so that it provides ten subnets and maximizes the number of host addresses for each subnet. Answer the following questions:

- How many bits will need to be borrowed for subnets? $2^s = 2^7 = 128$ subnets ($s = 7$ bits).
- What is the new subnet mask? Borrowing 7 host bits = 255.255.254.0 or /23.
- What are the first four subnets?
 - 172.16.0.0
 - 172.16.2.0
 - 172.16.4.0
 - 172.16.6.0
- What is the range of host addresses for the four subnets?
 - 172.16.0.1–172.16.1.254
 - 172.16.2.1–172.16.3.254
 - 172.16.4.1–172.16.5.254
 - 172.16.6.1–172.16.7.254

Variable-Length Subnet Masking

VLSM affords the options of including more than one subnet mask within a network and of subnetting an already subnetted network address. VLSM offers the following benefits:

- **More efficient use of IP addresses:** Without the use of VLSMs, companies must implement a single subnet mask within an entire Class A, B, or C network number.

For example, consider the 172.16.0.0/16 network address divided into subnetworks using /24 masking. One of the subnetworks in this range, 172.16.14.0/24, is further divided into smaller subnetworks using /27 masking, as shown in [Figure 8-5](#). These smaller subnetworks range from 172.16.14.0/27 to 172.16.14.224/27. In [Figure 8-5](#), one of these smaller subnets, 172.16.14.128/27, is further divided using the /30 prefix, which creates subnets with only two hosts, to be used on the WAN links. The /30 subnets range from 172.16.14.128/30 to 172.16.14.156/30. In [Figure 8-5](#), the WAN links used the 172.16.14.132/30, 172.16.14.136/30, and 172.16.14.140/30 subnets from the range.

Subnet 172.16.14.0/24 is divided into smaller subnets.

- Subnet with one mask (/27).
- Then further subnet one of the unused /27 subnets into multiple /30 subnets.




Figure 8-5. Example of Variable-Length Subnet Mask for 172.16.0.0/16

- **Greater capability to use route summarization:** VLSM allows more hierarchical levels within an addressing plan and thus allows better route summarization within routing tables.

For example, in [Figure 8-5](#), subnet 172.16.14.0/24 summarizes all of the addresses that are further subnets of 172.16.14.0, including those from subnet 172.16.14.0/27 and from subnet 172.16.14.128/30.

- **Isolation of topology changes from other routers:** Another advantage to using route summarization in a large, complex network is that it can isolate topology changes from other routers. For example, when a specific link in the 172.16.27.0/24 domain is rapidly fluctuating between being active and inactive (called flapping), the summary route does not change. Therefore, no router that is external to the domain needs to keep modifying its routing table because of this flapping activity.

A Working VLSM Example

The example in [Figure 8-6](#) shows an address space defined by the subnet address 172.16.32.0/20. This address space is used for this portion of the enterprise network, and is generated from subnetting the 172.16.0.0/16 Class B network into multiple /20 subnets.

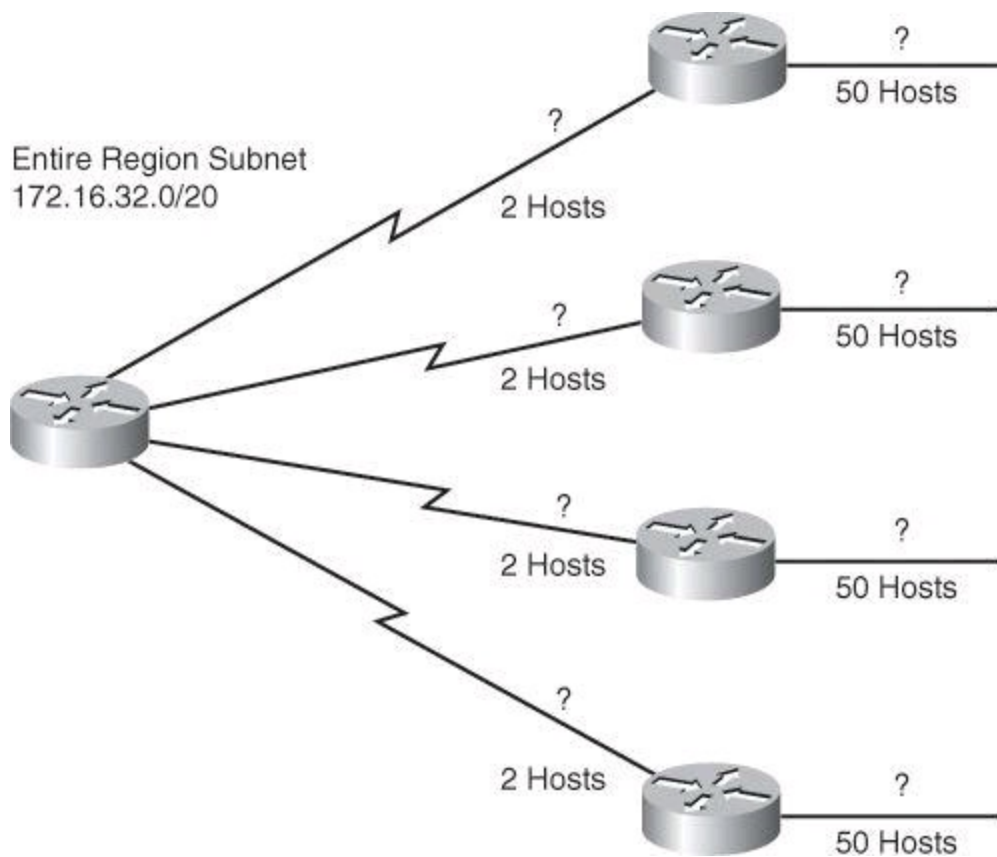


Figure 8-6. Working VLSM Example

How would you define the subnets and masks to assign IP addresses to the WAN links, which require 2 host IP addresses, and the LAN segments, each one showing a requirement of 50 hosts?

By using VLSM, you can further subnet an already subnetted address. Consider again the example address space shown in [Figure 8-6](#). Your region of the enterprise network has a subnet address of 172.16.32.0/20 and you need to assign addresses to multiple LANs, each with 50 hosts, within your region. With VLSM, you can further subnet address 172.16.32.0/20 to give you more network addresses and fewer hosts per network. For example, as shown in [Figure 8-7](#), if you subnet 172.16.32.0/20 to 172.16.32.0/26, you gain 64 (2^6) subnets, each of which could support 62 ($2^6 - 2$) hosts.

Subnetted Address: 172.16.32.0/20
 In Binary 10101100.00010000.00100000.00000000

VLSM Address: 172.16.32.0/26
 In Binary 10101100.00010000.00100000.00000000

	Network	Subnet	VLSM Subnet	Host
1st subnet:	172 . 16	.0010	0000.00	000000=172.16.32.0/26
2nd subnet:	172 . 16	.0010	0000.01	000000=172.16.32.64/26
3rd subnet:	172 . 16	.0010	0000.10	000000=172.16.32.128/26
4th subnet:	172 . 16	.0010	0000.11	000000=172.16.32.192/26
5th subnet:	172 . 16	.0010	0001.00	000000=172.16.33.0/26

Figure 8-7. VLSM for 172.16.32.0/20

To calculate the subnet addresses that are used on the WAN links, further subnet one of the unused

/26 subnets using a similar process.

The resulting address space and subnetting strategy is shown in [Figure 8-8](#). The subnet addresses that are used on the Ethernet LANs are those generated from subdividing the 172.16.32.0/20 subnet into multiple /26 subnets.

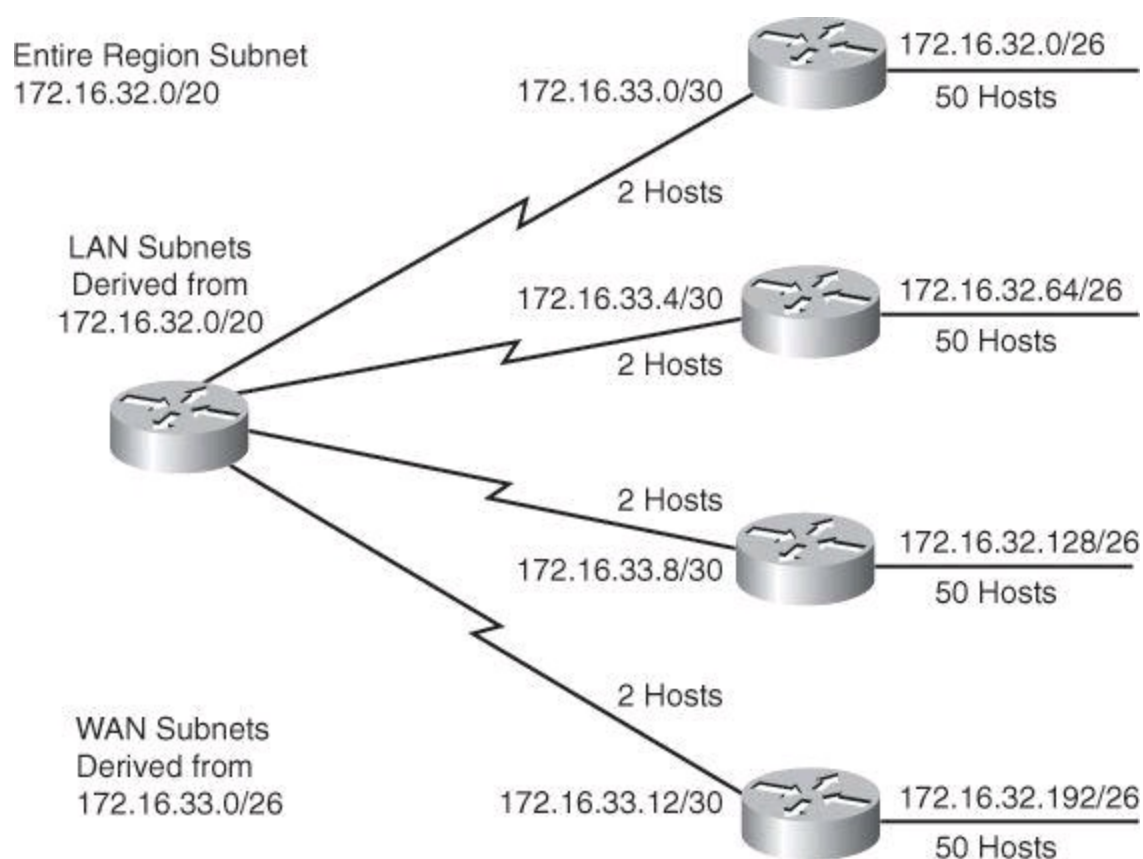


Figure 8-8. LANs and Point-to-Point Subnets for 172.16.32.0/20

The subnet addresses that are used on the WAN links are those generated from subdividing one of the /26 subnets, the 172.16.33.0/26 subnet, into multiple /30 subnets. This provides 16 (2^4) subnets and 2 ($2^2 - 2$) hosts for each of the WAN links.

ACL Wildcard Bits

One of the benefits of proper summarization is the efficient use of wildcard masks in ACLs for packet filtering purposes.

In ACLs, address filtering occurs when you use ACL address wildcard masking to identify how to check or ignore corresponding IP address bits. Wildcard masking for IP address bits uses the numbers 1 and 0 to identify how to treat the corresponding IP address bits, as follows:

- **Wildcard mask bit 0:** Match the corresponding bit value in the address.
- **Wildcard mask bit 1:** Do not check (ignore) the corresponding bit value in the address.

Note

A wildcard mask is sometimes referred to as an inverse mask. Be aware, however, that you should use regular subnet masks if you are configuring ACLs on a Cisco ASA.

[Figure 8-9](#) illustrates how wildcard bits are used to check the corresponding address bits.




Figure 8-9. Wildcard Bits: How to Check the Corresponding Address Bits

By carefully setting wildcard masks, you can permit or deny tests with one ACL statement. You can select a single IP address or any IP address.

Example: Wildcard Masking Process for IP Subnets

In [Figure 8-10](#), an administrator wants to test a range of IP subnets that is to be permitted or denied. Assume that the IP address is a Class B address (the first two octets are the network number), with 8 bits of subnetting (the third octet is for subnets). The administrator wants to use the IP wildcard masking bits to match subnets 172.30.16.0/24 to 172.30.31.0/24.

Match for IP subnets 172.30.16.0/24 to 172.30.31.0/24.

• Address and wildcard mask:

• 172.30.16.0 0.0.15.255

Network. Host

172.30.16.0

Wildcard Mask:	0	0	0	1	0	0	0	0	
	0	0	0	0	1	1	1	1	
	----- Match -----			----- Don't Care -----					
	0	0	0	1	0	0	0	0	= 16
	0	0	0	1	0	0	0	1	= 17
	0	0	0	1	0	0	1	0	= 18
				:					:
	0	0	0	1	1	1	1	1	= 31

Figure 8-10. Wildcard Bits to Match IP Subnets 172.30.16.0 to 172.30.31.0

To use one ACL statement to match this range of subnets, use the IP address 172.30.16.0 in the ACL, which is the first subnet to be matched, followed by the required wildcard mask.

First, the wildcard mask matches the first two octets (172.30) of the IP address using corresponding 0 bits in the first two octets of the wildcard mask.

Because there is no interest in an individual host, the wildcard mask ignores the final octet by using the corresponding 1 bit in the wildcard mask. For example, the final octet of the wildcard mask is 255 in decimal.

In the third octet, where the subnet address occurs, the wildcard mask of decimal 15, or binary 00001111, shown in bold, matches the high-order 4 bits of the IP address. In this case, the wildcard mask matches subnets starting with the 172.30.16.0/24 subnet. For the final (low-end) 4 bits in this octet, shown in gray, the wildcard mask indicates that the bits can be ignored. In these positions, the address value can be binary 0 or binary 1. Thus, the wildcard mask matches subnets 16, 17, 18, and so on up to subnet 31. The wildcard mask does not match any other subnets.

In [Figure 8-10](#), the address 172.30.16.0 with the wildcard mask 0.0.15.255 matches subnets 172.30.16.0/24 to 172.30.31.0/24.

In some cases, you must use more than one ACL statement to match a range of subnets. For example, to match 10.1.4.0/24 to 10.1.8.0/24, use 10.1.4.0 0.0.3.255 and 10.1.8.0 0.0.0.255.

From your previous knowledge acquired in ICND2 (CCNA), you will remember the following command syntax, which we will review in more detail in a moment. So, suppose that you want to block all IP traffic from the range of subnets shown in [Figure 8-10](#); the ACL could be this:

```
Router(config)# access-list 1 deny 172.30.16.0 0.0.15.255
```

The 0 and 1 bits in an ACL wildcard mask cause the ACL to either match or ignore the corresponding bit in the IP address. Working with decimal representations of binary wildcard mask bits can be tedious. For the most common uses of wildcard masking, you can use abbreviations. These abbreviations reduce how many numbers you are required to enter while configuring address test conditions.

Example: Wildcard Masking Process with a Single IP Address

In [Figure 8-11](#), instead of entering 172.30.16.29 0.0.0.0, you can use the string **host 172.30.16.29**. Using the abbreviation **host** communicates the same test condition to the Cisco IOS ACL Software. The command syntax using the abbreviation **host** is as follows:

```
Router(config)# access-list 2 permit host 172.30.16.29
```

- 172.30.16.29 0.0.0.0 matches all of the address bits.
- Abbreviate this wildcard mask using the IP address preceded by the keyword **host** (host 172.30.16.29).




Figure 8-11. ACL and the *host* Keyword

And this is the same as entering the following:

```
Router(config)# access-list 2 permit 172.30.16.29 0.0.0.0
```


Should you enter the full host mask of 0.0.0.0, the router will convert the command for the **host** abbreviation automatically.

Example: Wildcard Masking Process with a Match Any IP Address

In [Figure 8-12](#), instead of entering 0.0.0.0 255.255.255.255, you can use the word **any** by itself as the keyword. Using the abbreviation **any** communicates the same test condition to the Cisco IOS ACL Software. Both syntaxes, which accomplish the same results, are as follow, first in the long form, then using the keyword **any**:

```
Router(config)# access-list 2 permit 0.0.0.0 255.255.255.255
```

- 0.0.0.0 255.255.255.255 ignores all address bits.
- Abbreviate expression with the keyword any.



Wildcard Mask: 255.255.255.255
(Ignore All Bits)

Figure 8-12. ACL and the *any* Keyword

And this is the same as entering the following:

```
Router(config)# access-list 2 permit any
```

Tip

As a matter of fact, you can enter any valid host IP address followed by the wildcard mask of 255.255.255.255 and the router will convert your entry to the keyword **any**, as demonstrated by the following test:

```
YYZ(config)# access-list 3 deny 192.168.16.214 255.255.255.255  
YYZ(config)# do show access-list 3  
Standard IP access list 3  
 10 deny any  
YYZ(config)#
```

Using ACLs to Control Traffic

To configure numbered standard IPv4 ACLs on a Cisco router, you must create a standard IPv4 ACL and activate an ACL on an interface. The **access-list** command creates an entry in a standard IPv4 traffic filter list.

The following is the syntax for the **access-list** command:

```
Router(config)# access-list access-list-number {permit | deny | remark}  
source-  
address [mask] [log]
```

To create an ACL, use 1 to 99 for the *access-list-number* parameter. The first entry is assigned a sequence number of 10, and successive entries are incremented by 10. The default wildcard mask is 0.0.0.0 (only standard ACL). The **remark** argument lets you add a description to the ACL.

Caution

The **no access-list *access-list-number*** command removes the entire ACL.

The **log** option of the **access-list** command causes an informational logging message to be created about the packet that matches the ACE.

Tip

You control the level of messages that are logged to the console using the **logging console** command. You can also have the messages kept in buffer with the **logging buffered** command. The number of system error and debugging messages in the system logging buffer is determined by the configured size of the syslog buffer. This size of the syslog buffer is also set using the **logging buffered** command.

The logging message includes the ACL number, whether the packet was permitted or denied, the source address, and the number of packets. The message is generated for the first packet that matches, and then at five-minute intervals, including the number of packets permitted or denied in the previous five-minute interval. This five-minute interval prevents firewall resources from being consumed for repetitive packets from the same flow.

The **ip access-group** command links an existing ACL to an interface. Only one ACL per protocol, per direction, and per interface is allowed. The syntax for this command is as follows:

```
Router(config-if)# ip access-group access-list-number {in | out}
```

Note

To remove an IP ACL from an interface, first enter the **no ip access-group** command on the interface; then enter the global **no access-list** command to remove the entire ACL.

[Table 8-1](#) provides an example of the steps required to configure and apply a numbered standard ACL on a router.

Table 8-1. Numbered Standard ACL Configuration Procedure

Step	Action	Notes
1	Use the access-list global configuration command to create an entry in a standard IPv4 ACL: RouterX(config)# access-list 1 permit 172.16.0.0 0.0.255.255	Enter the global no access-list <i>access-list-number</i> command to remove the entire ACL. The example statement matches any address that starts with 172.16.x.x. Use the remark option to add a description to your ACL.
2	Use the interface configuration command to select an interface to which to apply the ACL: RouterX(config)# interface ethernet 1	After you enter the interface command, the CLI prompt changes from (config)# to (config-if)#.
3	Use the ip access-group interface configuration command to activate the existing ACL on an interface: RouterX(config-if)# ip access-group 1 out	To remove an IP ACL from an interface, enter the no ip access-group <i>access-list-number</i> command on the interface. This example activates the standard IPv4 ACL 1 on the interface as an out-bound filter.

Example: Numbered Standard IPv4 ACL—Deny a Specific Subnet

[Table 8-2](#) describes the command syntax that is presented in [Example 8-1](#) for the network topology shown in [Figure 8-13](#).

Table 8-2. Numbered Standard IPv4 ACL Example

access-list Command Parameter	Description
1	ACL number that indicates this ACL is a standard list.
deny	Indicates that traffic that matches the selected parameters is not forwarded.
172.16.4.0	IP address of the source subnet.
0.0.0.255	Wildcard mask; 0s indicate positions that must match, 1s indicate “don’t care” positions. The mask with 0s in the first three octets indicates those positions must match; the 255 in the last octet indicates a “don’t care” condition.
permit	Indicates that traffic that matches the selected parameters is forwarded.
any	Abbreviation for the IP address of the source. The abbreviation any indicates a source address of 0.0.0.0 and a wildcard mask of 255.255.255.255; all source addresses will match.

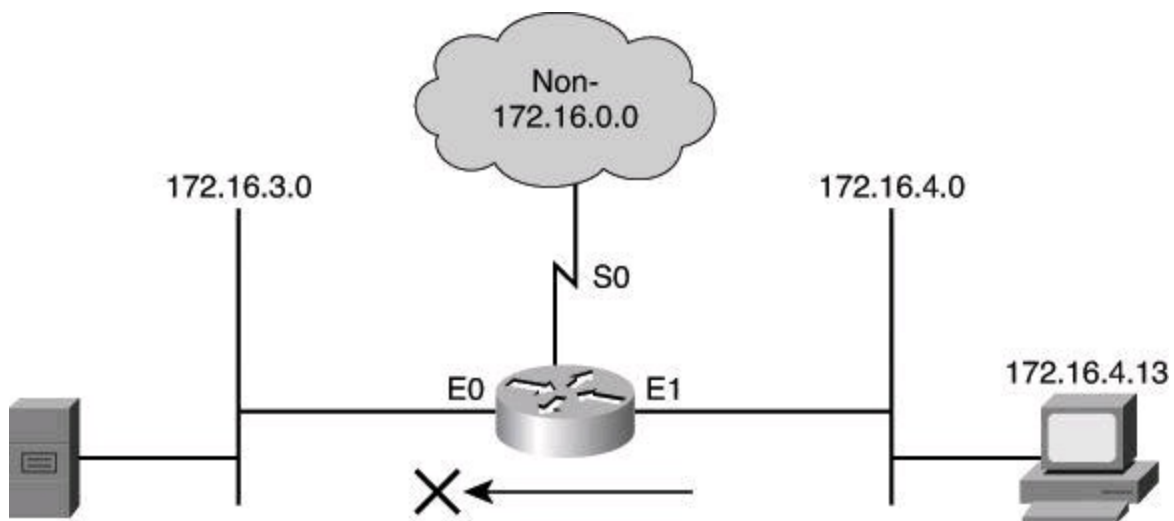


Figure 8-13. Numbered Standard IPv4 ACL

Example 8-1. Numbered Standard IPv4 ACL

[Click here to view code image](#)

```
r1(config)# access-list 1 deny 172.16.4.0 0.0.0.255
r1(config)# access-list 1 permit any
( implicit deny all = access-list 1 deny 0.0.0.0 255.255.255.255 )
r1(config)# interface ethernet 0
r1(config-if)# ip access-group 1 out
```

The ACL shown in [Example 8-1](#) is designed to block traffic from a specific subnet, 172.16.4.0, and to allow all other traffic to be forwarded out E0. However, traffic from subnet 172.16.4.0 is not filtered when leaving by S0.

vty Access

To control traffic into and out of the router (not through the router), you protect the router virtual ports. A virtual port (or virtual terminal line) is commonly referred to as vty. By default, there are five such virtual terminal lines, numbered vty 0 through vty 4. You can configure Cisco IOS Software images to support more than five vty ports.

Restricting vty access is primarily a technique for increasing network security and defining which addresses are allowed Telnet access to the router EXEC process.

The command used to filter Telnet traffic is as follows:

```
Router(config-line)# access-class access-list-number {in | out}
```

Filtering Telnet traffic is typically considered an extended IP ACL function because it filters a higher-level protocol. However, because you are using the **access-class** command to filter incoming or outgoing Telnet sessions by source address and apply filtering to vty lines, you can use standard IP ACL statements to control vty access.

The following example applies a standard ACL to control vty access. It permits any device on network 192.168.1.0 0.0.0.255 to establish a virtual terminal session with the router. Of course, the user must know the appropriate passwords to enter user mode and privileged

mode. This example assumes that the protocol used to exchange with the router is Telnet; however, other protocol can be filtered by the **access-class** command, such as Secure Shell (SSH).

```
R1 (config) # access-list 12 permit 192.168.1.0 0.0.0.255
!  
R1 (config) # line vty 0 4  
R1 (config-line) # access-class 12 in
```

Notice that identical restrictions have been set on every vty line (0 to 4) because you cannot control on which vty line a user will connect. The implicit **deny any** statement still applies to the ACL when it is used as an access class entry.

Numbered Extended IPv4 ACL

To provide more precise traffic-filtering control, use extended IPv4 ACLs, which are named or numbered 100 to 199 and 2000 to 2699, because extended IPv4 ACLs check for the source and destination IPv4 addresses. In addition, with an extended ACL statement, you can specify the protocol and, optionally, TCP or UDP application to filter traffic more precisely. To specify an application, you can configure either the port number or name of a well-known application.

[Table 8-3](#) shows some well-known port numbers of IP protocols that you can use with extended ACLs.

Table 8-3. Well-Known Port Numbers and IP Protocols

Well-Known Port Number (Decimal)	IP Protocol
20 (TCP)	FTP data
21 (TCP)	FTP control
22 (SSH)	Secure Shell
23 (TCP)	Telnet
25 (TCP)	SMTP
53 (TCP/UDP)	DNS
69 (UDP)	TFTP
80 (TCP)	HTTP

To configure numbered extended IPv4 ACLs on a Cisco router, create an extended IPv4 ACL and activate that ACL on an interface. Use the **access-list** command to create an entry to express a condition statement in a complex filter. [Table 8-4](#) explains the syntax of the command for configuring a numbered extended ACL, shown here:

```
Router (config) # access-list access-list-number {permit | deny}  
 protocol source source-wildcard [operator port] destination  
 destination-wildcard [operator port] [established] [log]
```

Table 8-4. Command Parameters for a Numbered Extended ACL

access-list Command Parameter	Description
<i>access-list-number</i>	This parameter identifies the list using a number in the ranges of 100 to 199 or 2000 to 2699.
<i>permit deny</i>	This parameter indicates whether this entry allows or blocks the specified address.
<i>protocol</i>	This parameter can be ip , tcp , udp , icmp , gre , or igrp .
<i>source</i> and <i>destination</i>	These parameters identify the source and destination IP addresses.
<i>source-wildcard</i> and <i>destination-wildcard</i>	In the wildcard mask, 0s indicate positions that must match, 1s indicate “don’t care” positions.
<i>operator</i> [<i>port</i> <i>app_name</i>]	The operator can be lt (less than), gt (greater than), eq (equal), or neq (not equal). The port number referenced can be either the source port or the destination port, depending on where in the ACL the port number is configured. As an alternative to the port number, well-known application names can be used (for example, telnet , ftp , smtp).
established	This option is for inbound TCP only. It allows TCP traffic to pass if the packet is a response to an outbound initiated session. This type of traffic has the acknowledgment (ACK) bits set. (See Example 8-3.)
log	This option sends a logging message to the console.

Note

The syntax of the **access-list** command that is presented in [Table 8-4](#) is representative of the TCP protocol form. Not all parameters and options are given. For the complete syntax of all forms of the command, refer to the appropriate Cisco IOS Software documentation available at Cisco.com.

To link an existing extended ACL to an interface, use the following command:

```
Router(config-if)# ip access-group access-list-number (in | out)
```

Only one ACL per protocol, per direction, and per interface is allowed.

[Table 8-5](#) describes the parameters of the **ip access-group** command.

Table 8-5. *ip access-group* Command Parameters

Parameter	Description
<i>access-list-number</i>	Indicates the number of the ACL that is to be linked to an interface
<i>in out</i>	Selects whether the ACL is applied as an input or output filter; <i>out</i> is default

[Example 8-2](#) shows the configuration for denying FTP access from subnet 172.16.4.0/24 to subnet 172.16.3.0/24, for the network depicted in [Figure 8-14](#). All other traffic is allowed. By applying the ACL to the E0 interface in an outbound direction, traffic is dropped just before it would be transmitted out the E0 interface.

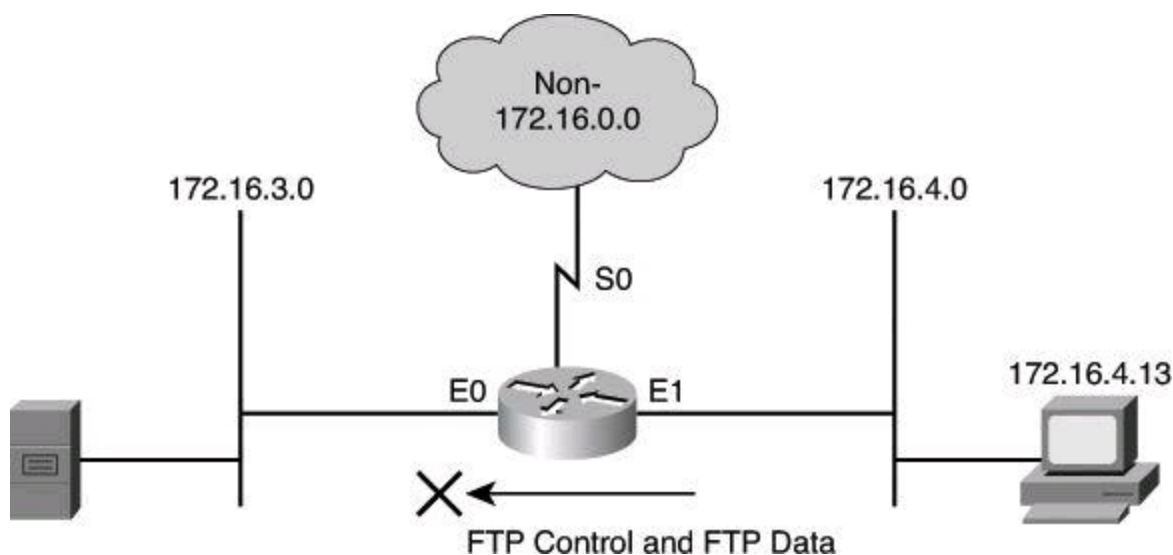


Figure 8-14. Numbered Extended IPv4 ACL Example

Example 8-2. Numbered Extended IPv4 ACL Example

[Click here to view code image](#)

```
r1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.255 172.16.3.0
0.0.0.255 eq 21
r1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.255 172.16.3.0
0.0.0.255 eq 20
r1(config)# access-list 101 permit ip any any
(implicit deny all)
(access-list 101 deny ip 0.0.0.0 255.255.255.255 0.0.0.0
255.255.255.255)
r1(config)# interface ethernet 0
r1(config-if)# ip access-group 101 out
```

Notice the use of wildcard masks to define how to check or ignore the corresponding IP address bits. In the example in [Figure 8-14](#), the first access control entry (ACE) matches source addresses in the 172.16.4.0 subnet. The way to check only the first 3 bytes on this prefix is to follow the prefix with the wildcard mask 0.0.0.255. Zeroes (0) in the wildcard mask tell the ACL to check the corresponding bit of packet addresses with the prefix. Ones (1) in the wildcard mask tell the ACL not to check the corresponding bit of packet addresses with the prefix.

In [Example 8-2](#), notice the wildcard mask. As explained earlier in the chapter, using wildcard masks is a valuable tool for creating effective and efficient ACLs, especially in CIDR environments.

In [Example 8-3](#), the **established** parameter of the extended ACL allows responses to traffic that originate from the mail host, 200.1.1.2, to return inbound on the serial 0 interface. A match occurs if the TCP datagram has the ACK bit or reset (RST) bit set, which indicates that the packet belongs to an existing connection. Without the **established** parameter in the ACL statement, the mail host could receive SMTP traffic but not send it.

Caution

ACLs using the keyword **established** are not a substitute for a stateful firewall. The ACL only checks whether the ACK (acknowledgement) flag bit is turned on in the TCP header, without reference to other prior transmission. In other words, as long as the ACK bit is turned on and the other filtering criteria mentioned in the ACL entry are valid, the router will execute the action. It will not check to determine whether a proper TCP three-way handshake was done. It is therefore easy to fool the router using a packet-crafting tool.

Example 8-3. Using the *established* Keyword

[Click here to view code image](#)

```
Router(config)# access-list 102 permit tcp any host 200.1.1.2
established
Router(config)# access-list 102 permit tcp any host 200.1.1.2 eq smtp
Router(config)# interface serial 0
Router(config-if)# ip access-group 102 in
```

Displaying ACLs

When you finish the ACL configuration, use the **show** commands to verify the configuration. Use the **show access-lists** command to display the contents of all ACLs, as shown in [Example 8-4](#). By entering the ACL name or number as an option for this command, you can display a specific ACL. To display the contents of only the IP ACLs, use the **show ip access-list** command.

Example 8-4. *show access-lists* Command Output

[Click here to view code image](#)

```
Router# show access-lists
Extended IP access list OUTBOUND
10 permit tcp 10.0.0.0 0.255.255.255 any eq www (67 matches)
20 permit tcp 10.0.0.0 0.255.255.255 any eq 443 (1190 matches)
30 permit udp 10.0.0.0 0.255.255.255 any eq domain
40 deny ip any any (41 matches)
```

Fast switching uses the route cache to store information about packet flows. The first packet in a flow is looked at thoroughly by the router, which performs Layer 3-to-Layer 2 mapping, selects the best route, the outbound interface, and its ACL, and decides whether the packet is allowed to leave. The decision taken on the first packet is stored in the route cache so that any subsequent packets from the same flow can be fast-switched or sent to Cisco Express Forwarding (CEF) without being looked at by the CPU, the interface ACL, and so forth.

If you enable CEF, a Layer 3 IP switching technology, and then create an ACL that uses the **log** keyword, the packets that match the ACL are not CEF switched. They are fast switched. ACL logging disables CEF.

The **show ip interfaces** command displays IP interface information and indicates whether any IP ACLs are set on the interface. In the **show ip interfaces** command output shown in [Example 8-5](#), IP ACL OUTBOUND has been configured on the FastEthernet 0/0 interface as an inbound ACL. No outbound IP ACL has been configured on the FastEthernet interface.

Example 8-5. *show ip interfaces* Command Output

[Click here to view code image](#)

```
r1# show ip interfaces FastEthernet 0/0
FastEthernet 0/0 is up, line protocol is up
  Internet address is 10.0.2.1/24
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound access list is OUTBOUND
  Proxy ARP is enabled
  <text omitted>
```

Enhancing ACLs with Object Groups

ACLs provide basic security to the network by permitting or blocking certain types of traffic. ACLs use IP addresses, protocols, and ports to filter network traffic. In some networks, the number of ACLs can become quite large and difficult to manage. The Object Groups for ACLs feature simplifies this problem. By using the Object Groups for ACLs feature, the administrator can group users, devices, or protocols into object groups and create access control entries. Each ACE can then permit or deny a group of users access to a group of servers or services.

The benefits of using Object Groups for ACL include the following:

- Increased performance when network traffic is heavy.
- Reduced storage in NVRAM compared to conventional ACLs.
- Separate ownership of the components of an ACE. For example, you can create an ACE

where each department within an organization can control its group membership. You can also create an ACE to permit or deny the departments to contact each other.

- Allows you to create an object group that contains other object groups. For example, you can create an ENG-ALL address group, which contains the ENG-EAST and ENG-WEST address groups.

In [Example 8-6](#), an access rule results in eight access control entries in an ACL.

Example 8-6. Example of an ACL Without Object Group

[Click here to view code image](#)

```
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.12 eq www
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.12 eq ftp
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.13 eq www
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.13 eq ftp
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.12 eq www
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.12 eq ftp
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.13 eq www
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.13 eq ftp
```

With object groups, as shown in [Example 8-7](#), you can modularize the sources in this example using object group SOURCES, which is a network object group. Another network object group called DESTINATIONS is used to modularize the destination IP addresses, and a service object group called APPLICATIONS is used to modularize the TCP ports in the example. The resulting ACL has one line and still matches all the sources, all the destinations, and all the services. Notice how object groups extend the conventional ACL syntax to support object group–based ACLs and also add new keywords along with the source and destination addresses and ports.

Example 8-7. Example of an ACL Using Object Groups

[Click here to view code image](#)

```
object-group network SOURCES
 host 10.6.252.65 host 10.6.252.66
object-group network DESTINATIONS
 host 171.8.2.12 host 171.8.2.13
object-group service APPLICATIONS
 tcp www tcp ftp
access-list 100 deny object-group APPLICATIONS object-group SOURCES
object-group
 DESTINATIONS
```

When configuring object group–based ACLs, all features that use or reference conventional ACLs are compatible with object group–based ACLs. Feature interactions for conventional ACLs are the same with object group–based ACLs.

These rules also apply when creating object–group based ACLs:

- You can add to, delete from, or change objects in an object group membership list

dynamically (meaning without deleting and redefining the object group).

- You can add to, delete from, or change objects in an object group membership list without redefining the ACL ACE that is using the object group (meaning changing the object group without deleting the ACE and then redefining the ACE after the change).
- You can add objects to groups, delete them from groups, and then ensure that the changes are properly functioning within the object group–based ACL without reapplying the ACL to the interface.

The following components can be modularized in network object groups:

- Any IP address—includes a range from 0.0.0.0 to 255.255.255.255 (specified using the **any** command)
- Host IP addresses
- Hostnames
- Other network object groups
- Ranges of IP addresses
- Subnets

The following components can be modularized in service object groups:

- Source and destination protocol ports (such as Telnet or Simple Network Management Protocol [SNMP])
- ICMP types (such as echo, echo-reply, or host-unreachable)
- Top-level protocols (such as TCP, UDP, or Encapsulating Security Payload [ESP])
- Other service object groups

ACL Considerations

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

- **Base your ACLs on your security policy:** Unless you anchor the ACL in a comprehensive security policy, you cannot be certain that it will effectively control access in the way that access needs to be controlled.
- **Write it out:** Never sit down at a router and start to develop an ACL without first spending some time in design. The best ACL developers suggest that you write out a list of things that you want the ACL to accomplish. Start with something as simple as “This ACL must block all SNMP access to the router except for the SNMP host at 172.19.1.15.”
- **Set up a development system:** Whether you use your laptop PC or a dedicated server, you need a place to develop and store your ACLs. Word processors or text editors of any kind are suitable, as long as you can save the files in ASCII text format. Build a library of your most commonly used ACLs and use them as sources for new files. ACLs can be pasted into the router running configuration (requiring console or Telnet access) or can be stored in a router configuration file. The system that you choose should support Trivial FTP (TFTP) to make it easy to transfer any resulting configuration files to the router.
- **Access list comments:** During the configuration of the ACL, consider adding comments explaining the goal of this ACL.

Note

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

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

You should consider several caveats when working with ACLs:

- **ACLs by themselves are stateless:** ACLs are a stateless access control mechanism. As such, they may not fully support application layer behavior that requires a stateful inspection engine. Consider implementing the Cisco IOS Zone-based Policy Firewall if you need stateful inspection. Also, because ACLs check only Layer 3 and Layer 4, they can't stop packets with a malicious Layer 7 payload.
- **Only one ACL per interface, per protocol stack, per direction:** Only one ACL per protocol, per direction, and per interface is allowed. Multiple ACLs are permitted per interface, but each must be for a different protocol or different direction.
- **Implicit deny all:** All Cisco ACLs end with an implicit **deny all** statement. Although you might not actually see this statement in your ACLs, they do exist. Every ACL should have at least one **permit** statement. Otherwise, all traffic is denied. You should create the ACL before applying it to an interface. With most versions of Cisco IOS Software, an interface that has an empty ACL applied to it permits all traffic.
- **Standard ACL limitation:** Because standard ACLs are limited to packet filtering on only source addresses, you might need to create extended ACLs to implement your security policies.
- **Order of specific statements:** Certain ACL statements are more specific than others; therefore, you need to place them higher in the ACL. For example, blocking all UDP traffic at the top of the list negates the blocking of SNMP packets lower in the list. Take care that statements at the top of the ACL do not negate any statements found lower in the list.
- **Directional filtering:** Cisco ACLs have a directional filter that determines whether they examine inbound packets (ingress traffic, traffic coming toward the interface) or outbound packets (egress traffic, traffic going away from the interface). Always double-check the direction of data that your ACL is filtering.
- **Modifying ACLs:** Cisco IOS Software Release 12.2 and earlier always appends new statements to an existing ACL to the bottom of the ACL. Because of the inherent top-down statement evaluation order of ACLs, these new entries can render the ACL unusable. When a new statement does render the ACL unusable, you must create a new ACL with the correct statement ordering, delete the old ACL, and assign the new ACL to the router interface. If you are using Cisco IOS Software Release 12.3 and later, you can use

sequence numbers to ensure that you are adding a new statement into the ACL in the correct location. The ACL is processed top-down based on the sequence numbers of the statements (lowest to highest).

- **Special packets:** Router-generated packets, such as routing table updates, are not subject to outbound ACL statements applied to interfaces on the source router. If your security policy requires filtering these types of packets, inbound ACLs on adjacent routers or other router filter mechanisms using ACLs must do the filtering task.
- **Extended ACL placement:** If you use extended ACLs on routers too far from the source that you need to filter, packets flowing to other routers and interfaces might be adversely affected. Always consider placing extended ACLs on routers as close as possible to the source that you are filtering.
- **Standard ACL placement:** Because standard ACLs filter packets based on the source address, placing these ACLs too close to the source can adversely affect packets destined to other destinations. Always place standard ACLs as close to the destination as possible.

Configuring ACLs for Threat Control Using Cisco Configuration Professional

Rules define how the router will respond to a particular kind of traffic. Using Cisco Configuration Professional (CCP), you can create access rules that cause the router to block certain types of traffic while permitting other types, create NAT rules that define the traffic that is to receive address translation, and create IPsec rules that specify which traffic is to be encrypted.

Rules in Cisco Configuration Professional

Cisco Configuration Professional uses the high-level concept of rules to define how the router will respond to a particular kind of traffic. ACLs are rules, but in CCP, not all rules are ACLs. Rules that can be created using CCP include the following:

- ACLs
- NAT rules
- IPsec rules
- Network Admission Control (NAC) rules
- Firewall rules
- Quality of service (QoS) rules
- Unsupported rules
- Externally defined rules

Using CCP, you can create access rules that cause the router to block certain types of traffic while permitting other types, NAT rules that define the traffic that is to receive address translation, and IPsec rules that specify which traffic is to be encrypted. CCP also provides default rules that are used in guided configurations, and that you can examine and use when you create your own access rules. CCP also allows you to view rules that were not created using CCP, called external rules, and rules with syntax that CCP does not support, called unsupported rules.

On the other hand, ACLs can be used for multiple functions. An ACL is nothing more than a traffic classification mechanism that works at Layers 3 and 4. As such, it can be used to classify traffic. The

action that then applies to traffic depends on where you apply the ACL. If you apply the ACL to router interfaces, the action is to permit or deny traffic traversing the interface in a particular direction. However, if you apply the ACL to a class map in Modular Policy Framework, the action might be to inspect the traffic that matches the ACL. This chapter focuses on ACLs applied to interfaces for filtering purposes.

Working with ACLs in CCP

The ACL Summary window in CCP (choose **Configure > Router > ACL > ACL Summary** to access it), shown in [Figure 8-15](#), enables you to view a summary of the rules in the router's configuration and to navigate to other windows to create, edit, or delete rules.

The screenshot shows the Cisco Configuration Professional (CCP) interface. The breadcrumb navigation is **Configure > Router > ACL > ACL Summary**. The left sidebar shows a tree view with **ACL Summary** selected under the **ACL** folder. The main content area displays the **Cisco CP Rules (ACLs) Summary** table.

Category	No. of Rules	Description
Access Rules	0	From running config; Supported.
NAT Rules	1	From running config; Supported.
IPSec Rules	0	From running config; Supported.
NAC Rules	0	From running config; Supported.
Firewall Rules	0	From running config; Supported.
QoS Rules	0	From running config; Supported.
Unsupported Rules	0	From running config; Not supported; Created outside Cisco CP.
Externally-defined Rules	2	From running config; Supported; Created outside Cisco CP.
Cisco CP Default Rules	11	Cisco CP-provided defaults

Figure 8-15. CCP ACL Summary

ACL Configuration Scenario Using CCP

To demonstrate the configuration of ACLs using CCP and the router CLI, we will step through the configuration scenario depicted in [Figure 8-16](#). In this example, we will configure an outbound policy for the internal network, allowing only certain applications and protocols to originate traffic from the inside network.

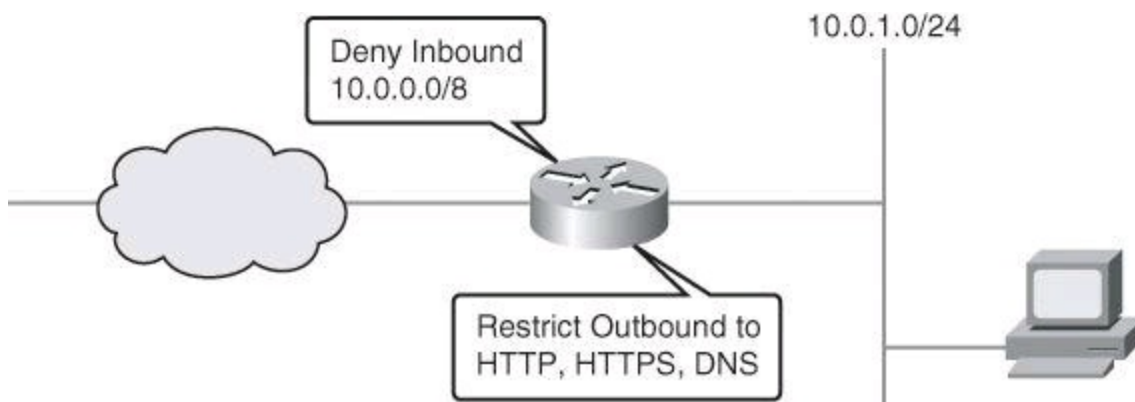


Figure 8-16. ACL Configuration Scenario

The configuration steps are as follows:

- Step 1.** Configure the ACL.
- Step 2.** Apply the ACL to interfaces.
- Step 3.** Verify and monitor the ACL.

ACL Editor

The ACL Editor is most commonly used to define the traffic that you want to permit or deny entry to your LAN or exit from your LAN, but ACLs can be used for other purposes as well, such as identifying traffic for NAT or for encryption.

The upper portion of the ACL Editor window, shown in [Figure 8-17](#), lists the access rules that have been configured on this router. The lower portion of the window lists the rule entries associated with the selected rule. A rule entry consists of criteria that incoming or outgoing traffic is compared against, and the action to take on traffic matching the criteria. If traffic does not match the criteria of any of the entries in this box, it is dropped.

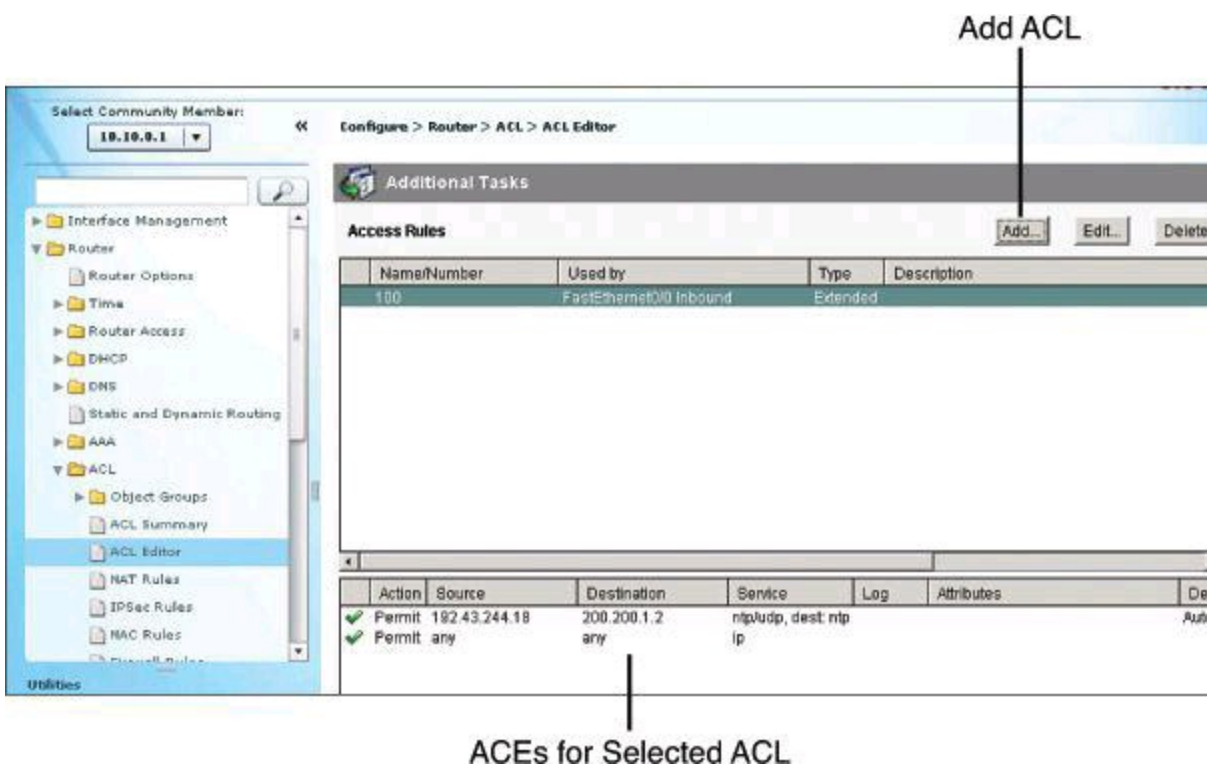


Figure 8-17. ACL Editor

To add an ACL, click the **Add** button in the top-right corner of the window. This opens the Add a Rule dialog box, discussed in the following section.

Adding Rules

The Add a Rule dialog box, shown in [Figure 8-18](#), lets you add a rule. (A similar dialog box, Edit a Rule, is displayed if you select a rule in the Access Rules pane and click Edit.) In the Add a Rule dialog box, you can define a name and add a description for the rule; add, change, reorder, or delete rule entries (also known as access control entries, or ACEs); and associate the rule to an interface.

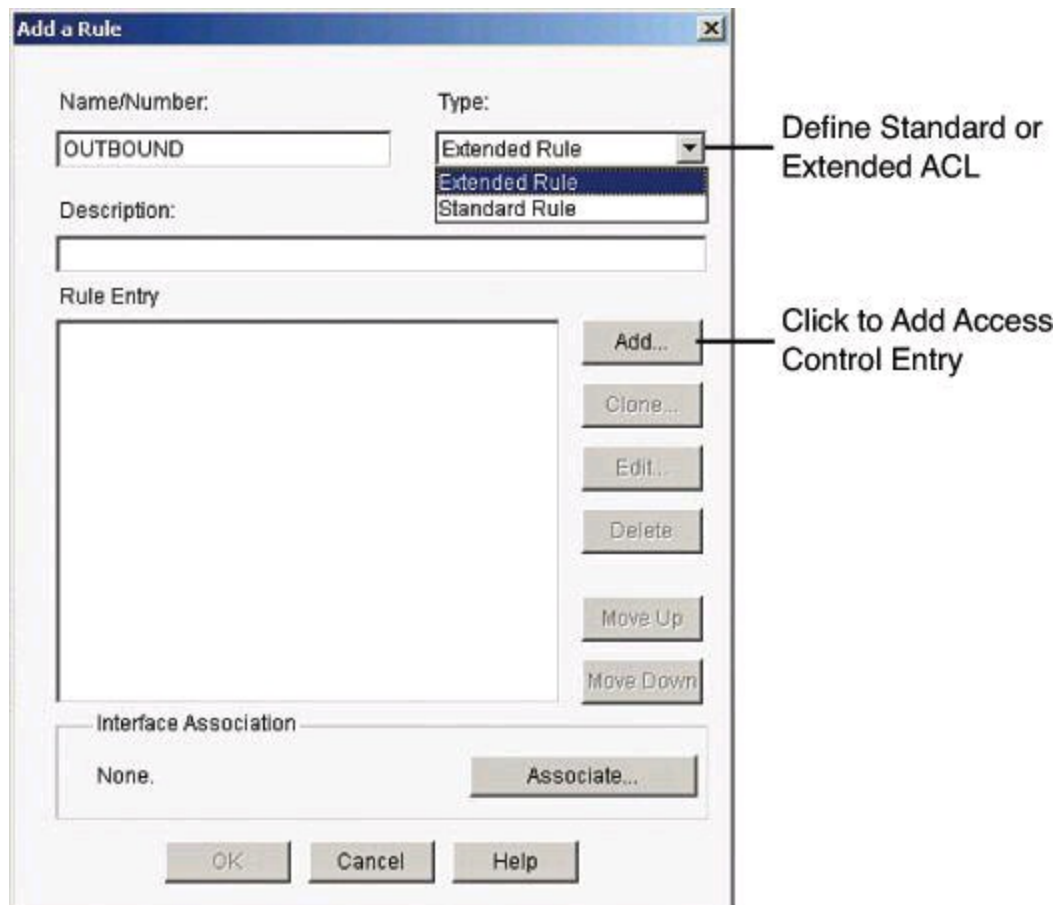


Figure 8-18. Adding a Rule with CCP

The Rule Entry pane shows the entries that make up the rule. You can add, edit, and delete entries by using the buttons to the right of the pane. You can also reorder entries to change the order in which they are evaluated. Observe the following guidelines when creating rule entries:

- There must be at least one **permit** statement in the list; otherwise, all traffic will be denied.
- A **permit all** or **deny all** entry in the list must be the last entry. Remember, a default **deny all** entry is implicit and will not show in the list.
- Standard entries and extended entries cannot be mixed in the same rule.
- No duplicate entries can exist in the same rule.

Applying the log Keyword on Explicit permit or deny ip any any

The **log** keyword on the final **permit ip any any** can prove useful for testing valid configurations. After the testing period is over, you can apply the **log** keyword, changing the

permit ip any any log to a **deny ip any any log**. This will still allow you to log the data required to characterize the threat(s) while still blocking the traffic.

ACEs can be added by clicking the Add button in the Rule Entry section of the dialog box.

The Add an Extended Rule Entry dialog box, shown in [Figure 8-19](#), lets you configure the logic of the specific ACE. An extended rule entry allows you to permit or deny traffic based on its source and destination and on the protocol and service specified in the packet. In this dialog box you will complete the following sections:

- **Action:** Select the action you want the router to take when a packet matches the criteria in the rule entry. The choices are Permit and Deny.
- **Description:** Add a description for the ACE.
- **Source Host/Network:** Specify the source IP address criteria that the traffic must match. The fields in this area of the dialog box change based on the value of the Type field. They are IP Address, Wildcard Mask, Hostname, or Network Object Group.
- **Destination Host/Network:** Specify the source IP address criteria that the traffic must match. The fields and information to complete are similar to those of the Source Host/Network section, but using the components to match a destination address.
- **Protocol and Service:** Select the protocol and service, if applicable, that you want the entry to apply to. The information that you provide differs from protocol to protocol. Click the protocol to see what information you need to provide. For TCP and UDP, you can use logical operators such as these in the Service field:
 - =: The rule entry applies to the value that you enter in the field to the right.
 - !=: The rule entry applies to any value except the one that you enter in the field to the right.
 - <: The rule entry applies to all port numbers lower than the number you enter.
 - >: The rule entry applies to all port numbers higher than the number you enter.
- **Range:** The entry applies to the range of port numbers that you specify in the fields to the right

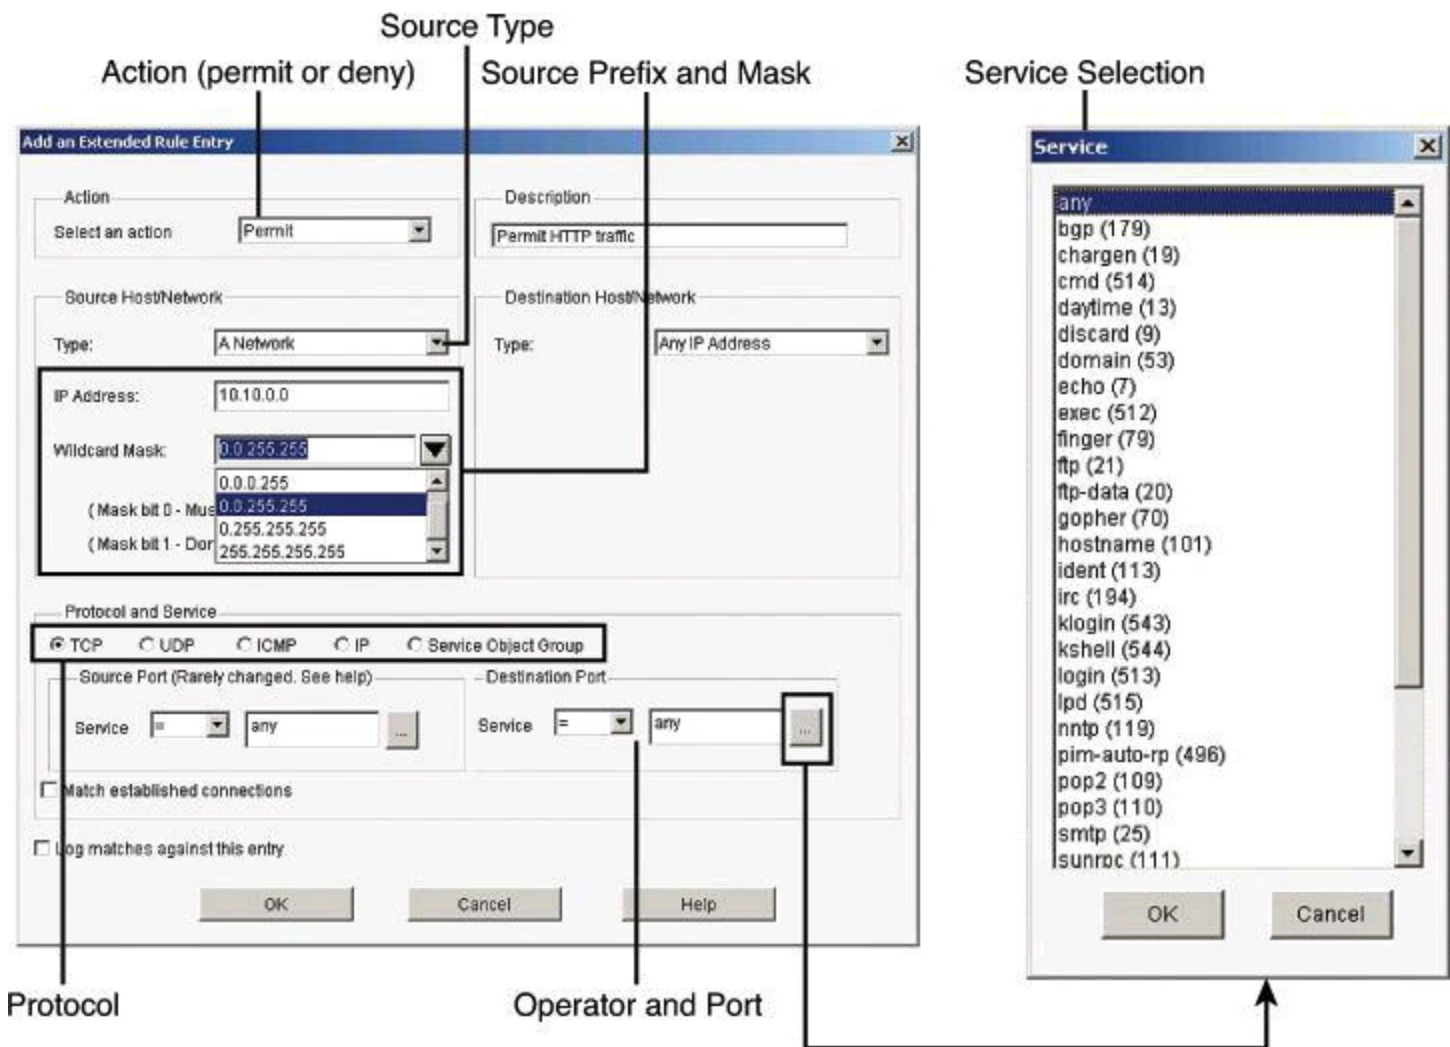


Figure 8-19. Adding an Extended Rule

If you do not remember the name or number, click the ... button and select the value you want from the Service dialog box.

Associating Rules with Interfaces

ACLs and other rules, by themselves, do nothing. You must apply them to interfaces and other router components. To apply an ACL to an interface, select a rule and click the **Associate** button to select the interface and direction in the Associate with an Interface dialog box, as shown in [Figure 8-20](#). If the Associate button is not enabled, you can associate the rule with an interface by double-clicking the interface in the Interfaces and Connections window and using the Associate tab. The following are the options:

- **Select an Interface:** Select the interface to which you want this rule to apply.
- **Specify a Direction:** If you want the router to check ingress packets arriving at the interface, click **Inbound**. The router checks for a match with the rule before routing it; the router accepts or drops the packet based on whether the rule states **permit** or **deny**. If you want the router to match egress packets, click **Outbound**.

Notice in [Figure 8-20](#) how important it is to use the correct terminology. In the example, the ACL is called **OUTBOUND**, as it is aimed at filtering traffic originating on the internal network on its way to external destinations. The ACL is applied inbound to the FastEthernet0/1.4 interface because it will match ingress traffic on the

interface, on its way out to the external network.

- **If another rule is already associated with the interface:** If an information box appears that states that another access rule is associated with the interface and direction you specified, you either can cancel the operation or can continue by appending the rule entries to the rule that is already applied to the interface or by disassociating the rule with the interface and associating the new rule.

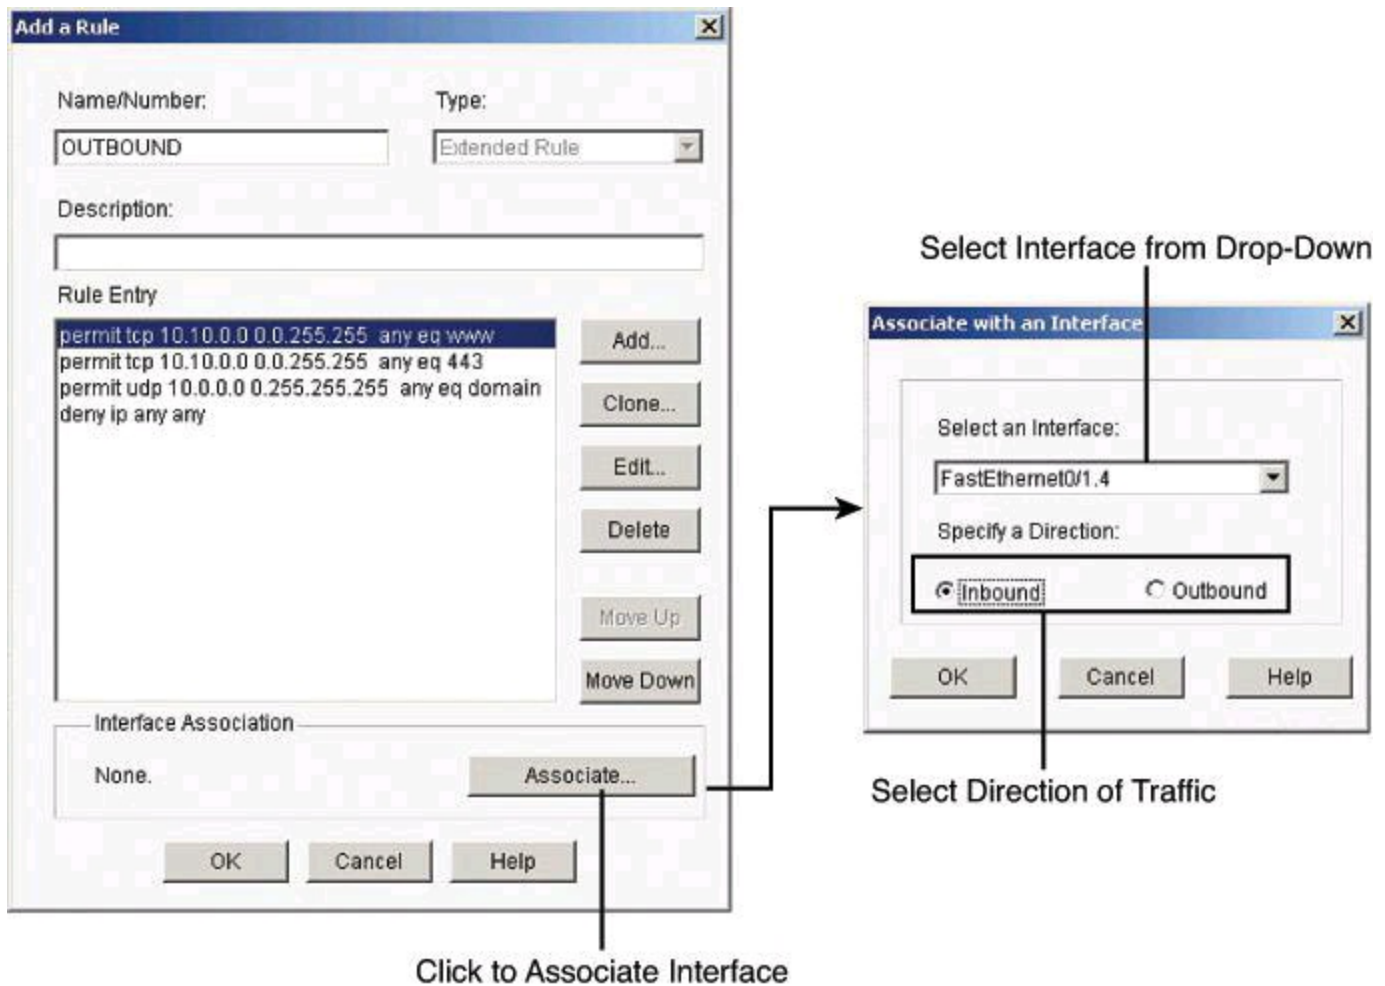


Figure 8-20. Associating Rules with Interfaces

Enabling Logging with CCP

Activity on your access control rules is monitored through the creation of log entries. If logging is enabled on the router, whenever an access rule that is configured to generate log entries is invoked—for example, if a connection is attempted from a denied IP address—then a log entry is generated and can be viewed in Monitor mode.

The first step to viewing firewall activity is to enable logging on the router. To enable logging, follow these steps:

- Step 1.** From the toolbar, choose **Configure > Router > Logging** to open the Logging dialog box, shown in [Figure 8-21](#).

Select Logging Destination and Level

Figure 8-21. Enabling Logging

Step 2. Click **Edit**.

Step 3. In the Syslog window that just opened, check the **Logging to Buffer** check box.

In addition to enabling logging, you must identify the access rules that you want to generate log entries. To configure access rules for generating log entries:

Step 1. From the toolbar, choose **Configure > Router > ACL**.

Step 2. Click **ACL Editor**. Each access rule appears in the upper table on the right side of the screen, as shown in [Figure 8-17](#). The lower table shows the specific source and destination IP addresses and the services that are permitted or denied by the rule.

Step 3. In the upper table, choose the rule that you want to modify and click **Edit**. The Edit a Rule dialog box appears.

Step 4. The Rule Entry field shows each of the source IP/destination IP/service combinations that are permitted or denied by the rule. Click the rule entry that you want to configure to generate log entries, and click **Edit**.

Step 5. In the Edit an Extended Rule Entry dialog box, check the **Log Matches Against this Entry** check box, as shown in [Figure 8-22](#).

Enable Logging for This ACE

Figure 8-22. Selecting ACEs that Will Generate Log Entries

Step 6. Click **OK** to close the dialog boxes that are open. The rule entry that you just modified will now generate log entries whenever a connection is attempted from the IP address range and services that define the rule entry.

Step 7. Repeat Step 4 through Step 6 for each rule entry that you want to configure to generate log entries.

Monitoring ACLs with CCP

Once your logging configuration is complete, choose **Monitor > Router > Logging** and select the **Syslog** tab in the Logging dialog box, as shown in [Figure 8-23](#). The table at the bottom shows each generated router log entry, including the time and the reason that the log entry was generated. For ACL log entries, be sure to select at least Informational from the Select a Logging Level to View drop-down menu. Notice that the table is not automatically refreshed or cleared; you need to click the Update button or the Clear button to update the view. The Search button is also important, because this table shows all syslog messages, not only those related to ACLs from the security perspective.

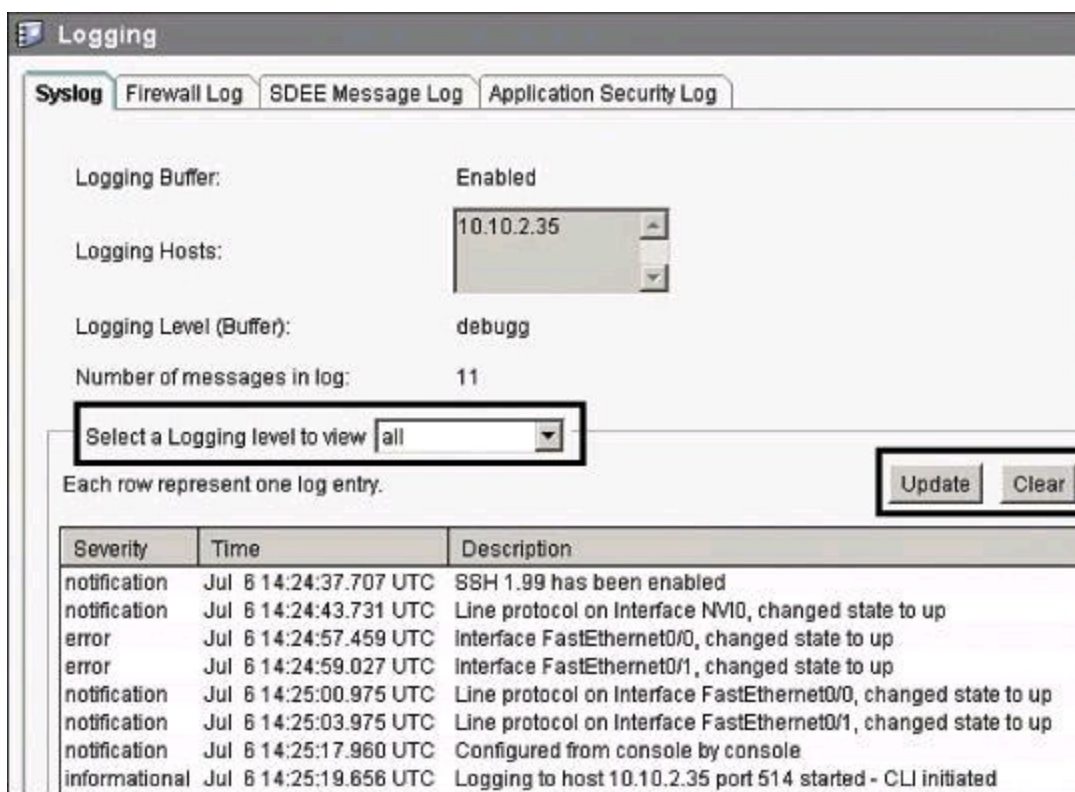
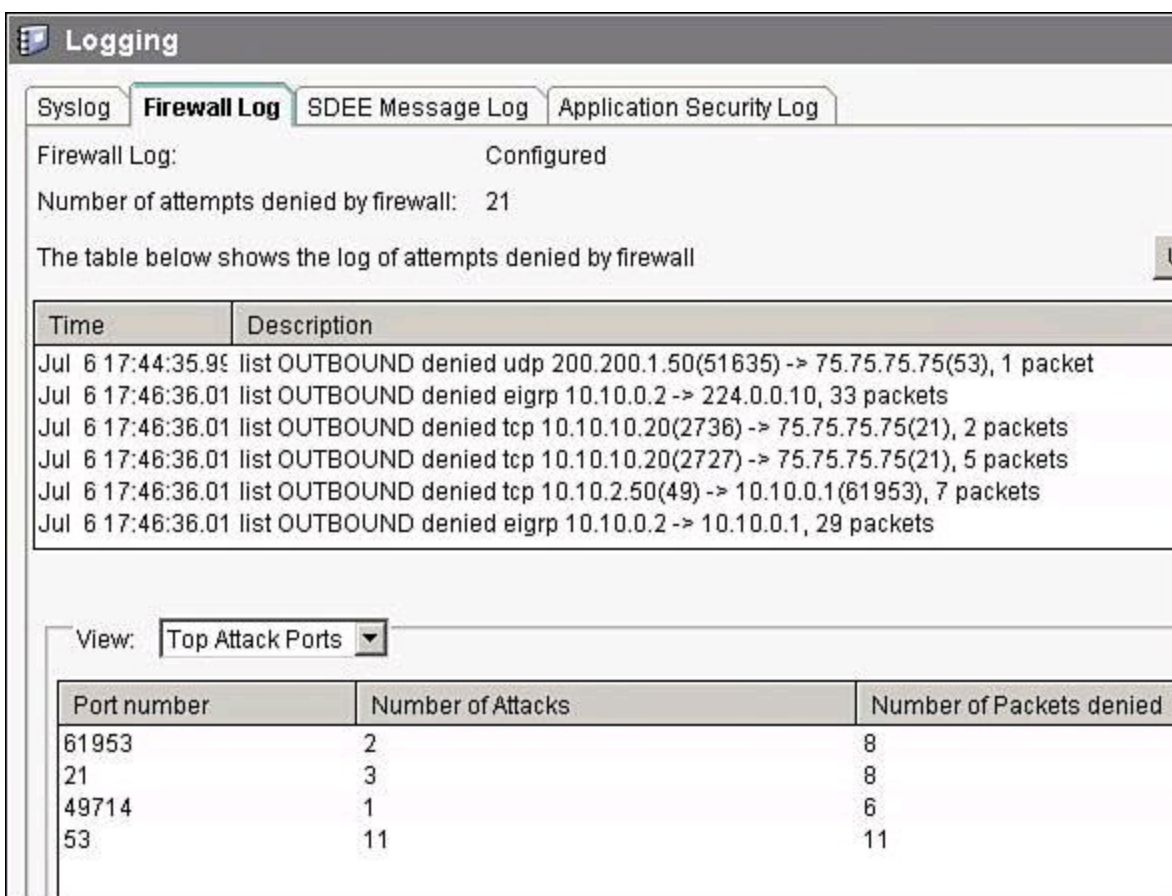


Figure 8-23. Using CCP to View Log Entries

The Firewall Log tab of the Logging dialog box, shown in [Figure 8-24](#), also allows you to see logged ACEs, this time from the security perspective. The log entries shown in the top part of this tab are determined by log messages generated by the firewall. The Number of Attempts Denied by Firewall entry shows the number of connection attempts rejected by the firewall. The top-attacks table below the View drop-down menu displays the top attack entries. For ACLs, traffic matching a deny entry is considered an attack from the perspective of the firewall.



Note

Logged ACLs will display events on the Firewall Log tab even if a zone-based firewall is not configured on the router.

Configuring an Object Group with CCP

You can create two types of ACL object groups:

- **Network object groups:** Can contain hostnames, host IP addresses, subnet masks, range of IP addresses, and other existing network object groups
- **Service object groups:** Can contain top-level protocols, such as TCP, UDP, and TCP-UDP; ICMP types; source and destination protocol ports; and other existing service object groups

[Figure 8-25](#) illustrates an example of the creation of a network object group, the steps for which follow:

Step 1. Choose **Configure > Router > ACL > Object Groups > Network Object Groups** to open the Network Object Group summary page.

Step 2. Click **Create** to open the Create Network Object Group dialog box.

Step 3. Enter the group name and description, specify the parameters in the Network Object Group Members area, and then click **Add**. The parameters that you entered in the left pane are added to the right pane.

Step 4. Click **OK** to send the configured group information to the router.

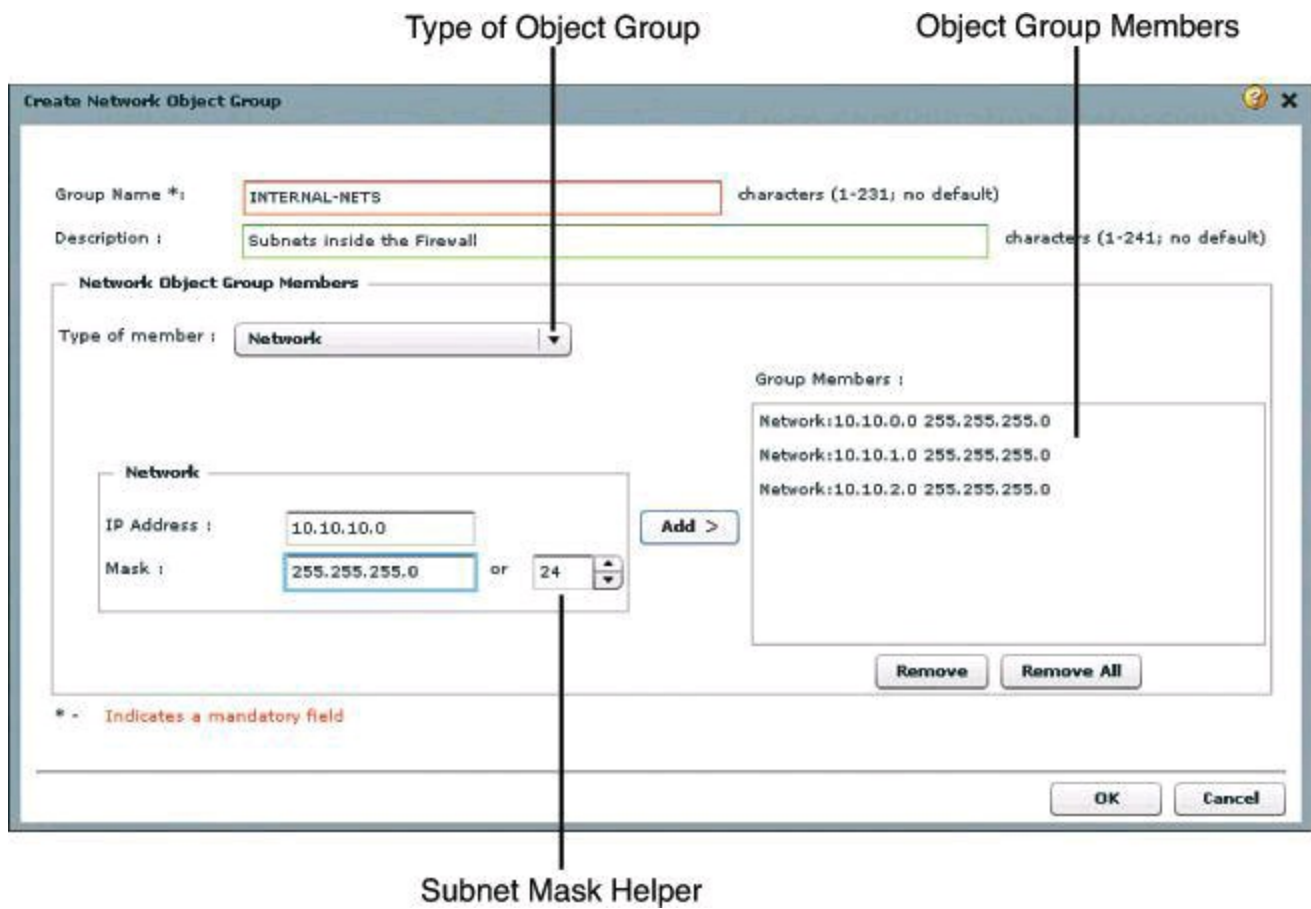


Figure 8-25. Configuring an Object Group Using CCP

[Example 8-8](#) shows the CLI equivalent of the configuration done in [Figure 8-25](#).

Example 8-8. Configuring an Object Group in the CLI

[Click here to view code image](#)

```

Router# config t
Router(config)# object-group network INTERNAL-NETS
Router(config-network-group)# description Subnets inside the Firewall
Router(config-network-group)# 10.10.0.0 255.255.255.0
Router(config-network-group)# 10.10.1.0 255.255.255.0
Router(config-network-group)# 10.10.2.0 255.255.255.0
Router(config-network-group)# 10.10.10.0 255.255.255.0

```

After you create network object groups and/or service object groups, you can create an ACL that can permit or deny traffic to these object groups, as shown in [Figure 8-26](#).

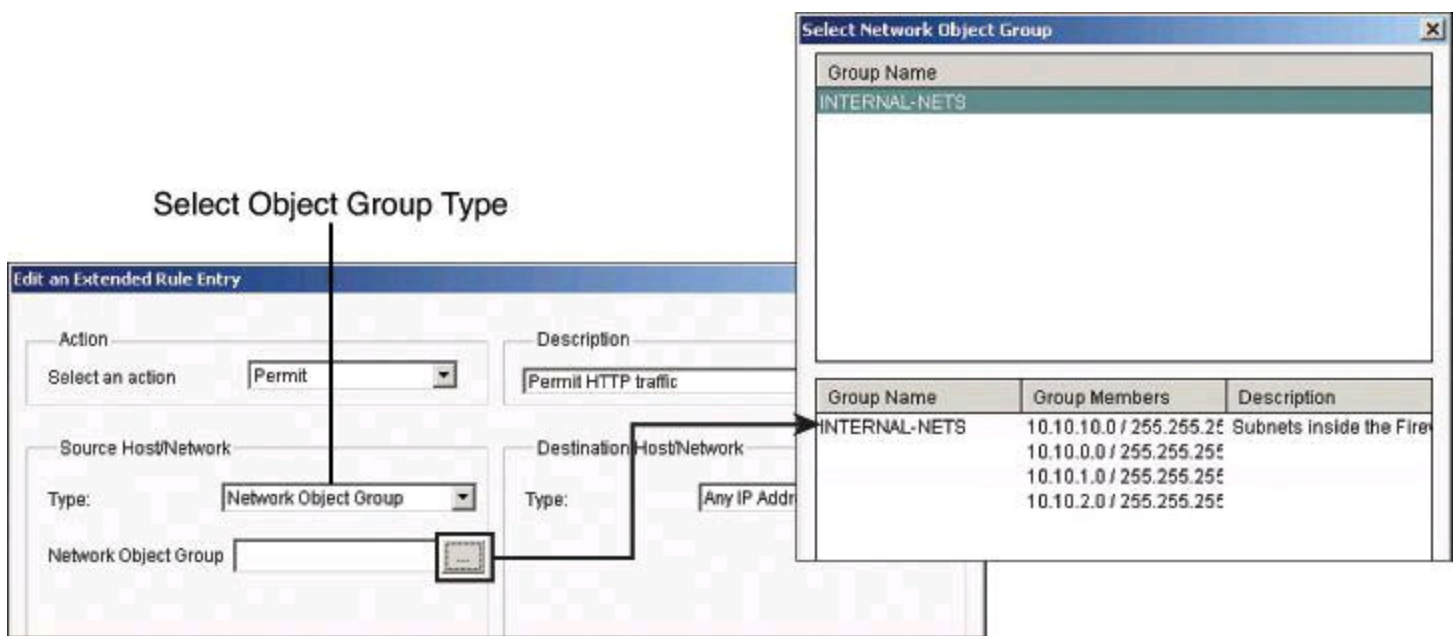


Figure 8-26. Assigning Object Groups to ACLs

Use this procedure to create an ACL that can permit or deny traffic to the configured object groups.

Step 1. Choose **Configure > Router > ACL > ACL Editor** to open the Additional Task Rules dialog box.

Step 2. Click **Add** to open the Add a Rule dialog box.

Step 3. Add a name and description for the rule in the appropriate fields, and then click **Add** to open the Add an Extended Rule Entry dialog box.

Step 4. From the Action field, choose the action you want to configure. The options are Permit and Deny.

Step 5. If you want to use object groups as an ACL source, from the Source Host/Network pane, do the following:

- Choose **Network Object Group** from the Type field.
- Click the **...** (more) button—located beside the Network Object Group field—to open the Select Network Object Group dialog box.
- From the Select Network Object Group dialog box, select the network object group, and then click **OK**.

Step 6. If you want to use an object group as an ACL destination, from the Destination Host/Network pane, do the following:

- Choose **Network Object Group** from the Type field.
- Click the **...** (more) button—located beside the Network Object Group field—to open the Select Network Object Group dialog box.
- From the Select Network Object Group dialog box, select the network object group, and then click **OK**.

Step 7. If you want to use an object group as an ACL protocol and port, from the Protocol and Service pane, do the following:

- Click the **Service Object Group** radio button.

- Click the ... (more) button—located beside the Service Object Group field—to open the Select Service Object Group dialog box.
- From the Select Service Object Group dialog box, select the service object group, and then click **OK**.

Step 8. Click **OK** in the Add an Extended Rule Entry dialog box.

Using ACLs in IPv6 Environments

Packet filtering in IPv6 is similar to packet filtering in IPv4. A strategy to prevent common attacks between the two protocol stacks, such as reconnaissance and spoofing attacks, typically starts with ACLs trying to match malicious traffic.

IPv6-specific attacks, such as the one depicted in [Figure 8-27](#), require targeted filtering and IPv6-aware ACLs. IPv6 ACLs can help mitigate the following threats, among others:

- Header extension threats; for instance, amplification attacks based on Routing Header (RH 0)
- Threats based on misuse and abuse of IPv6 ICMP
- Reconnaissance based on multicast IPv6 addresses
- Threats that exploit tunneling solutions such as those used in IPv6 migration environments

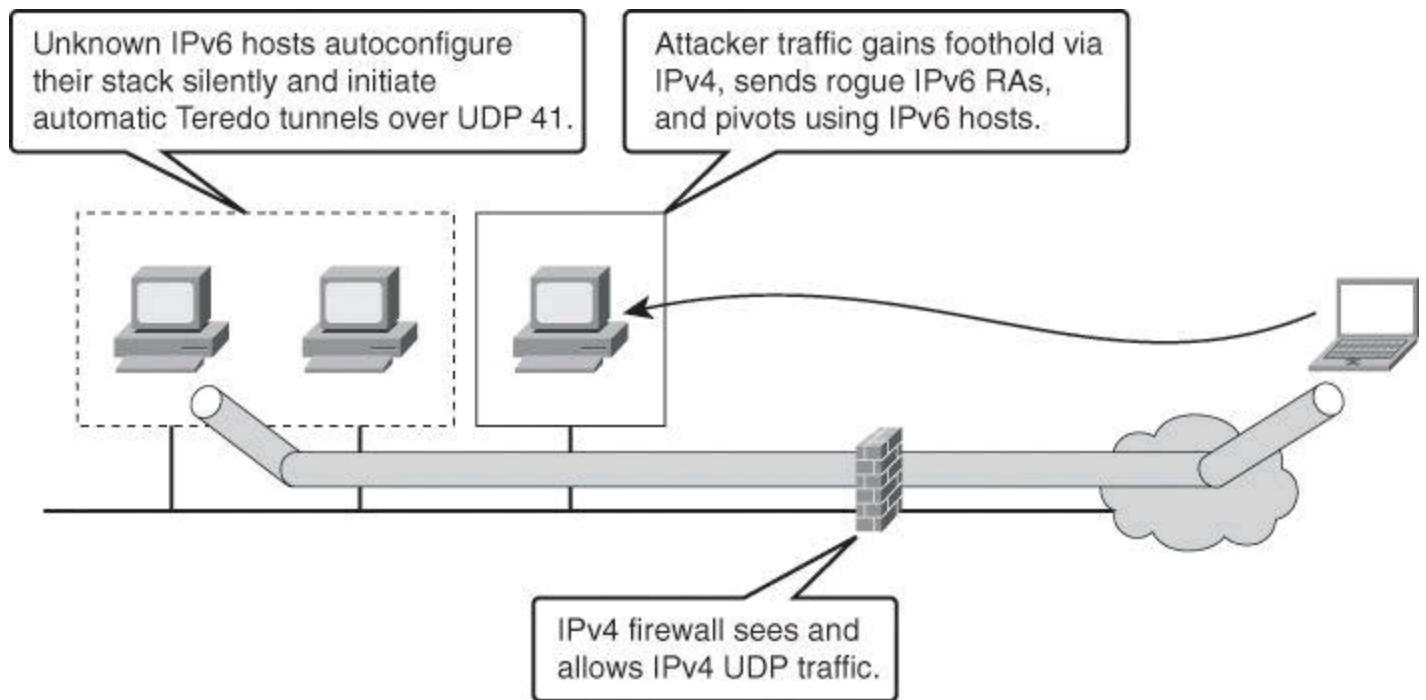


Figure 8-27. Examples of IPv6 Potential Attacks

Some threats require a combined approach using both protocol stacks, specifically those that exploit dual-stack environments. The scenario in [Figure 8-27](#) depicts a Teredo tunnel launched from hosts that have been compromised via IPv6 but use IPv4 tunnels. Infrastructure filtering and firewalling for IPv4 would mitigate and contain the IPv6 threat.

The standard ACL functionality in IPv6 is similar to standard ACLs in IPv4. ACLs determine which traffic is blocked and which traffic is forwarded at router interfaces and allow filtering based on source and destination addresses and some Layer 4 information, inbound and outbound to a specific interface. Each ACL has an implicit **deny** statement at the end. IPv6 ACLs are defined and

their deny and permit conditions are set using the **ipv6 access-list** command with the **deny** and **permit** keywords in global configuration mode. An empty ACL means that traffic is allowed. Reflexive ACLs and time-based ACLs are also available in IPv6. An IPv6 ACL can match the following IPv6 headers:

- **routing:** Matches any route header
- **mobility:** Matches any mobility header
- **dest-option-type:** Matches any destination option header
- **auth:** Matches IPsec's AH
- **undetermined-transport:** Matches any packet whose Layer 4 protocol cannot be determined (fragmented or unknown extension header) (available only with the **deny** command)

Each IPv6 ACL contains implicit **permit** rules to enable IPv6 neighbor discovery, in addition to the traditional rule that implements a default **deny** policy. The resulting implicit rules for IPv6 ACLs are shown in [Example 8-9](#).

Example 8-9. IPv6 ACL Implicit Entries

[Click here to view code image](#)

```
permit icmp any any nd-na
permit icmp any any nd-na
deny ipv6 any any
```

The IPv6 neighbor discovery process makes use of the IPv6 network layer service. Therefore, by default, IPv6 ACLs implicitly allow IPv6 neighbor discovery packets to be sent and received on an interface. Neighbor discovery is similar to ARP for IPv4.

To add some context to these implicit rules, they implement what IPv4 implemented for the Address Resolution Protocol (ARP). In IPv4, ARP, which is equivalent to the IPv6 neighbor discovery process, makes use of a separate data link layer protocol. Therefore, by default, IPv4 ACLs implicitly allow ARP packets to be sent and received on an interface. The same is true for neighbor discovery in IPv6.

The user can override these rules by placing a **deny ipv6 any any** statement within an ACL.

Historically, in IPv4, administrators add the implicit **deny any** rule, along with the **log** keyword, to the end of an ACL. This is effective for monitoring and auditing ACLs, and also for logging denied packets. Doing the same in IPv6 requires adding all three implicit rules. Not doing so would result in neighbor discovery issues. In the example in [Figure 8-28](#), a single **deny ipv6 any any** line is added to log denied packets. The two implicit rules that allow neighbor discovery in [Example 8-9](#) are never evaluated, as all traffic is denied before it reaches the neighbor discovery rules. The solution to this problem is to explicitly add all three implicit lines, in the right order, at the end of your ACLs.

[Figure 8-28](#) shows the topology we will use for the following named IPv6 ACL example.

Prefix: 2001:db8:2c80:1000::/64




Figure 8-28. Example of IPv6 ACLs

The syntax of IPv6 ACLs is similar to that of IPv4 ACLs. However, this time, the ACEs match IPv6 sources, destinations, and ports. Also notice that wildcard masks are no longer used, and the used of the /x nomenclature is pervasive across the syntax. One noticeable difference with IPv4 ACL commands is the use of the **ipv6 traffic filter** command to apply the ACL to interfaces, instead of the traditional **ip access-group** command used in IPv4 ACLs. Proper planning is required in using masks to match summarizable address blocks, given the more complex structure of IPv6 IP addresses.

The following is the syntax of IPv6 ACLs. Refer to Cisco IOS documentation for the detailed information:

[Click here to view code image](#)

```
permit protocol {source-ipv6-prefix/prefix-length | any | host
source-ipv6-address | auth} [operator [port-number]] {destination-
ipv6-prefix/prefix-length | any | host destination-ipv6- address |
auth} [operator [port-number]] [dest-option-type [doh-number | doh-
type]] [dscp value] [flow-label value] [fragments] [log] [log-
input] [mobility] [mobility-type [mh-number | mh-type]] [reflect
name [timeout value]] [routing] [routing-type routing-number]
[sequence value] [time-range name]
```

or

[Click here to view code image](#)

```
permit protocol {source-ipv6-prefix/prefix-length | any | host
source-ipv6-address | auth} [operator [port-number]] {destination-
ipv6-prefix/prefix-length | any | host destination-ipv6- address |
auth} [operator [port-number]] [dest-option-type [doh-number | doh-
type]] [dscp value] [flow-label value] [fragments] [log] [log-
input] [mobility] [mobility-type [mh-number | mh-type]][routing]
[routing-type routing-number] [sequence value] [time-range
name] [undetermined-transport]
```

[Example 8-10](#) illustrates an IPv6 ACL for ICMP that implements RFC 4890, which describes recommendations for filtering ICMPv6 messages in firewalls.

In networks supporting IPv6, ICMPv6 plays a fundamental role with a large number of functions, and a correspondingly large number of message types and options. ICMPv6 is essential to the functioning of IPv6, but there are a number of security risks associated with uncontrolled forwarding of ICMPv6 messages. Filtering strategies designed for the corresponding protocol, ICMP, in IPv4 networks are not directly applicable, because these strategies are intended to accommodate a useful auxiliary protocol that may not be required for correct functioning.

[Example 8-10](#) illustrates a restrictive policy in which the required and necessary ICMPv6 message types are allowed and any unmatched types are denied. This is not a guarantee of protection against ICMPv6-based attacks, but it is a good starting point.

Example 8-10. RFC 4890 ICMP ACL

[Click here to view code image](#)

```
ipv6 access-list RFC4890
permit icmp any any echo-reply permit icmp any any echo-request permit
icmp any any
  1 3
permit icmp any any 1 4
permit icmp any any packet-too-big permit icmp any any time-exceeded
permit icmp any
  any parameter-problem permit icmp any any mld-query
permit icmp any any mld-reduction permit icmp any any mld-report permit
icmp any any
  nd-na
permit icmp any any nd-ns
permit icmp any any router-solicitation
```

Summary

This chapter presented the use cases and benefits of ACLs in general. It described the building blocks and operational framework of ACLs. Prior to digging into ACLs, we reviewed how summarizable address blocks in the context of CIDR and VLSM environments work, and we then evaluated how ACL wildcard masks allow for threat mitigation in those environments.

This chapter also presented design considerations when deploying ACLs in general and demonstrated the use of Cisco Configuration Professional and the CLI to deploy and verify a threat containment strategy using ACLs and to correlate ACL log and alarm information in order to monitor their impact and effectiveness. You learned how to configure object groups to streamline the implementation of ACLs for threat control. This chapter finished by showing you how to configure ACLs in IPv6 environments, highlighting the operational differences with IPv4 ACLs.

References

For additional information, refer to these Cisco.com resources:

“Identifying Incidents Using Firewall and Cisco IOS Router Syslog Events,”
<http://www.cisco.com/web/about/security/intelligence/identify-incident-via-syslog.html>

“IP Access List Entry Sequence Numbering,”
http://www.cisco.com/en/US/docs/ios/12_2s/feature/guide/fsaclseq.html

“Understanding Access Control List Logging,”
<http://www.cisco.com/web/about/security/intelligence/acl-logging.html>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

1. Which type of threat cannot be mitigated using ACLs?
 - a. Malicious payload
 - b. DoS TCP SYN attack
 - c. ICMP message filtering
 - d. Source IP address spoofing
2. What is the range of ACL numbers for a standard access list? (Choose two.)
 - a. 100–199
 - b. 1–99
 - c. 1300–1999
 - d. 0–99
3. Which wildcard mask would match for IP subnets 172.30.16.0/24 to 172.30.31.0/24?
 - a. 0.0.31.255
 - b. 0.0.0.255
 - c. 0.0.255.255
 - d. 0.0.15.255
4. Which of the following ACLs would allow users to access their Post Office Protocol server?
 - a. **access-list 101 permit tcp any any eq 25**
 - b. **access-list 101 permit tcp any eq 25 any**
 - c. **access-list 101 permit tcp any eq 110 any**
 - d. **access-list 101 permit tcp any any eq 110**
5. Which of the following is a caveat regarding ACLs?
 - a. Standard ACLs must be used to implement security policies.
 - b. Beware of the explicit **deny all** at the bottom of ACLs.
 - c. Place less specific statements higher in ACLs.
 - d. Ensure that statements at the top of ACLs do not negate statements found lower in the list.
6. Which of the following are rules that can be configured by Cisco Configuration Professional? (Choose all that apply.)
 - a. NAT rules
 - b. Firewall rules
 - c. Unsupported rules
 - d. Sysopt rules
7. Using which keyword in the **access-list** command is a best practice method for understanding what is happening to traffic as it passes through the firewall?
 - a. **established**
 - b. **eq**

c. log

d. matches

8. Which of the following are benefits of using ACL object groups? (Choose all that apply.)

a. Optimized logging

b. Increased performance

c. Reduced storage

d. Separates ownership of the components of an ACE

9. Using which keyword in the **access-list** command designates that the packet is a response to an outbound initiated session?

a. established

b. eq

c. log

d. matches

10. Which of the following are implicit IPv6 ACLs entries? (Choose all that apply.)

a. permit icmp any any nd-na

b. deny ipv6 any any

c. permit icmp any any nd-ns

d. deny ipv6 icmp any any

Chapter 9. Firewall Fundamentals and Network Address Translation

This chapter teaches firewall concepts, technologies, and design principles. At the end of this chapter, you will be able to do the following:

- Explain the operations of the different types of firewall technologies
- Describe firewall technologies that historically have played, and still play, a role in network access control and security architectures
- Introduce and describe the function and building blocks of Network Address Translation
- List design considerations for firewall deployment
- Describe guidelines for firewall ruleset creation

Introducing Firewall Technologies

A firewall protects network devices from intentional, hostile intrusions that could threaten information assurance (availability, confidentiality, and integrity) or lead to a denial-of-service (DoS) attack. A firewall can protect a hardware device or a software program running on a secure host computer. This chapter introduces the firewall technologies that Cisco uses in routers and security appliances.

Firewall Fundamentals

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

In principle, as shown in [Figure 9-1](#), a firewall is a pair of mechanisms that perform these two separate functions, which are set by policies:

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

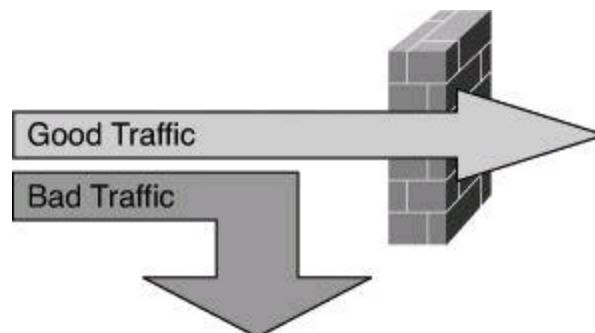


Figure 9-1. Firewall: Enforcing Access Control

A firewall can be defined as follows:

A system or group of systems that enforces an access control policy between two networks.

Because this definition is very generic, almost anything can be considered a firewall. Many network access technologies can be used to build a firewall:

- Packet-filtering routers
- LAN switches
- Complex systems integrating many hosts into a firewall system

Note

Depending on which definition you look at regarding firewall technologies, proxy servers and gateways may or may not be included. Proxy gateways, discussed later in this chapter, are also referred to as application layer firewalls. The Cisco IronPort Web Security Appliance is an example of a proxy server.

Firewalls mean different things to different organizations, and each organization has unique requirements. Nevertheless, all firewalls usually share some common properties:

- **Must be resistant to attacks:** Compromise of the firewall system should be very unlikely, because it would enable an attacker to disable the firewall or change its access rules.
- **Must be the only transit point between networks:** All traffic between networks must flow through the firewall. This requirement prevents a hacker from using a backdoor connection to bypass the firewall and violate the network access policy.
- **Enforces the access control policy of an organization:** The access control policy should define what the firewall permits or denies.

By performing network access control, you can use a firewall as a protective measure against the following:

- **Exposure of sensitive hosts and applications to untrusted users:** A firewall hides most of the functionality of a host and permits only the minimum required connectivity to a host. Complexity is thus reduced, and many possible vulnerabilities are not exposed.
- **Exploitation of protocol flaws:** You can program a firewall to inspect protocol messages and verify their compliance with the protocol, whether it is Layer 3, Layer 4, or a higher-layer protocol. The firewall limits what attackers can send to their target, preventing the delivery of malformed packets that are used in an attempt either to crash a system or to gain access to an application.
- **Malicious data:** A firewall can detect and block malicious data sent to clients or servers inside the application stream, thereby stopping it from infecting the server or the client. Because firewalls are located on critical interconnection points of the network, enforcing the network access policies is simple, scalable, and robust. Sometimes a small number of firewalls can handle most of the network access control needs of an organization.

Firewalls are often misunderstood, and false assumptions are often made about their capabilities.

Although it is true that firewalls would not be necessary if host and application security could be made extremely robust, many organizations use firewalls as a replacement for host or application security. Such an attitude is extremely dangerous because it can completely ignore host and application security even in extreme cases, such as connecting a sensitive server inside an Internet firewall.

Today, firewalls are such a mainstream technology that they are often considered a panacea for many security issues. While you should be aware of the benefits of the firewall model, you should also be aware of the many limitations that firewalls have and how to mitigate some of these limitations:

- Because firewalls are used in critical points of the network, their misconfiguration can have disastrous consequences. Firewalls are often a single point of failure when it comes to security; a single mistake in a configuration rule or firewall code can compromise the network access policy.
- Many modern applications require that the firewall handle, in addition to the primary control connection, secondary connections that are created to carry, for example, data. A typical example is an application that opens dynamic sessions from the outside to the inside after an initial client request initiated from the inside to the outside. Multimedia applications such as those for audio streaming and videoconferencing are examples where the user opens one session from the inside to the outside to request a feed. To support the streaming, however, additional sessions are opened from the outside to the inside, and by default the firewall rejects those new incoming requests. Once firewall vendors have a chance to study the new protocol, they can create a rule that forces the firewall to peek inside the payload of the original outgoing packet to gain information about the additional sessions that will be created and to prepare for those new incoming sessions.
- End users, when faced with a restrictive firewall, might find their own methods of bypassing it. For example, inside users may use wireless broadband from the protected network to an Internet service provider (ISP) and create a backdoor connection to the protected network.
- Because firewalls are commonly placed at chokepoints, the design of the firewall model must ensure network performance is not affected by inspecting all the traffic.
- Tunneling unauthorized data over authorized connections (covert channels) is simple and generally impossible to detect. This activity usually requires the help of someone on the trusted side of the firewall.

Firewalls in a Layered Defense Strategy

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

A layered defense strategy combines different types of firewalls in layers to add depth to the information defense of an organization, as shown in [Figure 9-2](#). For example, traffic that comes in from the untrusted network first encounters a packet filter on the outer router, located in the perimeter security. The traffic next goes to the screened host firewall or bastion host system, which applies

more rules to the traffic and discards suspect packets. This bastion host system is also part of the perimeter security. The traffic then goes to an interior screening router, part of the core network security. The traffic moves to the internal destination host only after this routing. This type of demilitarized zone (DMZ) setup, located between the perimeter and the core, is called a *screened subnet configuration*. Endpoint traffic moves in the network after it has been properly authenticated by sophisticated Identity services.

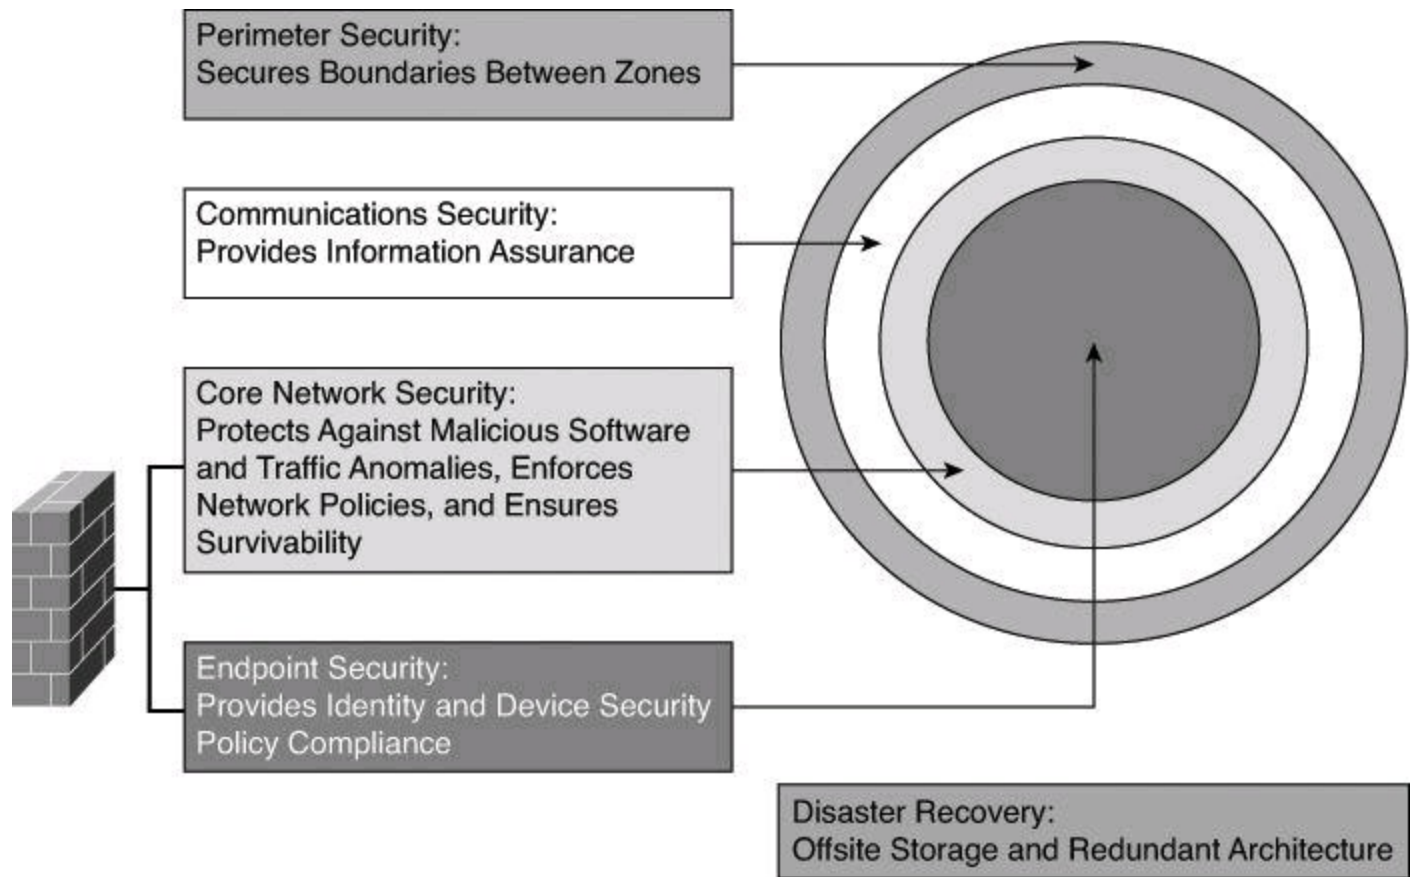


Figure 9-2. Layered Defense Strategy

Two recurring terms in this chapter are bastion host and packet filter, which are defined as the following:



- **Bastion host:** A computer that is expected to be attacked and therefore is hardened. An example of a bastion host is firewall software installed on a workstation that is already running a commonly available operating system. Before the firewall is put into production, the workstation needs to be hardened to protect against the potential vulnerability that the operating system has.
- **Packet filter:** Typically, a router configured with access lists used to filter out unwanted traffic.

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

- A significant number of intrusions come from hosts within the network. For example, firewalls often do little to protect against viruses downloaded through email.
- Firewalls do not protect against rogue modem or rogue wireless access point installations. In addition, and most important, a firewall is not a substitute for informed administrators and users.
- Firewalls do not replace backup and disaster recovery mechanisms resulting from attack or hardware failure. An in-depth defense also includes offsite storage and redundant hardware topologies.

Defense in depth and diversity of defense are related topics. Defense in depth calls for multiple levels of defense; diversity of defense calls for using different types of technologies in that defense.

Key Topic

An example of implementing both diversity of defense and defense in depth is using a perimeter router as a packet filter and using a stateful firewall to segment the unprotected segment from the protected segment.

Static Packet-Filtering Firewalls

Packet-filtering firewalls work primarily at the network layer of the Open Systems Interconnection (OSI) model, or the IP layer of TCP/IP, as shown in [Figure 9-3](#). Packet-filtering firewalls are generally considered Layer 3 devices, but they typically have the capability to permit or deny traffic based on Layer 4 information, such as protocol, and source and destination port numbers, in addition to the Layer 3 source and destination IP address. Packet filtering uses rules and ACLs to determine whether to permit or deny traffic based on source and destination IP addresses, protocol, source and destination port numbers, and packet type. Packet-filtering firewalls are usually part of a router firewall.

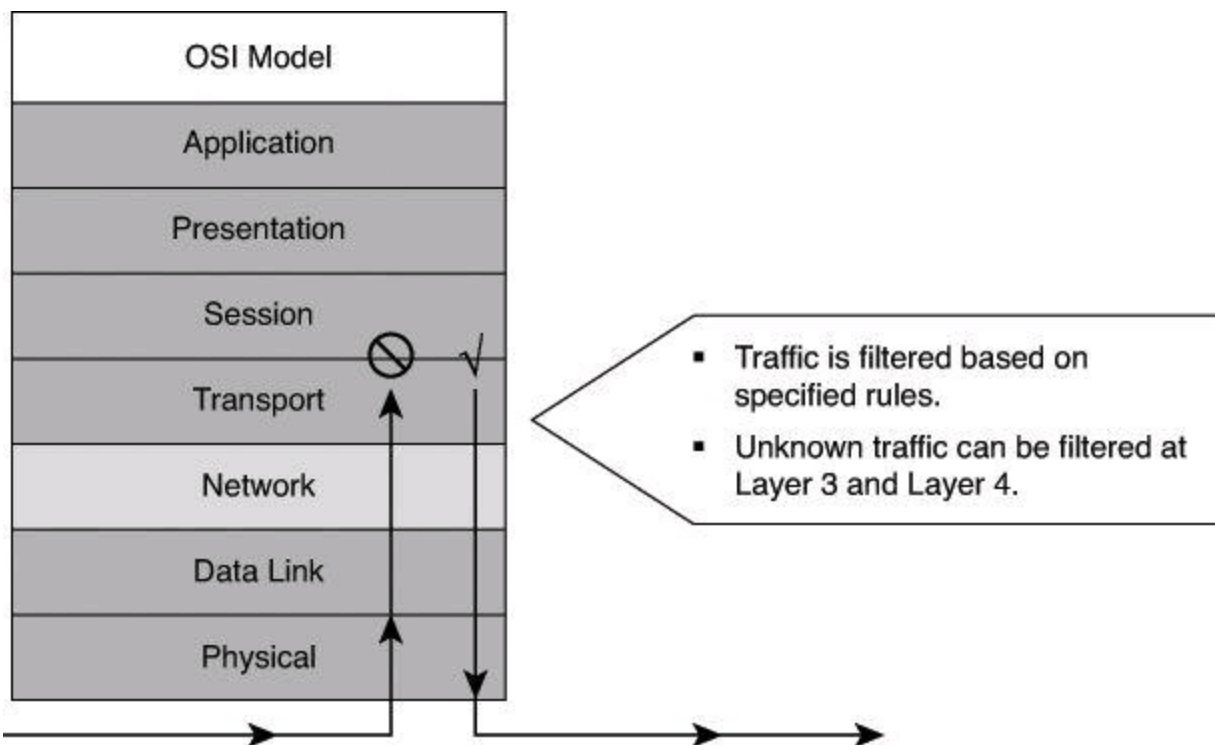


Figure 9-3. How Static Packet Filters Map to the OSI Model

Recall that services rely on specific ports to function; for example, Simple Mail Transfer Protocol (SMTP) servers listen to port 25 by default. Because packet-filtering firewalls filter traffic according to static packet header information, they are sometimes referred to as *static filters*. By restricting certain ports, you can restrict the services that rely on those ports. For example, blocking port 25 on a specific workstation prevents an infected workstation from broadcasting email viruses across the Internet.

Packet-filtering firewalls are similar to packet-filtering routers but with some differences in implementation. Packet filters are very scalable, application independent, and have high performance standards; however, they do not offer the complete range of security solutions required in modern networks. For example, a packet filter does not have the capability to understand dynamic protocols upon which a client request requires additional incoming connections to be initiated from the outside, toward the inside client.

Any device that uses ACLs can perform packet filtering. Cisco IOS router configurations commonly use ACLs, not only as packet-filtering firewalls but also to select specified types of traffic that is to be analyzed, forwarded, or influenced in some way. Later in this chapter, you will see examples of both ingress and egress filtering done with ACLs.

[Figure 9-4](#) shows a simple packet-filtering example using a Cisco router.

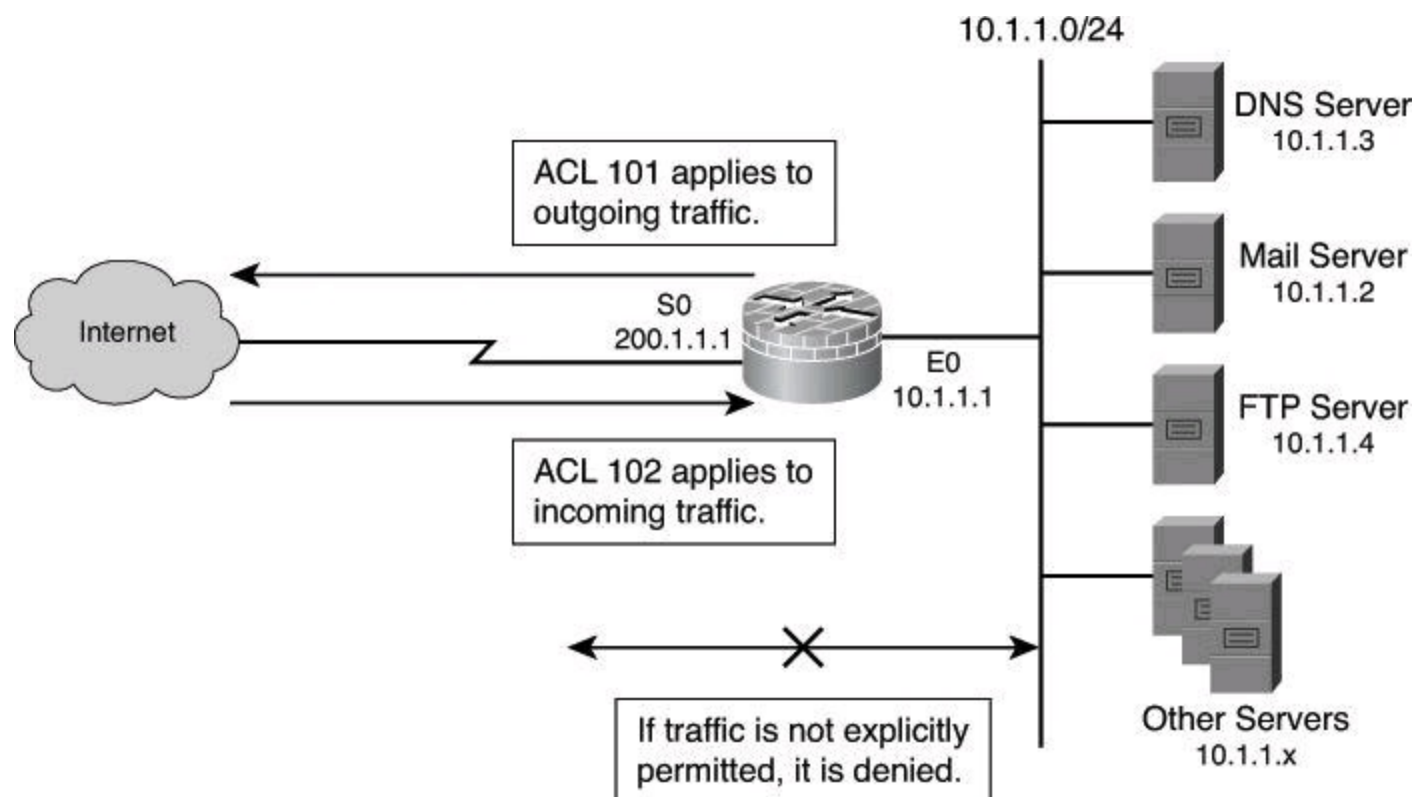


Figure 9-4. Static Packet Filter in Action

In most network topologies, you need to protect the Ethernet interface connecting to the internal (inside) network, while the serial interface that connects to the Internet (outside) is unprotected. In [Figure 9-4](#), the internal addresses that the firewall must protect are in the 10.1.1.0/24 subnet (on the Ethernet interface). The IP address of the Ethernet 0 interface is 10.1.1.1/24.

The particular network security policy shown in [Figure 9-4](#) (ACL 101) allows all users from the

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

Packet-filter firewalls (or packet filters) use a simple policy table lookup that permits or denies traffic based on the following possible criteria:

- Source IP address
- Destination IP address
- Source port number
- Destination port numbers
- Synchronize/start (SYN) packet receipt

The firewalls are extremely fast because they do little computation. The rules are extremely easy to implement because they require little security expertise. Router manufacturers easily embed packet-filtering logic in silicon and, consequently, packet filtering is a feature of most routers. Packet-filtering firewalls are relatively inexpensive. Even if other firewalls are used, implementing packet filtering at the router level affords an initial degree of security at the network and transport layers.

Packet filters do not represent a complete firewall solution. However, they are a key element of a secure architecture.

The following are disadvantages of packet filters:

- Packet filtering is susceptible to IP spoofing. Hackers send arbitrary packets that fit ACL criteria and pass through the filter.
- Packet filters do not filter fragmented packets well. Because fragmented IP packets carry the Layer 4 header in the first fragment and packet filters may filter based on information found in the Layer 4 header, fragments after the first fragment are passed unconditionally. A decision to use packet filters assumes that the filter of the first fragment accurately enforces the policy.
- Complex ACLs are difficult to implement and maintain correctly.
- Packet filters cannot dynamically filter certain services. For example, sessions that use dynamic port negotiations are difficult to filter without opening access to a whole range of ports.
- Packet filters are stateless. They examine each packet individually rather than in the context of the state of a connection.

Application Layer Gateways

Application layer firewalls, also called proxy firewalls or application gateways, provide a higher level of security than packet-filtering firewalls because they allow the greatest level of control. Application gateways operate on Layers 3, 4, 5, and 7 of the OSI model, as shown in [Figure 9-5](#).

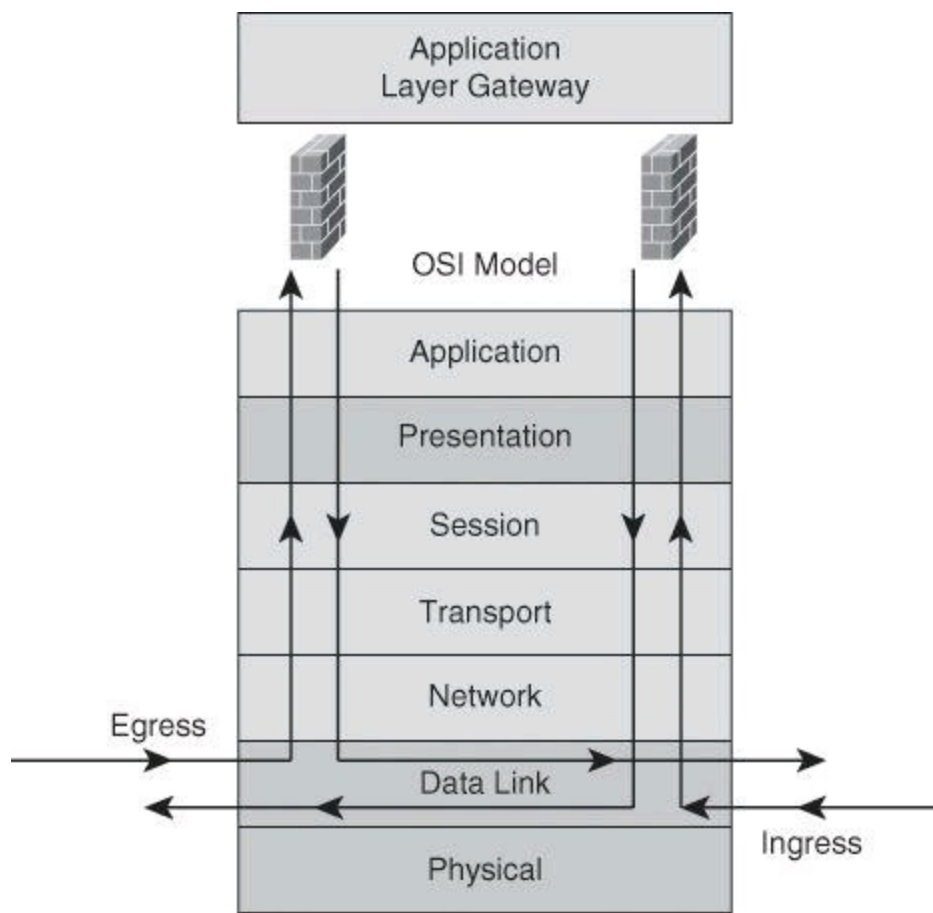


Figure 9-5. Application Layer Gateway in Action

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

Proxy firewalls act as intermediaries between networks to determine whether to allow the communication to proceed. No direct connection exists between an outside user and internal network resources, because the original connection stops in the proxy and a new connection is set between the proxy and the outside destination. For this reason, the only IP address of the network that is visible from the Internet is the IP address of the outside interface of the proxy. The client connects to the proxy server and submits an application layer request. The application layer request includes the true destination and the data request itself. The proxy server analyzes the request and can filter or change the packet contents. The server makes a copy of each incoming packet, changes the source address, and sends the packet to the destination address. The destination server replies to the proxy server, and the proxy server passes the response back to the client.

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

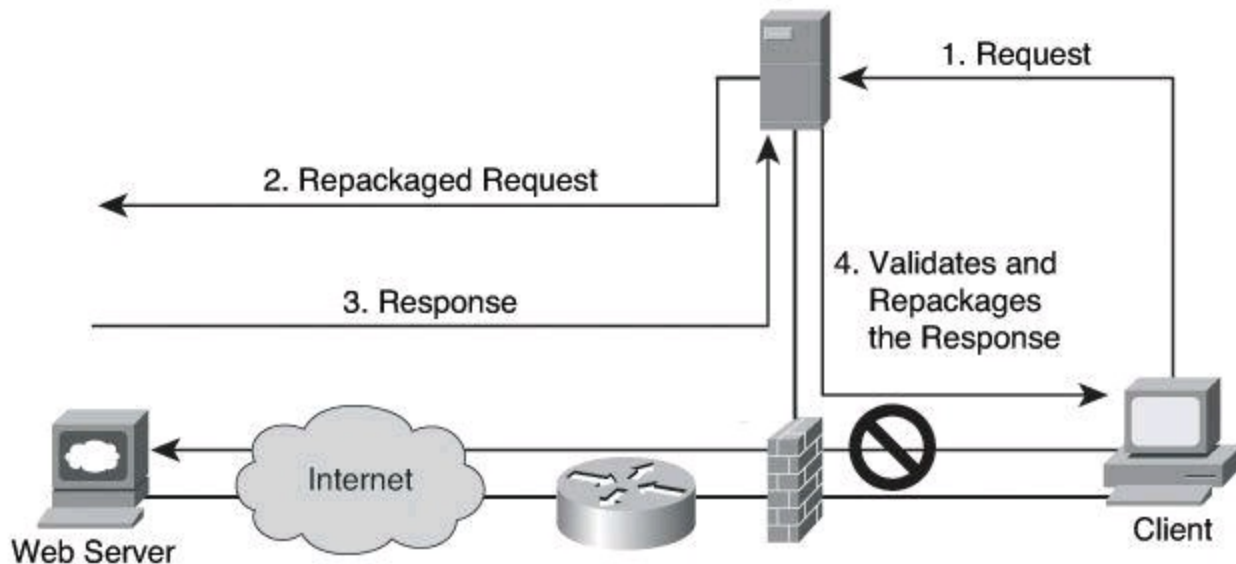
Application layer firewalls provide several advantages:

- **Application layer firewalls authenticate individuals, not devices:** These firewalls typically allow you to authenticate connection requests before allowing traffic to an internal or external resource. This process enables you to authenticate the user requesting the connection instead of authenticating the device.
- **Application layer firewalls make it is harder for hackers to spoof and implement DoS attacks:** An application layer firewall enables you to prevent most spoofing attacks, and DoS attacks are limited to the application firewall itself. The application firewall can detect DoS attacks, and thus reduce the burden on your internal resources.
- **Application layer firewalls can monitor and filter application data:** You can monitor all data on a connection, so you can detect application attacks such as malformed URLs, buffer-overflow attempts, unauthorized access, and more. You can even control the commands or functions you allow an individual to perform based on the authentication and authorization information.
- **Application layer firewalls can provide detailed logging:** Using application layer firewalls, you can generate detailed logs and monitor the actual data that the individual is sending across a connection. This logging can prove extremely useful if a hacker finds a new type of attack, because you can monitor what the hacker does and how the machine does it, and then address the attack. Besides using logging for security purposes, you can use it for management purposes by keeping track of who is accessing what resources, how much bandwidth is used, and how often a user accesses the resources.

The topology in [Figure 9-6](#) represents a typical proxy server deployment. A client inside the network is requesting access to a website. The client browser uses a proxy server for all HTTP requests. Perfect examples of proxy gateways performing application layer filtering are the Cisco IronPort Web Security Appliance (WSA) and the Cisco IronPort Email Security Appliance (ESA). Network security policies force all client connections to go through the proxy server. As shown in [Figure 9-6](#), the browser connects to the proxy server to make requests. Client-side DNS queries and client-side routing to the Internet are not needed when using a proxy server. The client has to reach only the proxy server to make the request.

Cisco IronPort Web Security Appliance (WSA)

Proxy Server:
Dedicated Application Layer
Gateway for HTTP



1. A client's browser forwards its web request to the IronPort WSA.
2. After checking the policies to confirm that the client's request was acceptable, the WSA initiates a connection to the web server. In the process, its session goes through the perimeter firewall, which is most likely set up to permit outbound HTTP traffic only if the source IP address is that of the WSA.
3. The web server sends its reply addressed to the WSA.
4. The WSA performs policy checks, such as malware detection. Acceptable responses are passed to the client.

Figure 9-6. Proxy Server Communication Process

When the proxy server receives the request from a client, it performs user authentication according to the rules applied to it and uses its Internet connection to access the requested website. It forwards only packets that match the firewall rules. On the return route, the proxy server analyzes the packet, including the Layer 5 and Layer 7 header and payload, to ensure that the server allows the content of the reply back in (as an example, checking whether the payload carries hidden malware) before forwarding the packet to the client.

In spite of how an application layer firewall works, with its thorough inspection of the request and response, this firewall provides only a limited number of services, such as HTTP and FTP; however, this type of firewall provides the highest level of filtering for those specific protocols.

The main limitation of application layer firewalls is that they are process intensive because the server evaluates a significant amount of information embedded in many packets. This type of technology requires many CPU cycles and a lot of memory to process every packet that needs inspection, which sometimes creates throughput problems. In addition, the detailed logging can create disk space problems. To address these issues, you can use one of two solutions:

- **Employ a context transfer protocol:** Using a context transfer protocol, where identity-specific information tracks users, enables you to perform only authentication and authorization; you cannot monitor data on the connection, only whether the user is

authorized to go on the Internet. This solution is not a real firewall per se, because packets are checked not for their content but for the validity of their source and destination addresses. This solution would be similar to a packet filter that has the capability to learn the source and destination information dynamically and match return traffic based on that information.

- **Monitor only key applications:** With this solution, you limit the application layer firewall to process only certain application types (such as email, Telnet, FTP, or web services) and then, perhaps, process only connections to specific internal resources. The problem with this approach is that you are not monitoring all applications and connections, and this creates a security weakness.

Application layer firewalls typically do not support all applications, such as multimedia or peer-to-peer file sharing applications (to name a few). Instead, they are generally limited to one or a few connection types, typically common applications such as email, Telnet, FTP, and web services. Therefore, an application layer firewall cannot monitor data on all applications: it monitors data only on applications it intrinsically understands.

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

Dynamic or Stateful Packet-Filtering Firewalls

Stateful packet filters, or stateful firewalls, are the most versatile and therefore the most common firewall technologies in use. Stateful filtering provides dynamic packet-filtering capabilities to firewalls. Stateful inspection is a firewall architecture that is classified at the network layer, although for some applications it can analyze traffic at Layer 4 and Layer 5, too, as shown in [Figure 9-7](#). Some stateful firewalls can analyze traffic up to Layer 7 under special circumstances and additional configuration.

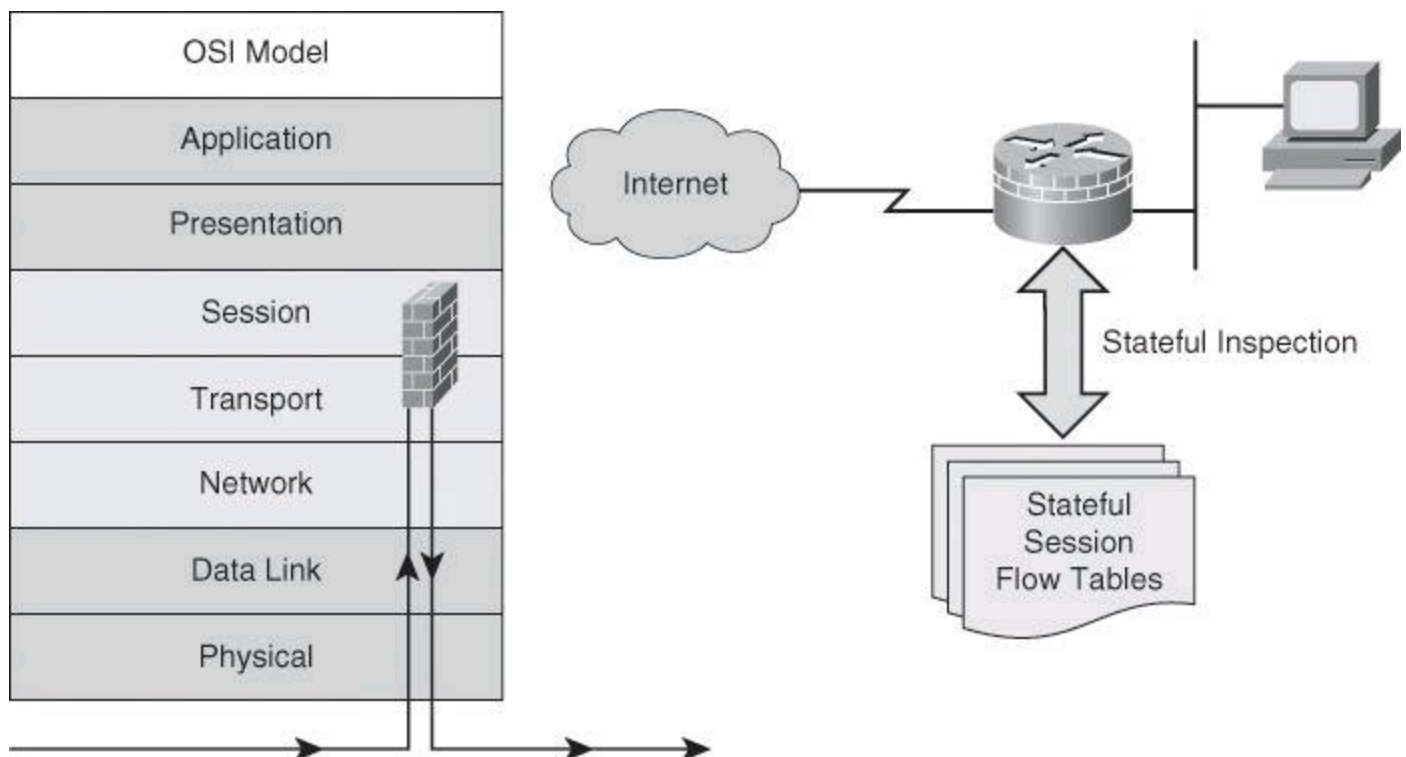


Figure 9-7. Stateful Packet-Filtering Firewall

Unlike static packet filtering, which examines a packet based on the information in its header, stateful inspection tracks each connection traversing all interfaces of the firewall and confirms that the session is valid. Stateful packet filtering maintains a state table and allows modification of the security rules on-the-fly. The state table is part of the internal structure of the firewall and tracks all sessions and inspects all packets passing through the firewall. If packets have the expected properties that are predicted by the state table, the firewall allows them to pass. The state table changes dynamically according to traffic flow.

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

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

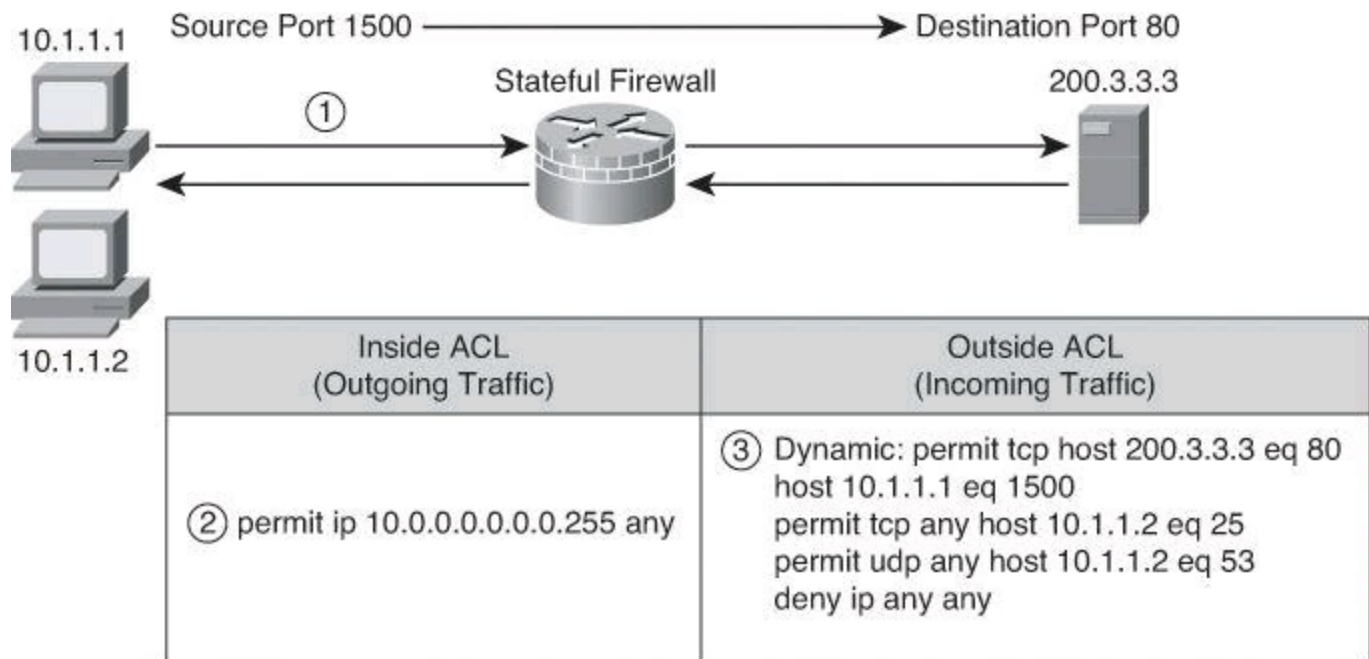
The stateful session flow table contains the source and destination addresses, port numbers, UDP connection information and TCP sequencing information, and additional flags for each TCP connection associated with a particular session. This information creates a connection object against which the firewall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

More advanced stateful firewalls include the capability to parse FTP port commands and update the state table to allow FTP to work transparently through the firewall. Advanced stateful firewalls can also provide TCP sequence number randomization, and DNS query and response matching to ensure that the firewall allows packets to return only in response to queries that originate from inside the network. These features reduce the threat of TCP RST flood attacks and DNS cache poisoning. Some stateful firewalls can also check the validity of protocol commands to ensure that intrusive and dangerous commands are not admitted on the network.

Packets inside the network must make their way to the outside network. This can possibly expose internal IP addresses to potential hackers. That is why most firewalls incorporate Network Address Translation (NAT) with stateful inspection and proxy servers for added security.

Stateful firewalls keep track of the state of a connection: they know whether the connection is in an initiation, data transfer, or termination state. This information is useful when you want to deny the initiation of connections from external devices (default behavior of a firewall) but allow your inside users to establish connections to these external devices and permit the responses to come back through the stateful firewall.

[Figure 9-8](#) shows a successfully established HTTP TCP session that leads to a dynamic ACL rule entry on the outside interface that permits response packets from the web server to the client.



The outside ACL has three permanent entries permitting outside traffic to initiate a connection on 10.1.1.2. Our firewall also has an ACL on its inside interface, permitting traffic from inside users to anywhere outside.

1. An inside user initiates a connection.
2. The user's connection is tested against the inbound ACL on the inside interface.
3. If the traffic is permitted to progress from inside to outside, the firewall inserts a dynamic entry in the inbound ACL of the outside interface to allow the reply traffic to come in to reach the inside user.

Figure 9-8. Stateful Packet Filtering

Stateful packet-filtering firewalls are good to use for the following applications:

- **As a primary means of defense:** In most situations, a stateful firewall is used as a primary means of defense by filtering unwanted, unnecessary, or undesirable traffic.
- **As an intelligent first line of defense:** Networks use routing devices that support a stateful function as a first line of defense or as an additional security boost on perimeter routers.
- **As a means of strengthening packet filtering:** Stateful filtering provides more stringent control over security than packet filtering does, without adding too much cost.
- **To improve routing performance:** Stateful packet-filtering devices perform better than packet filters or proxy servers. Stateful firewalls do not require a large range of port numbers to allow returning traffic back into the network, providing that the firewall is familiar with the protocol and its behavior. The state table determines whether a packet is returning traffic. If it is not returning traffic, the filtering table filters the traffic.
- **As a defense against spoofing and DoS attacks:** Stateful packet filtering works on packets and connections. In particular, stateful firewalls track the state of the connection in the state table listing every connection or connectionless transaction. By determining whether packets belong to an existing connection or are from an unauthorized source, stateful firewalls allow only traffic from connections that are listed in the table. As an example, during the three-way handshake of a TCP session, the firewall tracks the flag and

therefore can predict what the following exchange between the client and the server will be. When the firewall removes a connection from the state table (for instance, because of the connection termination following the TCP four-way handshake as a goodbye), the firewall does not allow any more traffic from that device. In addition, the stateful firewall can log more information than a packet-filtering firewall can, including when a connection was set up, how long it was up, and when it was torn down. This logging makes connections harder to spoof.

Stateful firewalls have the following limitations:

- **Stateful firewalls cannot prevent application layer attacks:** For example, your network might allow traffic to port 80 on a web server. Your stateful firewall examines the destination address in the Layer 3 packet and the destination port number in the segment. If there is a match, the stateful firewall allows the incoming and outgoing traffic. One problem with this approach is that the stateful firewall does not examine the actual contents of the HTTP connection.
- **Not all protocols have a state:** UDP and Internet Control Message Protocol (ICMP) do not have a state, so a stateful firewall does not know whether the session is initiating, transmitting, or finishing, as it would with TCP. For example, UDP has no defined process for how to set up, maintain, and tear down a connection. The firewall can provide only limited support for these protocols compared to the thorough information the firewall has regarding, as an example, a TCP connection. Routers define UDP connections on an application-by-application basis.
- **Some applications open multiple connections:** On earlier (more basic) stateful firewalls, if the client was inside the network and the server was outside the network, both stateful and packet-filtering firewalls had problems dealing with the data connection that the FTP server established to the client. You needed to open a whole range of ports to allow this second connection. Cisco IOS Firewall does not have this problem. The firewall monitors FTP control traffic and discovers information exchanged between the client and the server, which it uses to prepare the data path established from the server in the direction of the client. In this scenario, the firewall is not filtering the traffic at the application layer; it is only reading some purposeful information to prepare for the data. Here we are not talking about application inspection, discussed earlier, where the firewall reads the payload of Layer 7 to assess the FTP commands or some strings found in the FTP payload.
- **Stateful firewalls do not authenticate users by default:** Stateful firewall technology itself does not mandate user authentication. Add-on functionality is used to provide a feature such as proxy authentication.

Other Types of Firewalls

Over the years, variations of standard stateful firewalls have emerged. Some examples of those variations, which provide additional or restrictive features, are deep packet inspection (DPI) firewalls and Layer 2 firewalls.

Application Inspection Firewalls, aka Deep Packet Inspection

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

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

The application inspection function also monitors sessions to determine the port numbers for secondary channels. Many protocols open secondary TCP or UDP ports. The initial session on a well-known port negotiates dynamically assigned port numbers. The application inspection function monitors these sessions, identifies the dynamic port assignments, and permits data exchange on these ports for the duration of the specific session.

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

- **Transport layer mechanism:** From a transport layer perspective, the application inspection firewall acts like a stateful firewall by examining information in the headers of Layer 3 packets and Layer 4 segments. For example, the application inspection firewall looks at the TCP header for SYN, RST, ACK, FIN, and other control codes to determine the state of the connection.
- **Application layer mechanism:** The application inspection firewall checks the conformity of commands within a known protocol. For example, when the application inspection firewall checks the SMTP message type, it allows only acceptable message types on Layer 7 (such as, DATA, HELO, MAIL, NOOP, QUIT, RCPT, and RSET). In addition, the application inspection firewall checks whether the command attributes that are used (for example, length of a message type) conform to the internal rules. These rules often trust the RFC of a specific protocol. Sometimes, application layer firewalls provide protocol support for HTTP, and the application inspection firewall can determine whether the content is really an HTML website or a tunneled application, such as Instant Messaging or GoToMyPC. In the case of a tunneled application, the application inspection firewall would block the content or terminate the connection. Future development will provide application inspection support for more protocols on an application inspection firewall.

There are several advantages of an application inspection firewall:

- Application inspection firewalls are aware of the state of Layer 4 and Layer 5 connections.
- Application inspection firewalls check the conformity of application commands.
- Application inspection firewalls have the capability to check and affect Layer 7.
- Application inspection firewalls can prevent more kinds of attacks than stateful firewalls can. For example, application inspection firewalls can stop an attacker from trying to set up a virtual private network (VPN) tunnel (triggered from inside the network) through an application firewall by way of tunneled HTTP requests.

Transparent Firewalls (Layer 2 Firewalls)

Cisco IOS routers, Cisco ASA Adaptive Security Appliance Software, Cisco Firewall Services Module, and Cisco ASA Services Module offer the capability to deploy a security appliance in a secure bridging mode as a Layer 2 device to provide rich Layer 2 through 7 security services for the protected network, as illustrated in [Figure 9-9](#). This capability enables businesses to deploy security appliances into existing network environments without the need to readdress the network. Although the security appliance can be invisible to devices on both sides of a protected network, as shown in [Figure 9-9](#), administrators can use an exposed IP address to manage the appliance.




Figure 9-9. Transparent Firewalling: Firewall Interfaces All in the Same Subnet

Note

Layer 2 firewalls, also known as transparent firewalls, are sometimes referred to as bumps in the wire or as stealth firewalls.

Tip

Additional information on transparent firewalls is available in *CCNP Security FIREWALL 642-618 Official Cert Guide* (Cisco Press).

NAT Fundamentals

As you read earlier in this book, IPv6 will not require NAT. However, at the time of this writing, more than 99 percent of Internet traffic is still running on IPv4. NAT enables private IPv4 internetworks that use nonregistered IPv4 addresses to connect to the Internet. NAT, defined originally in RFC 1631 and enhanced in RFC 3022, operates in routers and other devices to provide address simplification and conservation. Usually, NAT connects two networks and translates the

private (inside local) addresses in the internal network into public addresses (inside global) before packets are forwarded to another network, as shown in [Figure 9-10](#). As part of this functionality, you can configure NAT to use only one address for the entire network to the outside world. Using only one address effectively hides the internal network, thus providing additional security.

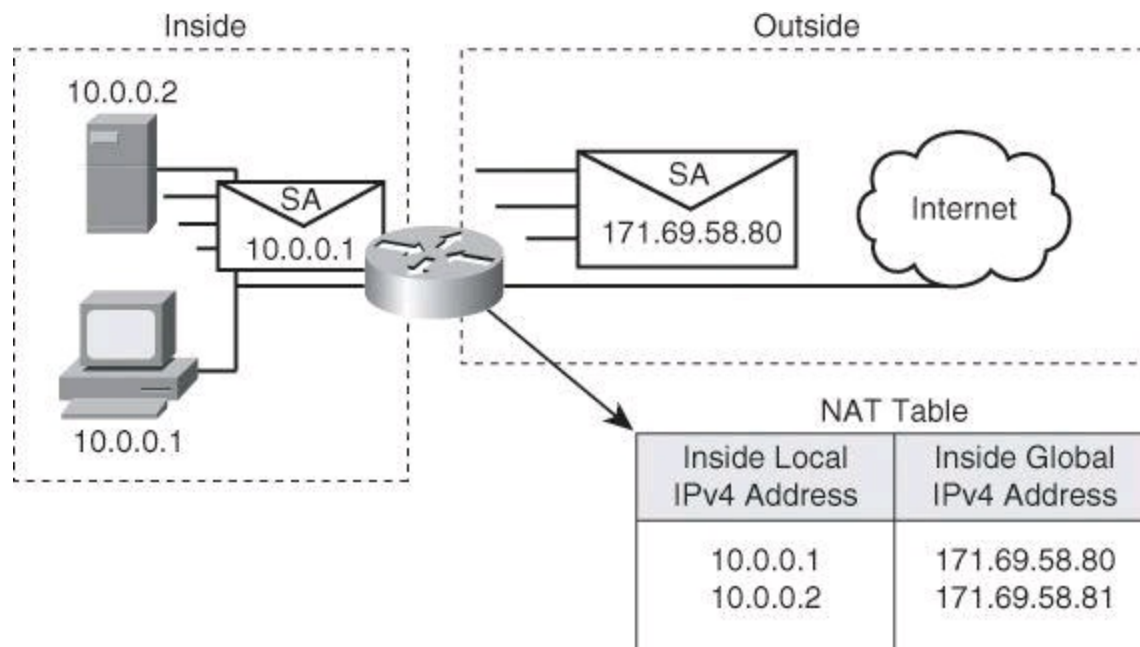


Figure 9-10. Example of Network Address Translation

[Chapter 3](#) mentioned RFC 1918 and the private address space. These private addresses provide a very large pool of addresses to be used inside the corporate network. These addresses, which are not routable, need to be translated for outbound traffic. A common practice is for these inside addresses to be translated by a single public IP address, often the IP address of the outside-facing interface of the firewall. This technique of translating many inside users to one single outside IP address will be discussed later in this chapter.

Multiple device types can be configured to perform NAT services. Although firewalls are most common, routers and Layer 3 switches are also capable of deploying this service.

Cisco defines the following list of NAT terms:

- **Inside local address:** The IPv4 address that is assigned to a host on the inside network. The inside local address is likely not an IPv4 address assigned by a service provider, but rather one that falls within reserved private IPv4 address spaces.
- **Inside global address:** A legitimate IPv4 address that is assigned by the network interface card (NIC) or service provider that represents one or more inside local IPv4 addresses to the outside world.
- **Outside local address:** The IPv4 address of an outside host as it appears to the inside network. Not necessarily a public address, the outside local address is allocated from a routable address space. Outside NAT (NATting the address of an outside entity when it appears on the inside network) is beyond the scope of this book.
- **Outside global address:** The IPv4 address that is assigned to a host on the outside network by the host owner. The outside global address is allocated from a globally routable address or network space.

NAT offers the following benefits:

- Eliminates the need to readdress all hosts that require external access, saving time and money.
- Conserves addresses through application port-level multiplexing. With NAT overloading, internal hosts can share a single registered IPv4 address for all external communications. In this type of configuration, relatively few external addresses are required to support many internal hosts, thus conserving IPv4 addresses. NAT overload, a term used by Cisco, is often referred to as PAT, which will be covered next.

One of the main features of NAT is dynamic Port Address Translation (PAT), which is also referred to as “NAT overload” in Cisco IOS configuration. PAT allows you to translate multiple internal addresses into a few external addresses, or even a single external address, essentially allowing the internal addresses to share one external address, as shown in [Figure 9-11](#). With NAT overload, as shown in [Figure 9-11](#), as long as the local host uses a unique source port, the router translates only the IP address and not the port. However, in the third entry, host 10.6.1.6 uses source port number 2031, which is already in use by host 10.6.1.2. In this situation, the router translates both source address and source port to create a unique combination that will be used to allow the return traffic to be mapped and translated back to the proper inside host.

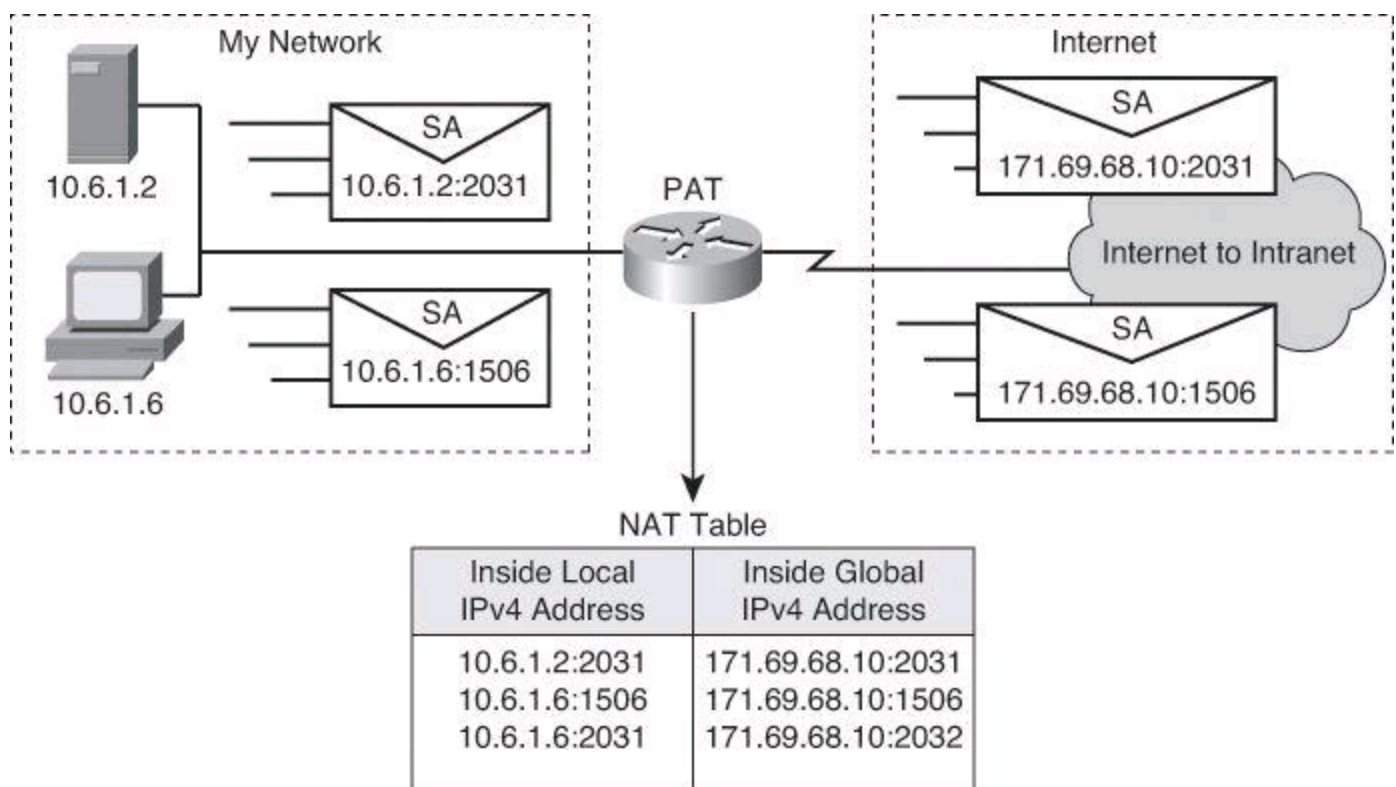


Figure 9-11. Example of Port Address Translation (aka NAT Overload) on Cisco IOS Router

PAT uses unique source port numbers on the inside global IPv4 address to distinguish between translations. Because the port number is encoded in 16 bits, the total number of internal addresses that NAT can translate into one external address is, theoretically, as many as 65,536 minus the 1023 reserved ports, so about 64,000 ports. Port numbers from 0 to 1023 are used for well-known ports and thus typically aren't used for source ports.

PAT attempts to preserve the original source port. If the source port is already allocated, PAT attempts to find the first available port number. It starts from the beginning of the appropriate port

group, 0 to 511, 512 to 1023, or 1024 to 65535. If PAT does not find an available port from the appropriate port group and if more than one external IPv4 address is configured, PAT moves to the next IPv4 address and tries to allocate the original source port again. PAT continues trying to allocate the original source port until it runs out of available ports and external IPv4 addresses.

Dynamic PAT is typically used when translating private addresses into routable Internet addresses, for outbound network traffic. You can configure static or dynamic inside source translation; however, dynamic PAT is typically used because there are not enough global addresses.

Example of Translating an Inside Source Address

[Figure 9-12](#) illustrates a router that is translating a source address inside a network into a source address outside the network. The steps for translating an inside source address are as follows:

Step 1. The user at host 1.1.1.1 opens a connection to Host B.

Step 2. The first packet that the router receives from host 1.1.1.1 causes the router to check its NAT table.

- If a static translation entry has been configured, the router goes to Step 3.
- If no static translation entry exists, the router determines that the source address 1.1.1.1 (SA 1.1.1.1) must be translated dynamically. The router then selects a global address from the dynamic address pool and creates a translation entry (in the example, 2.2.2.2).

Step 3. The router replaces the inside local source address of host 1.1.1.1 with the translation entry global address and forwards the packet.

Step 4. Host B receives the packet and responds to host 1.1.1.1 by using the inside global IPv4 destination address 2.2.2.2 (DA 2.2.2.2).

Step 5. When the router receives the packet with the inside global IPv4 address, the router performs a NAT table lookup by using the inside global address as a key. If dynamic PAT is configured, the lookup includes the destination port as well, which acts as a tiebreaker. The router finds the original host IP address and translates back to the inside local address 1.1.1.1, forwarding the packet to host 1.1.1.1.

Step 6. Host 1.1.1.1 receives the packet and continues the conversation. The router performs Steps 2 through 5 for each packet.

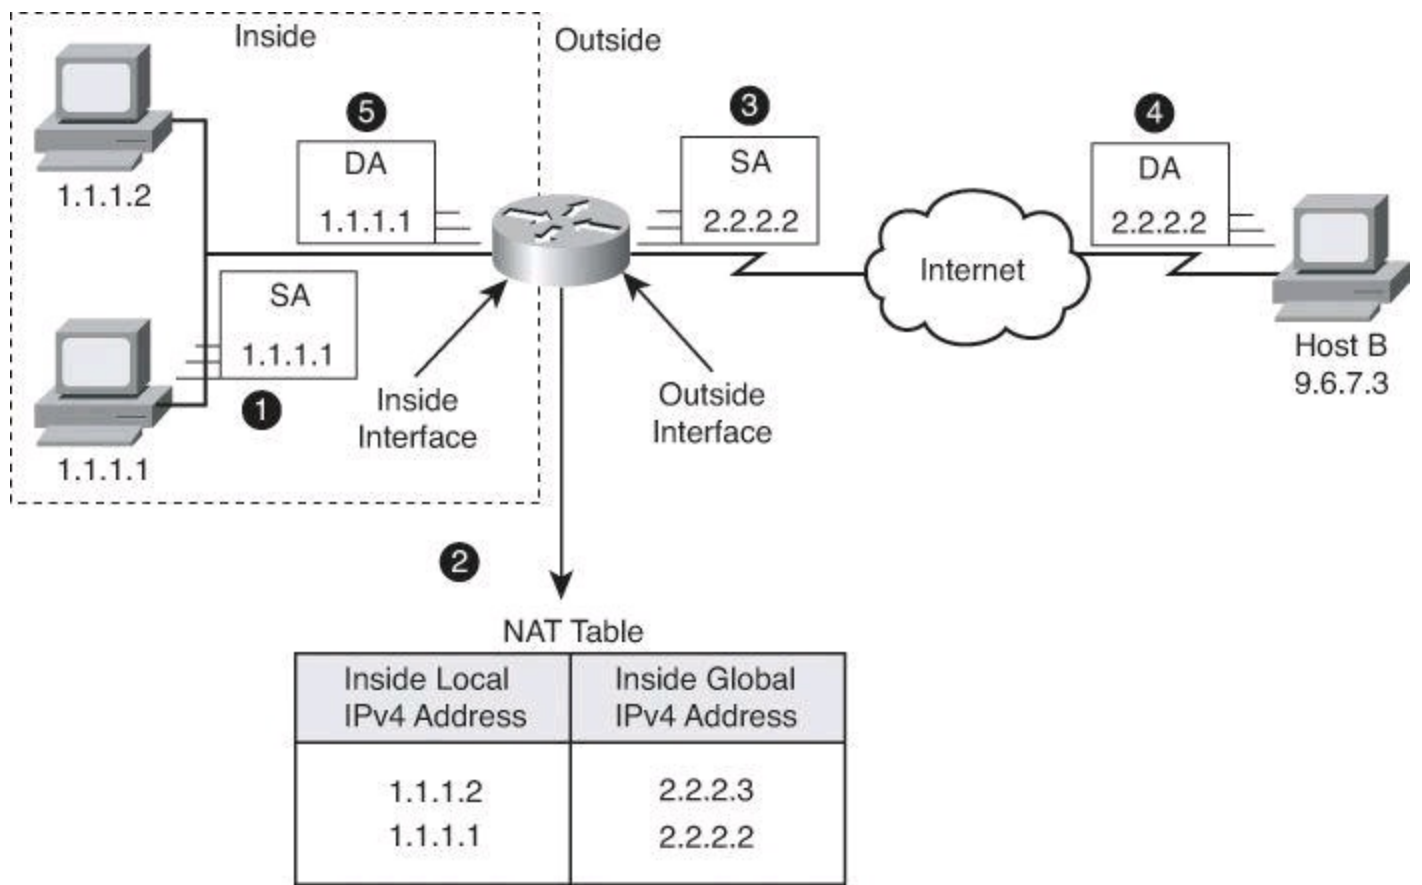


Figure 9-12. Translating Inside Source Address

Inbound access requires a different type of NAT. For servers such as host 1.1.1.1 in [Figure 9-13](#), inbound access from the Internet requires a unique and persistent global IPv4 address. The service must have a known and unique identifier (IP address) for DNS servers to translate the service URL into the same global IPv4 address and make the service reachable.

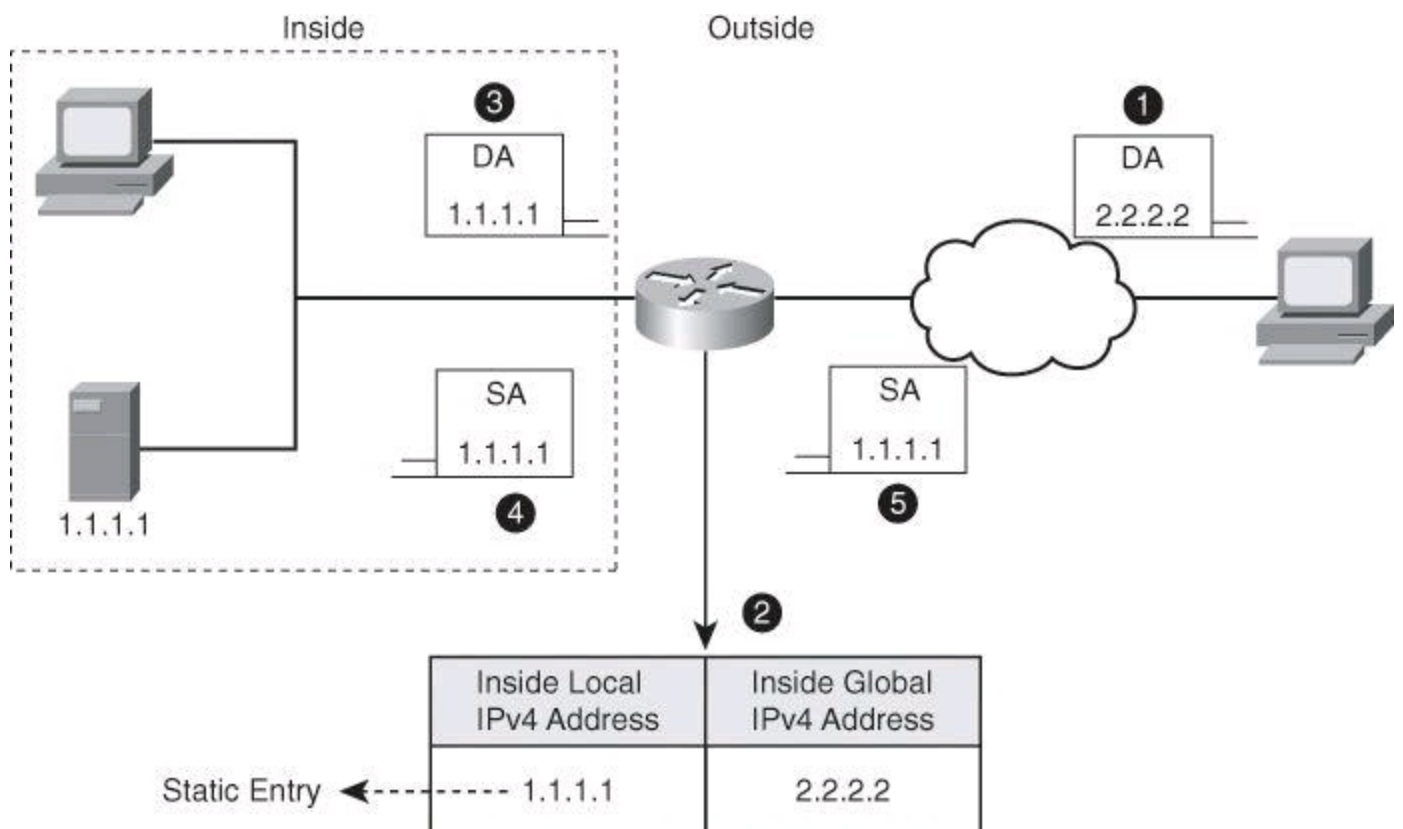


Figure 9-13. Static Translation

For this reason, static NAT is used to manually configure the translation between the known inside local address, in the example 1.1.1.1, and the known global address, in the example 2.2.2.2. This is a static NAT entry that makes all packets be translated:

- A client sends a request, shown in step 1 of [Figure 9-13](#), to 2.2.2.2 which is the public address of the server located inside our network. Upon receiving the packet, shown in Step 2, the router translates the destination address of 2.2.2.2 to 1.1.1.1.
- The router forwards the packet, with the real destination IP address, to the server, as shown in Step 3.
- The server returns its reply with source IP address of 1.1.1.1, shown in Step 4.
- The router looks at its NAT table and finds a static entry for 1.1.1.1 and translates it to its global address of 2.2.2.2 before forwarding the packet to the client, as shown in Step 5.
- Outbound packets, originated by the server 1.1.1.1, find the matching static entry as well, and all packets are translated accordingly.

NAT Deployment Choices

The deployment modes in NAT operations are as follows:

- **Static NAT:** Maps an unregistered IPv4 address to a registered IPv4 address. One-to-one static NAT is particularly useful when a device must be accessible from outside the network.
- **Dynamic NAT:** Maps an unregistered IPv4 address to a registered IPv4 address from a group of registered IPv4 addresses. Dynamic NAT is “many-to-many” translation: in this scenario, many inside users are translated to many global addresses. These global addresses are grouped together in a NAT pool. Dynamic NAT is useful for outbound client access, when you have fewer outside global IP addresses than inside local hosts. Each inside user will have a public address that will be assigned and reserved to it for a period of time for NAT.
- **Dynamic PAT (NAT overload):** Maps multiple unregistered IPv4 addresses to a few or even a single registered IP address (many-to-one) by tracking port numbers and using different ports, if needed. NAT Overload is also known as PAT, and is a form of dynamic translation.
- **Policy NAT:** Uses ACLs to perform NAT differently based on ACL criteria (for instance source addresses, destination addresses, and ports). A group of hosts might translate to a global address pool when sending traffic to the Internet, while translating to a nonoverlapping private address when sending traffic to remote offices.
- **Static PAT:** Uses ACLs to perform translation by allowing a specific UDP or TCP port on a global address to be translated to a specific port on a local address. That is, both the address and the port numbers are translated. This is referred to as port forwarding on home routers, such as Linksys models.

PAT is not implemented in exactly the same way in Cisco IOS and Cisco ASA. In Cisco IOS, PAT translates the source port only if needed. With Cisco ASA, the source port of an outbound connection is *always* translated to a randomly generated port number.

Firewall Designs

Ultimately, firewalls are referred to as *firewall systems* because they implement access control using multiple technologies. These technologies might even be implemented in different physical devices. When planning your firewall system, it is important to enable and use the firewall technologies that match your security requirements.

Typically, stateful packet filtering is the building-block technology. Stateful packet filters are good central access control components. Access rules are created for Layer 3 and Layer 4 criteria to define which traffic classes can traverse the firewall. These access rules define the first level of trust relationships that are enforced through the firewall.

As you design the access rules, other firewall technologies will be necessary. Application layer inspection expands matching criteria and access control actions to upper layers, providing an additional level of security. Proxy services, including user authentication, allow for user-based rules that result in more granular and context-aware access control.

Best practices documents are a composite effort of security practitioners. This partial list of best practices is generic and serves only as a starting point for your own firewall security policy:

- Position firewalls at key security boundaries, separating security domains with different levels of trust.
- Firewalls are the primary security device, but it is unwise to rely exclusively on a firewall for security.
- Deny all traffic by default and permit only services that are needed.
- Implement various firewall technologies, matching your application mix and security policy requirements.
- Ensure that physical access to the firewall is controlled.
- Regularly monitor firewall logs. Cisco Security Manager and other Cisco management tools are available for this purpose.
- Practice change management for firewall configuration changes.

Remember that firewalls primarily protect from technical attacks originating from the outside. Inside attacks tend to be nontechnical in nature.

Firewall Policies in a Layered Defense Strategy

Firewall access rules are the building blocks for access control. An enterprise firewall could have hundreds if not thousands of rules. Access rule creation, management, and auditing are critical to implementing a sound security posture effectively and efficiently. Too many redundant rules might result in performance degradation. Ineffective or obsolete rules might result in trust exploitation and security breaches.

When defining access rules, multiple criteria can be used as a starting point:

- **Rules based on service control:** Determine the types of Internet services that can be accessed, inbound or outbound. The firewall may filter traffic based on IP address and TCP port number. It may provide proxy software that receives and interprets each service request before passing it on. It may host the server software itself, such as a web or mail service. An example of a rule based on service control would be: allow HTTP, allow HTTPS, deny everything else.
- **Rules based on direction control:** Determine the direction in which particular service requests may be initiated and allowed to flow through the firewall. An example of a rule based on direction control would be: allow HTTP outbound, but not inbound.
- **Rules based on user control:** Control access to a service according to which user is attempting to access it. This feature is typically applied to users inside the firewall perimeter (local users). It may also be applied to incoming traffic from external users. The latter requires some form of secure authentication technology, such as is provided in IP Security (IPsec). An example of a rule based on user control would be: allow campus VLANs HTTP access, deny it for wireless VLANs.
- **Rules based on behavior control:** Control how particular services are used. For example, the firewall may filter email to eliminate spam, or it may enable external access to only a portion of the information on a local web server. An example of a rule based on behavior control would be open negotiated FTP ports after learning them during connection setup.

Typically, a combination of these criteria is used to define restrictive policies that permit only the required traffic, while blocking anything else.

Access rules are typically implemented using ACLs. That is the case with Cisco IOS and Cisco ASA firewalls. For that reason, access rules are effectively ordered lists of **permit** and **deny** statements, applied to specific interfaces, as represented in [Figure 9-14](#).

Action	Source Address	Destination Address	Service
DENY	10.1.1.0 /24	192.168.1.0 /24	Any
PERMIT	10.1.1.0 /24	172.16.0.0 /24	SSH
PERMIT	10.1.1.0 /24	ANY	HTTP
DENY	ANY	ANY	ANY

Figure 9-14. Firewall Access Rule Structure: Top-Down Process

The ruleset is evaluated sequentially from the top, and the rule that first matches the new connection is evaluated, permitting or denying the connection. When a rule matches a connection, all subsequent rules are ignored.

Because rules use ACLs, there is an implicit deny-all rule at the end of the list. If a connection does not match any of the rules, by default, it will be denied.

Firewall Rules Design Guidelines

When creating firewall rules, use the following guidelines:

- Use a restrictive approach, as opposed to a permissive approach, for all interfaces and all

directions of traffic. This translates into simple rulesets that permit the required traffic flows and deny everything else.

- The previous guideline implies that you should take a somewhat paranoid stance and assume that malicious traffic could come from any security domain, even the most trusted ones. This stance results in outbound rules becoming as important as inbound rules, even though they apply to “trusted” internal traffic.
- A tradeoff should be made between access control, performance, and rule maintenance. Specific and granular rules provide more control over the allowed or denied traffic, but also result in longer rulesets and more change-management and maintenance challenges. Observe the use of **any** keywords matching sources, destinations, and ports, as they may allow unwanted traffic.
- For a given access requirement, more specific rules should be pushed to the top, while more generic rules should be pushed to the bottom. The top-down processing order demands this approach to prevent shadow rules when creating exceptions to a given access policy.
- Filtering impossible packets is a common practice. An example is antispoofing rules that are aimed at blocking private IPv4 address spaces in the source address of inbound Internet packets.
- Auditing and change management are crucial. Obsolete rules tend to clutter the ruleset and cause breaches to the access policy.

Lack of a formal methodology for rule creation and maintenance often results in issues that affect the security posture or the ability of the firewall to process traffic.

These issues are, among others, based on consistency, completeness, and compactness. Consistency issues result in potential security breaches.

The following is a list of potential consistency issues with firewall rules:

- **Promiscuous rules:** Allow more access than necessary to meet access requirements. An example is the use of **any** keywords in ACLs, allowing access from all IP addresses when access is only necessary from a single subnet. Another example is a rule that allows access to a system on multiple ports when only a subset of those ports is required.
- **Redundant rules:** Result when a given rule duplicates a portion of the access that is permitted or denied by another rule. An example is a rule that allows a single IP address to access a server on a particular TCP port when an existing rule already allows all IP addresses to access that port. Another example is a rule that allows access to a port on a host when other rules allow the same type of access due to the membership of the port in a service group.
- **Shadowed rules:** Result when the incorrect ordering of rules in the ruleset completely prevents the execution of one or more firewall rules. An example is configuring a rule that denies access to a particular website below a rule that allows access to all websites. Another example is placing a rule that allows access to a server on a single port from a single IP address below a rule that blocks all access to the server from all IP addresses. The existence of the first rule would shadow the second rule, rendering it useless.

- **Orphaned rules:** These are rules that exist in the firewall ruleset but never match traffic passing through the firewall. An example is a rule designed to allow access to a database server that incorrectly specifies a nonexistent destination IP address. Another example is a rule designed to allow HTTP access to a server that no longer hosts a website.

Tip

Here's a trick used by David Chapman, a technical reviewer on this book: to discover orphan rules, reset access list counters on the first of each month. At the end of the month, check which ACL entries still have their counters at zero, and flag these entries for removal.

Completeness issues might also result in security breaches:

- **Errors in rule specification:** Include errors that are made in the translation between business requirements and technical or product-specific firewall rules. Sometimes, the rules simply do not meet business requirements due to a misunderstanding about business requirements. Other times, firewall administrators fail to specify a rule necessary to meet business requirements. Specifying the incorrect port for a service is an example, such as creating a rule for port 8080 when the business requirement called for standard port 80 for the HTTP service.
- **Data entry errors:** Include errors made when converting a technical rule definition into the firewall policy format and entering that rule into the firewall rule base. An example would be mistyping a port number, such as creating a rule for port 23 when the technical specification indicated port 53. Another example would be mistyping a source or destination address or failing to input a rule included in the technical specification.

Compactness issues typically result from obsolete rules, too much granularity, or an informal approach to defining the order of rules within the ruleset. These issues affect performance and, therefore, the ability to provide firewall services effectively. The firewall evaluates the rules inefficiently, often matching traffic flows after hundreds if not thousands of rules have been unnecessarily evaluated.

Summary

In this chapter, you have learned that a firewall is a set of rules designed to enforce an access control policy between two networks. You learned how to create firewall rules to implement your security policies. You also learned that firewall systems are a combination of packet filtering, stateful inspection, application layer inspection, and other technologies. You saw that NAT is a complementary technology to firewalling and that it allows for dynamic and static translations.

You saw that change management, ruleset auditing, and monitoring are critical firewall design considerations and that common issues found in firewall rule design include promiscuous rules, redundant rules, shadowed rules, and orphaned rules.

References

For additional information, refer to these resources:

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, "[Answers to Chapter Review Questions](#)."

- 1.** Which of the following are firewalls? (Choose all that apply.)
 - a. Static packet filters
 - b. URL filters
 - c. Application layer gateways
 - d. Layer 2 switches
- 2.** Stateful packet-filtering firewalls are good to use for all the following applications except:
 - a. As a primary means of defense
 - b. As an intelligent first line of defense
 - c. To sanitize web traffic of malware
 - d. As a means of strengthening packet filtering
 - e. To improve routing performance
- 3.** Match the following names and characteristics:

Names

 - a. Packet-filtering firewalls
 - b. Application layer gateways
 - c. Stateful packet filters
 - d. Application inspection firewalls

Characteristics

 1. Work primarily at the network level of the OSI model
 2. Are the most common firewalls
 3. Monitor sessions to determine the port numbers for secondary channels
 4. Were the first application layer firewalls
- 4.** Which threats can be mitigated using ACLs? (Choose all that apply.)
 - a. Malicious payloads
 - b. Spam
 - c. ICMP message filtering
 - d. IP address spoofing
- 5.** Cisco IronPort Web Security Appliance is an example of which of the following?
 - a. Stateless packet filter

- b.** Application layer gateway
- c.** Stateful packet firewall
- d.** Intrusion detection system

6. Match the following translation definitions and characteristics:

Translation types

- a.** Static NAT
- b.** Dynamic NAT
- c.** Dynamic PAT
- d.** Policy NAT

Characteristics

- 1.** Translation depends on both source and destination
- 2.** Translation is many-to-many
- 3.** Translation is one-to-one
- 4.** Translation is many-to-one

7. Match the following bases of firewall access rules with the examples:

Basis of firewall access rules:

- a.** Service control
- b.** Direction control
- c.** User control
- d.** Behavior control

Firewall access rules examples

- 1.** Allow HTTP outbound, but not inbound
- 2.** Allow campus VLANs HTTP access, deny it for wireless VLANs
- 3.** Open negotiated FTP ports after learning them during connection setup
- 4.** Allow HTTP, allow HTTPS, deny everything else

Chapter 10. Cisco Firewalling Solutions: Cisco IOS Zone-Based Firewall and Cisco ASA

Two of the Cisco Firewall solutions, Cisco IOS Zone-Based Policy Firewalls and Cisco Adaptive Security Appliance, can be configured to perform basic security operations on a network. At the end of this chapter, you will be able to do the following:

- Introduce and describe the function, operational framework, and building blocks of Cisco IOS Zone-Based Firewalls
- Describe the functions of zones and zone pairs, as well as their relationship in hierarchical policies
- Describe Cisco Common Classification Policy Language for creating zone-based firewall policies
- List the default policies for the different combinations of zone types
- Demonstrate the configuration and verification of zone-based firewalls using Cisco Configuration Professional and the CLI
- Demonstrate the configuration of NAT services for zone-based firewalls
- Describe the Cisco ASA family of products, identifying key supported features
- Describe the building blocks of Cisco ASA configuration
- Describe the navigation options, features, and requirements of Cisco ASDM
- Describe the use of access control lists on Cisco ASA
- Describe the deployment of policies using the Cisco Modular Policy Framework
- Describe the configuration procedure to deploy basic outbound access control on Cisco ASA using Cisco ASDM

The zone-based policy firewall changes the original implementation of Cisco IOS Classic Firewall stateful inspection from the older interface-based model to a more flexible, more easily understood zone-based configuration model. The first section of this chapter focuses on the features of Cisco IOS Zone-Based Policy Firewalls and how to use Cisco Configuration Professional to configure them.

Cisco Adaptive Security Appliance (ASA) implements a rich set of security technologies and can be effectively implemented as a perimeter firewall using several deployment modes. The second section of this chapter introduces Cisco ASA functionality, features, and underlying technologies and demonstrates how to configure the Cisco ASA 5505 model for basic connectivity using Cisco Adaptive Security Device Manager (ASDM).

Cisco Firewall Solutions

Cisco offers multiple different firewall solutions, each geared to a different environment. Currently, Cisco Firewall offerings include

- Cisco IOS Firewall
- Cisco ASA 5500 Adaptive Security Appliances
- Cisco ASA 1000V Cloud Firewall

- Cisco Virtual Security Gateway for Nexus 1000V Series Switch
- Cisco Catalyst 6500 Series ASA Services Module
- Cisco Catalyst 6500 Series Firewall Services Module
- Cisco Small Business SA500 Series Security Appliances

In this chapter, we cover the two first firewall technologies: Cisco IOS Firewall and Cisco ASA firewall.

Cisco IOS Zone-Based Policy Firewall

Many different types of techniques are used for firewalling. Cisco IOS routers have been using many different techniques for firewalling over the years. Recently, Cisco introduced Zone-Based Policy Firewall as an alternative to the older technology called Context-Based Access Control.

Zone-Based Policy Firewall Overview

The original implementation of Cisco IOS Classic Firewall stateful inspection used an interface-based configuration model, in which a stateful inspection policy was applied to an interface. All traffic passing through that interface received the same inspection policy. This configuration model limited the granularity of the firewall policies and caused confusion of the proper application of firewall policies, particularly in scenarios when firewall policies must be applied between multiple interfaces.

Zone-based policy firewalls (sometimes referred to as ZBF, or zone-policy firewall [ZPF]) change the firewall from the older interface-based model to a more flexible, more easily understood configuration model where interfaces are assigned to zones, and an inspection policy is applied to traffic moving between the zones. To demonstrate this model, [Figure 10-1](#) shows three zones:

- **Untrusted:** Represents the Internet
- **DMZ:** Demilitarized zone, which contains the corporate servers accessed by the public
- **Trusted:** Represents the inside network

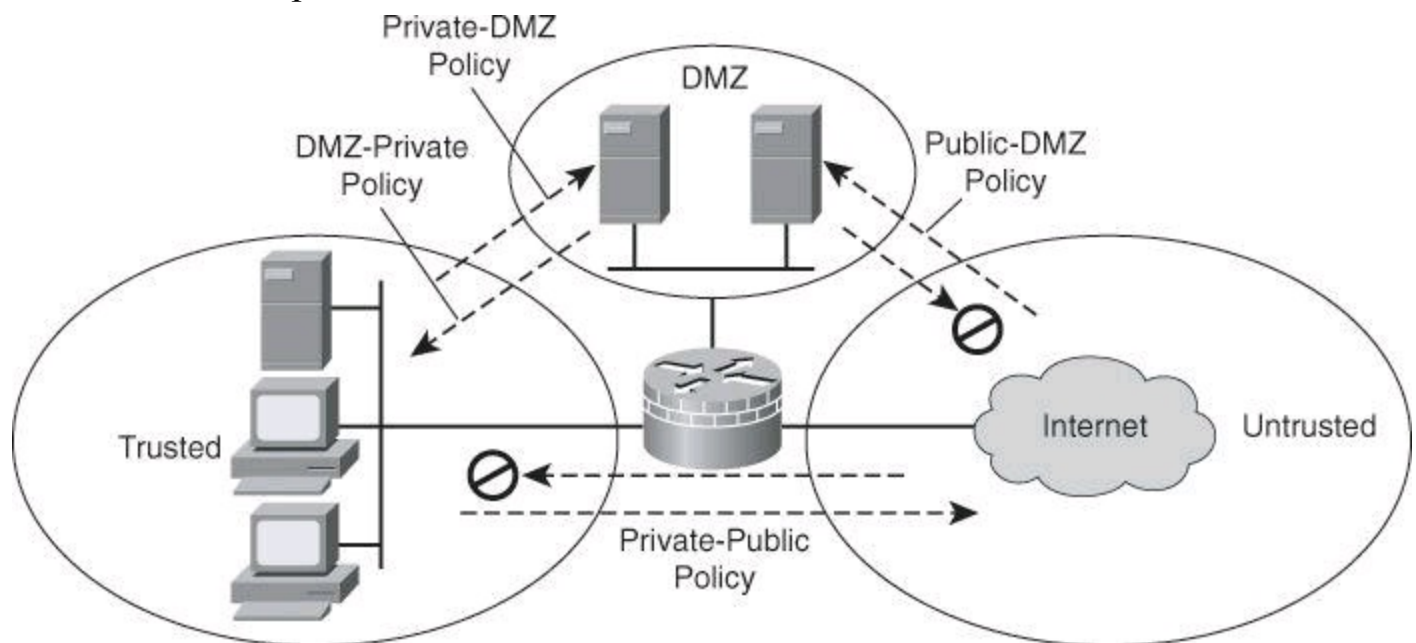


Figure 10-1. Cisco IOS Zone-Based Policy Firewall

Note

Notably missing from [Figure 10-1](#) is the self zone, the zone used for traffic generated by the router or destined to it, such as routing protocol traffic. The self zone is covered later in the chapter.

Interzone policies offer considerable flexibility and granularity, and thus enable you to apply different inspection policies to multiple host groups that are connected to the same router interface. The interzone policies in [Figure 10-1](#) are as follows:

- **Public-DMZ:** DMZ policy that sets the rules for traffic originating from the untrusted zone with the DMZ as destination
- **DMZ-Private:** Private policy that sets the rules for the traffic originating from the DMZ with the trusted zone as destination
- **Private-DMZ:** DMZ policy that sets the rules for the traffic originating from the trusted zone with the DMZ as destination.
- **Private-Public:** Public policy that sets the rules for the traffic originating from the trusted zone with the untrusted zone as destination

Rules are bidirectional. As an example, traffic originating from the trusted zone going to the DMZ would be allowed to return due to the Private-DMZ policy, which would permit the connection and allow the return traffic to come back in the trusted zone. The DMZ-Private policy is not involved for return traffic that originated from the trusted network and for which the reply is now going from the DMZ to the trusted network. The DMZ-Private policy is only needed for traffic originating from the DMZ and destined for the trusted network. This could be the case for email: all the incoming mail could end up in the DMZ to be analyzed by the Cisco IronPort Email Security Appliance (ESA). The ESA would forward to Microsoft Exchange Server, located on the trusted network, email it considered appropriate. The connection would originate from the ESA located in the DMZ and go to MS Exchange Server located on the trusted network. In this case, the DMZ-Private policy would be invoked to permit this traffic. This is likely the only permit action on the DMZ-Private policy: let email flow from the ESA to MS Exchange Server.

The firewall will have a rule that denies email to flow directly from the Internet to MS Exchange Server located on the trusted network.

Zone-based policy firewalls are configured with the Cisco Common Classification Policy Language (C3PL), which uses a hierarchical structure to define inspection for network protocols and the groups of hosts to which the inspection will be applied.

Cisco IOS Zone-Based Policy Firewalls support the following features:

A circular icon with a dotted border containing the text "Key Topic".

Key
Topic

- Stateful inspection
- Application inspection
- URL filtering

- Per-policy parameter
- Transparent firewall
- Virtual routing and forwarding aware firewall

First-generation firewalls, which were essentially packet filters, used only ACLs to control traffic. For this reason, it was relatively easy for attackers to breach a firewall, because no state data was examined. Second-generation firewalls, known as proxy firewalls, were concerned with the state of the connection, but they were application dependent. Third-generation firewalls, which were stateful packet filters, were developed to provide state tracking, application independence, and speed. Context-Based Access Control (CBAC), from the legacy Cisco IOS Firewall feature set, is an example of this type of third-generation firewall. CBAC performed traffic filtering using stateful filters, whereby selected outbound traffic created a temporary opening in the filter for replies to return through.

Note

CBAC can still be implemented on Cisco IOS Firewalls; however, interfaces that participate in a zone-based policy firewall cannot also participate in CBAC.

The legacy stateful inspection done by CBAC was very complicated because there was a combination of ACLs and inspection rules, all of which worked together to accomplish the stateful packet filtering done by the firewall.

The Cisco IOS Zone-Based Policy Firewall completely changes the way a Cisco IOS Firewall is configured.

The first major change to the firewall configuration is the introduction of a zone-based configuration. The classic Cisco IOS Firewall stateful inspection (CBAC) used an interface-based configuration model that you configured using the **ip inspect** command. This changed with the introduction of the Cisco IOS Zone-Based Policy Firewall, which does not use the stateful inspection (CBAC) commands but rather uses C3PL. You could use the two configuration models concurrently on routers, but not on the same interfaces; you cannot configure an interface as a security zone member and for **ip inspect** simultaneously.

Zones establish the security borders of your network. A zone defines a boundary where traffic is subjected to policy restrictions as it crosses to another region of your network. The default policy of a zone-policy firewall between zones is to “deny all.” If no policy is explicitly configured, all traffic moving between zones is blocked. This policy is a significant departure from the stateful inspection model, in which traffic is implicitly allowed unless it is explicitly blocked with an ACL entry.

The second major change is the introduction of the new classification policy language known as C3PL. The C3PL structure is similar to the Modular QoS CLI (MQC) structure in which class maps specify the traffic that is affected by the action that the policy map applies. C3PL will be discussed later in this chapter.

Note

Interface ACLs are still relevant and are applied before zone-based policy firewalls when

they are applied inbound. Interface ACLs are applied after zone-based policy firewalls when they are applied outbound.

Key benefits of zone-based policy firewall are as follows:

Key
Topic

- It is not dependent on ACLs.
- The router security posture is restrictive (which means block unless explicitly allowed).
- C3PL makes policies easy to read and troubleshoot.
- One policy affects any given traffic instead of needing multiple ACL and inspection actions.

Zones and Zone Pairs

A zone is a collection of networks that are reachable over a specific router interface and are designated to require the same security policies. Zone-based policy firewall access control policies then control access between two or more zones that are configured on the router, using a flexible configuration language that allows you to specify simple or complex access policies in a manageable manner.

Once zones are created, default rules apply. These default rules will be presented in more detail later in this chapter. Default rules generally allow unrestricted access between interfaces that belong to the same zone, while denying all interzone traffic. [Figure 10-2](#) shows two interfaces belonging to Zone 2. These two interfaces can pass traffic between them without any restrictions.

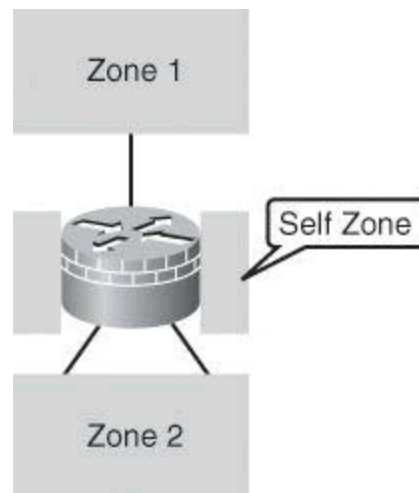


Figure 10-2. Interfaces Belong to Zone

Key
Topic

Note

An interface can belong to one zone and one zone only.

A zone pair allows you to specify a unidirectional firewall policy between two security zones, thus changing the default policy. The direction of the traffic is indicated by specifying a source zone and a destination zone. The source and destination zones of a zone pair must be previously created security zones.

To permit traffic between zone-member interfaces, you must configure a policy permitting (or inspecting) traffic between that zone and another zone.

Self Zone

If desired, you can select the default or self zone as either the source zone or the destination zone. The self zone is a system-defined zone. It does not have any interfaces as members. A zone pair that includes the self zone, along with the associated policy, applies to traffic directed to the router or traffic that is generated by the router. A zone pair does not apply to traffic through the router. Traffic to or from the self zone is by default permitted.

Zone-Based Topology Examples

A zone-based firewall approach results in multiple deployment options that will represent your security domains, which is a more intuitive way to implement firewall policies than a per-interface approach.

[Figure 10-3](#) shows a simple firewall with two security domains, suitable for remote office environments where outbound access is required only to the Internet or to a hub network. Trusted internal interfaces belong to an “inside” zone, while untrusted external interfaces belong to an “outside” zone.



Figure 10-3. Simple Firewall Topology with Two Security Domains

[Figure 10-4](#) shows a medium-sized location requiring inbound access to a few public servers, in addition to outbound user access. A third zone, the DMZ, includes the interfaces and assets that are related to the servers located in DMZ interfaces, which require restricted access from the outside, and perhaps a more granular permissive policy from the inside.

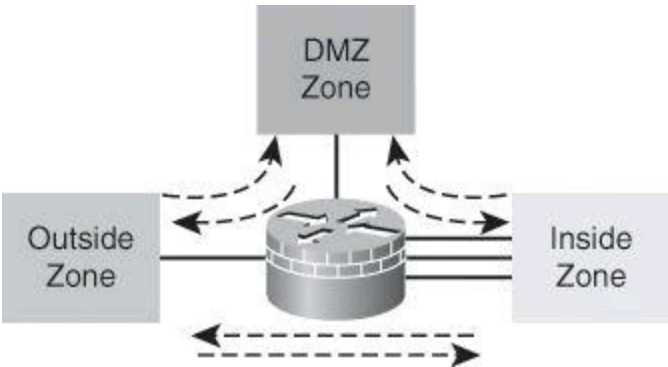


Figure 10-4. Medium-Sized Organization with Three Zones

If an interface doesn't belong to a zone, it can't communicate with an interface that is part of a

zone.



Introduction to Cisco Common Classification Policy Language

A class is a way to identify a set of packets based on their contents. A class is a particular type of traffic; it can be thought of as a precise flow of traffic. Examples of a class of traffic would be traffic for destination port 80 or traffic with a specific subnet source IP address. Normally, you define a class to then apply to the traffic identified by that class an action that reflects a particular policy. A class is designated through class maps.

An action is a specific functionality. It typically is associated with a traffic class. For example, **inspect**, **drop**, and **pass** are actions.

To create firewall policies, complete the following tasks:

Step 1. Define a match criterion (class map).

Step 2. Associate actions to the match criteria (policy map).

Step 3. Attach the policy map to a zone pair (service policy).

A packet arriving at the target (such as the input interface, output interface, or zone pair) is checked against the match criterion that is configured for a class map to determine whether the packet belongs to that class. If so, the policy map defines the action that applies to the specific class. The assignment of the policy to zone pairs is configured using a service policy.

There are three major components to C3PL, as shown in [Figure 10-5](#): class map, policy map, and service policy.

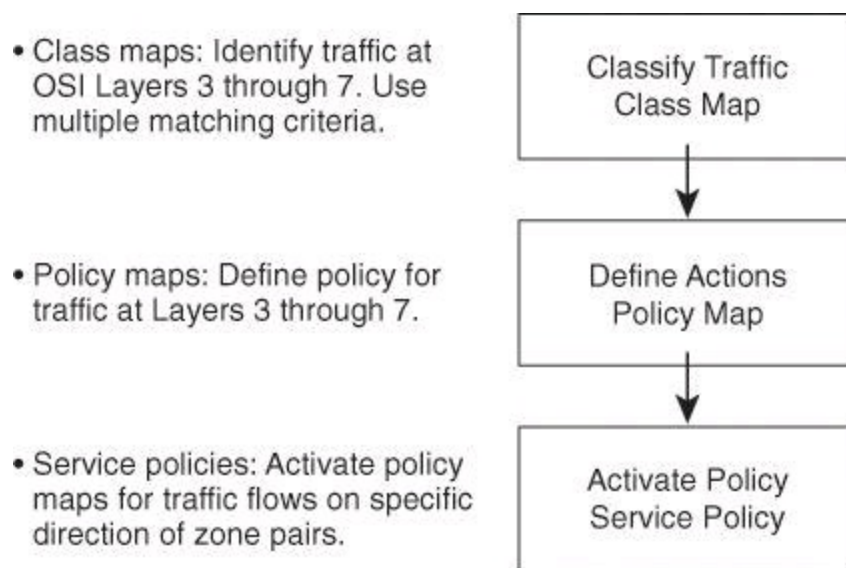


Figure 10-5. Components of Cisco Common Classification Policy Language

Cisco Common Classification Policy Language policies are modular, object oriented, and hierarchical in nature:

• **Modular and object oriented:** These traits give the firewall administrator the flexibility to create building-block objects such as class maps and policy maps, and reuse them within

a given policy and across policies.

- **Hierarchical:** This feature results in powerful policies that can be expanded to include customized inspection, application layer rules, and advanced inspection features. These advanced policies are outside the scope of this book, but are based on the hierarchical structure depicted in [Figure 10-6](#).

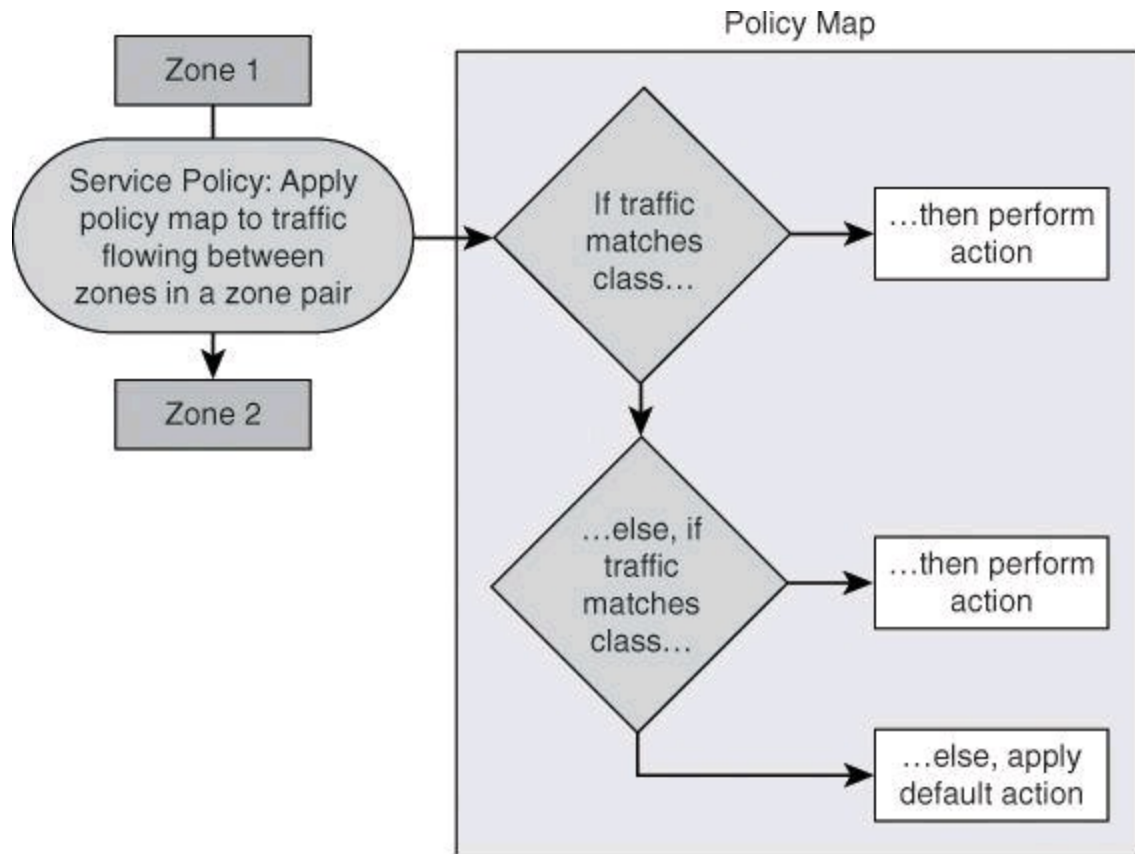


Figure 10-6. C3PL: If-Then-Else Structure

Any given policy map is then applied to traffic flows that move in a given direction between zones by using service policies. The policy map processes traffic flows in that direction by matching sessions against one or more pre-created class map objects and performing access rule actions that are defined for each class.

As shown in [Figure 10-6](#), policy maps can be thought as “if-then-else” statements, something similar to saying: “If a traffic flow matches the first class map, then the associated action is performed, and no other class map in the same policy is evaluated. Else, if the traffic flow matches the next class map in the configuration, then the associated action is performed.” If traffic is not matched by any class map in a policy, a default action is performed.

In [Figure 10-7](#), two class maps—Class 1 and Class 2—are reused across policies Policy 1 and Policy 2. The class maps identify two different traffic classes that will be subject to a different policy for different interface pairs.

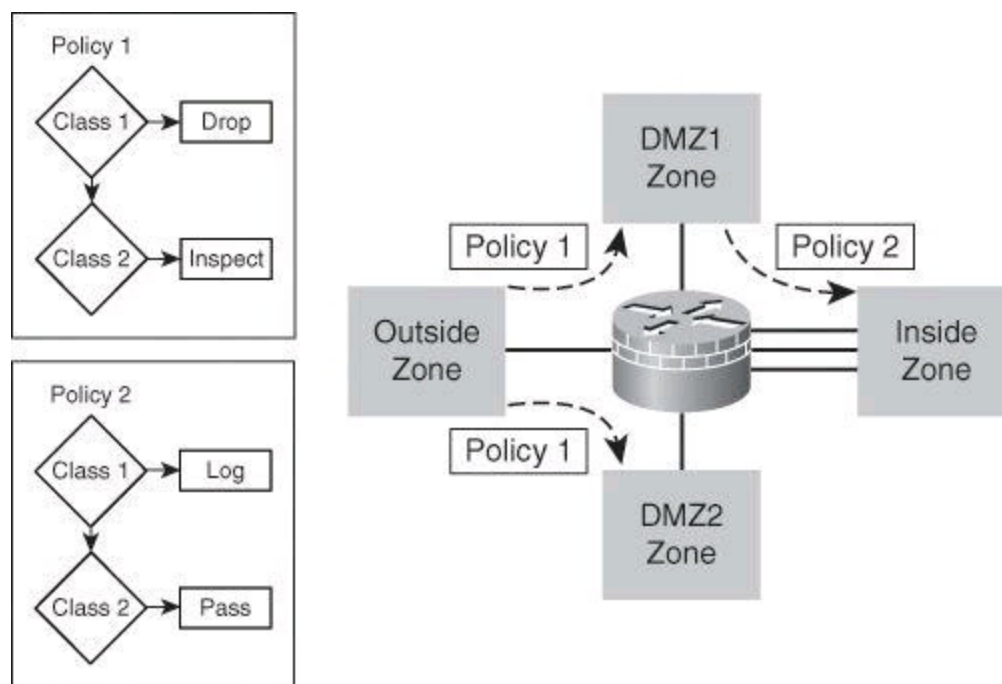


Figure 10-7. Modular Object-Oriented Configuration Design

By the same token, the policy-level objects can also be reused. In [Figure 10-7](#), the Policy 1 policy map is applied to two different zone pairs for the traffic source on the outside and moving toward the DMZ interfaces. However, only one of those DMZs is applied to the Policy 2 policy map, indicating a different set of policies for the same traffic in each instance.

The following are characteristics of class map objects that you should consider:

- Class maps that analyze Layer 3 and Layer 4 traffic sort the traffic based on the following criteria:
 - **Access-group:** A standard, extended, or named ACL can filter traffic based on source and destination IP addresses and source and destination ports.
 - **Protocol:** The class map can identify Layer 4 protocols such as TCP, UDP, and Internet Control Message Protocol (ICMP), and application services such as HTTP, Simple Mail Transfer Protocol (SMTP), and Domain Name System (DNS). Any well-known or user-defined service known to Port-to-Application Mapping (PAM) can be specified. PAM is used to support services that use port numbers that are different from the registered or well-known ports associated with a specific application: An example is an SMTP server listening at port 2525 instead of the standard port 25.
 - **Class-map:** A subordinate class map that provides additional match criteria can be nested inside another class map.
- Each class map can have multiple match statements, as shown in [Figure 10-8](#). The match type defines how multiple match statements are processed to match the class:
 - If **match-any** is specified, traffic must meet any one of the match criteria in the class map.
 - If **match-all** is specified, traffic must match all of the class map criteria to belong to that particular class.



Figure 10-8. Class Map with Multiple Match Statements

Other considerations apply to policy map creation and usage:

- Class maps are evaluated in the order in which they are configured and added to the policy map. Consider pushing the more specific traffic classes to the top of the policy, while pushing the less generic traffic classes to the bottom.
- Possible actions are **inspect**, **pass**, and **drop** (and, optionally, **log** for dropped packets). These actions will be explained in more detail later in this section.
- In addition to user-defined classes, a system-defined class map named *class-default* represents all packets that do not match any of the user-defined classes in a policy, as shown in [Figure 10-9](#). It is always the last class in a policy map. You can define explicit actions for this group of packets. If you do not configure any actions for class-default in an inspect policy, the default action is **drop**.

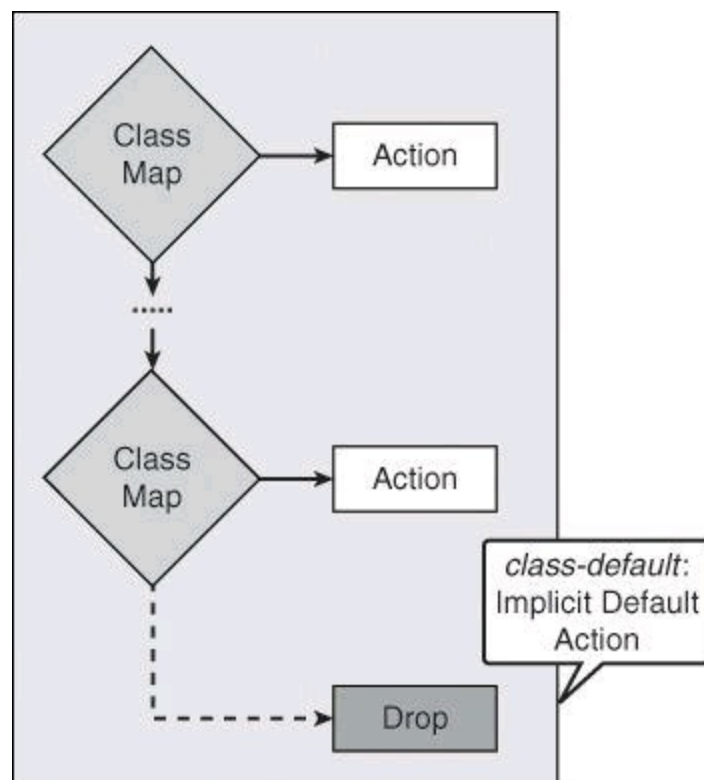


Figure 10-9. Implicit Default Action of the class-default Map

Zone-Based Policy Firewall Actions

The Cisco IOS Zone-Based Policy Firewall can take three possible actions when you configure it using CCP or the CLI:

- **inspect**: This action configures Cisco IOS stateful packet inspection.

- **drop:** This action is analogous to **deny** in an ACL. An additional **log** option can be added to **drop** to log dropped packets.
- **pass:** This action is analogous to **permit** in an ACL. The **pass** action does not track the state of connections or sessions within the traffic; **pass** allows the traffic only in one direction. A corresponding policy must be applied to allow return traffic to pass in the opposite direction.

Note

The **pass** action does not perform stateful inspection.

A circular icon with a dashed border containing the text "Key Topic".

Key
Topic

Note

Possible actions are **inspect**, **drop**, and **pass**, though some documentation might refer to **pass** as **permit**. Also remember that dropped packets can also be logged. So **log** can be considered as a fourth possible action, though only in combination with **drop**.

ZBF Terminology

When editing firewall policies, if you see:

Permit Firewall: Interpret this as “inspect according to class-map”

Permit ACL: Pass

Drop: Drop

Service Policy Zone Pair Assignments

Service policy assignments are unidirectional, as defined by the zone pair where the policy map is assigned. Remember, zone pairs define traffic behavior between zones in a unidirectional manner, based on source and destination zones. A policy map applied to a given zone pair will match traffic and perform action for that traffic in the direction defined by the zone pair. In other words, the policy map applies to traffic originating at the source zone and flowing toward the destination zone.

The final rule is that one and only one service policy can be applied to a given zone pair at any given time. If you apply a different service policy to a zone pair that already has one, you will overwrite the previous service policy.

A circular icon with a dashed border containing the text "Key Topic".

Key
Topic

Note

Zone pairs define traffic behavior between zones based on source and destination zones.

There can be one and only one service policy per zone pair. The policy map applies to traffic originating at the source zone and flowing toward the destination zone.

Zone-Based Policy Firewall: Default Policies, Traffic Flows, and Zone Interaction

The membership of the router network interfaces in zones is subject to several rules governing interface behavior, as is the traffic moving between zone member interfaces:

- A zone must be configured before you can assign interfaces to the zone.
- You can assign an interface to only one security zone.
- Traffic is implicitly allowed to flow by default among interfaces that are members of the same zone.
- To permit traffic to and from a zone member interface, a policy allowing or inspecting traffic must be configured between that zone and any other zone.
- Traffic cannot flow between a zone member interface and any interface that is not a zone member. You can apply **pass**, **inspect**, and **drop** actions only between two zones.
- Interfaces that have not been assigned to a zone function as classical router ports and might still use classical stateful inspection (CBAC) configuration.
- If you do not want an interface on the router to be part of the zone-based firewall policy, it might still be necessary to put that interface in a zone and configure a “pass all” policy (sort of a dummy policy) between that zone and any other zone to which traffic flow is desired.
- From the preceding rules it follows that if traffic is to flow among all the interfaces in a router, all the interfaces must be part of the zoning model (each interface must be a member of a zone).

[Table 10-1](#) shows a number of examples of different interface and configuration combinations.

Table 10-1. Zone-Based Policy Firewall: Rules for Application Traffic

Source Interface Member of Zone?	Destination Interface Member of Zone?	Zone Pair Exists?	Policy Exists?	Result
No	No	N/A	N/A	No impact of zoning/policy
Yes (zone 1)	Yes (zone 1)	N/A*	N/A	No policy lookup (pass)
Yes	No	N/A	N/A	Drop
No	Yes	N/A	N/A	Drop
Yes (zone 1)	Yes (zone 2)	No	N/A	Drop
Yes (zone 1)	Yes (zone 2)	Yes	No	Drop
Yes (zone 1)	Yes (zone 2)	Yes	Yes	Policy actions

* A zone pair must have different zones as the source and destination.

Zone-Based Policy Firewall: Rules for Router Traffic

The rules for a zone-based policy firewall are different when the router is involved in the traffic flow, whether as the source of traffic or the destination. A zone-based policy firewall is used to control router administration where the router is the destination. [Table 10-2](#) illustrates various scenarios that involve traffic in or out of the router.

Table 10-2. Zone-Based Policy Firewall: Rules for Router Traffic

Source Interface Member of Zone?	Destination Interface Member of Zone?	Zone Pair Exists?	Policy Exists?	Result
Router (self)	Yes	No	N/A	Pass
Router (self)	Yes	Yes	No	Pass
Router (self)	Yes	Yes	Yes	Policy actions
Yes	Router (self)	No	N/A	Pass
Yes	Router (self)	Yes	No	Pass
Yes	Router (self)	Yes	Yes	Policy actions

When an interface is configured to be a zone member, the hosts connected to the interface are included in the zone, but traffic flowing to and from the interfaces of the router is not controlled by the zone policies. Instead, all the IP interfaces on the router are automatically made part of the self zone when a zone-based policy firewall is configured. To limit IP traffic moving to the IP addresses of the router from the various zones on a router, policies must be applied to block, allow, or inspect traffic between the zone and the self zone of the router, and vice versa. If there are no policies between a zone and the self zone, all traffic is permitted to the interfaces of the router without being inspected.

If desired, you can define a policy using the self zone as either the source or destination zone. The self zone is a system-defined zone. It does not require any interfaces to be configured as members. A zone pair that includes the self zone, along with the associated policy, applies to traffic that is directed to the router or to traffic that the router generates. It does not apply to traffic traversing the router.

The following are additional rules for zone-based policy firewalls that govern interface behavior when the router is involved in the traffic flow:

- All traffic to and from a given interface is implicitly blocked when the interface is assigned to a zone, except traffic to or from other interfaces in the same zone, and traffic to any interface on the router.
- All the IP interfaces on the router are automatically made part of the self zone when a zone-based policy firewall is configured. The self zone is the only exception to the default deny-all policy. All traffic to any router interface is allowed until traffic is explicitly denied.

The following considerations should be weighted when designing Cisco IOS Zone-Based Policy Firewalls:

- An interface can be assigned to one zone and one zone only.
- An interface pair can be assigned one policy and one policy only.
- Consider default traffic flows for interfaces without zones, traffic flows between zones, and traffic flows to or from the router interfaces themselves.
- Inspection actions cannot be applied to the class-default class.
- The default policy action for unclassified traffic is **drop**.

Watch for Misconfiguration

It is considered a misconfiguration to put an **inspect** action from the self zone (the router) to another zone or, vice versa, to put an **inspect** action from a zone to the self zone. The results will vary.

Configuring Basic Interzone Policies Using CCP and the CLI

The CCP Firewall Wizard helps you to implement a firewall. The wizard walks you through creating the firewall by asking you for information about the interfaces on the router, whether you want to configure a DMZ network, and what rules you want to use in the firewall.

CCP Demo Mode

Recall from [Chapter 3](#) that a Demo Mode version of CCP is available for practice. For a helpful tutorial on CCP Demo Mode, check out [Parts I](#) and [II](#) of “Cisco Configuration Professional Demo Mode,” by Doug McKillip of Global Knowledge, the URLs for which are provided in the “References” section toward the end of this chapter.

The configuration scenario shown in [Figure 10-10](#) depicts a simple policy for a remote location that uses an untrusted network for corporate connectivity. We will configure two zones: the inside zone, which includes internal, trusted LAN interfaces, and the outside zone, which represents the untrusted network. The goal is to configure zone-based policy firewall rules that allow selected traffic to travel outbound from inside zone interfaces, while being inspected by the firewall. Inbound traffic, originating at the untrusted, outside interfaces, should be denied and logged.

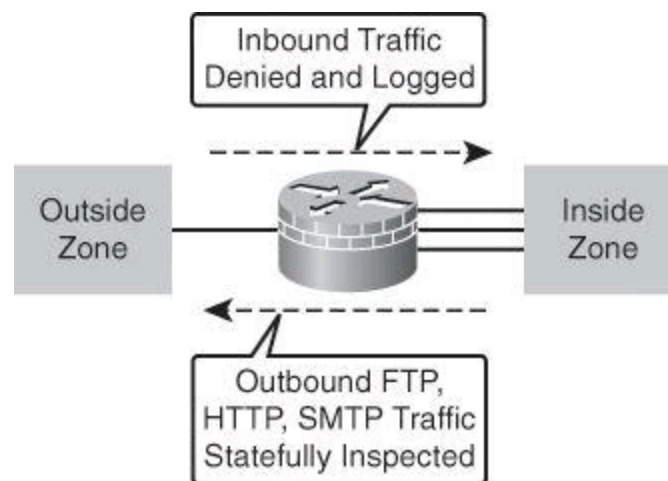


Figure 10-10. Cisco IOS Zone-Based Firewall Configuration Scenario

Using the scenario in [Figure 10-10](#), we will use CCP to configure and verify the firewall, and we will use the CLI to verify the configuration.

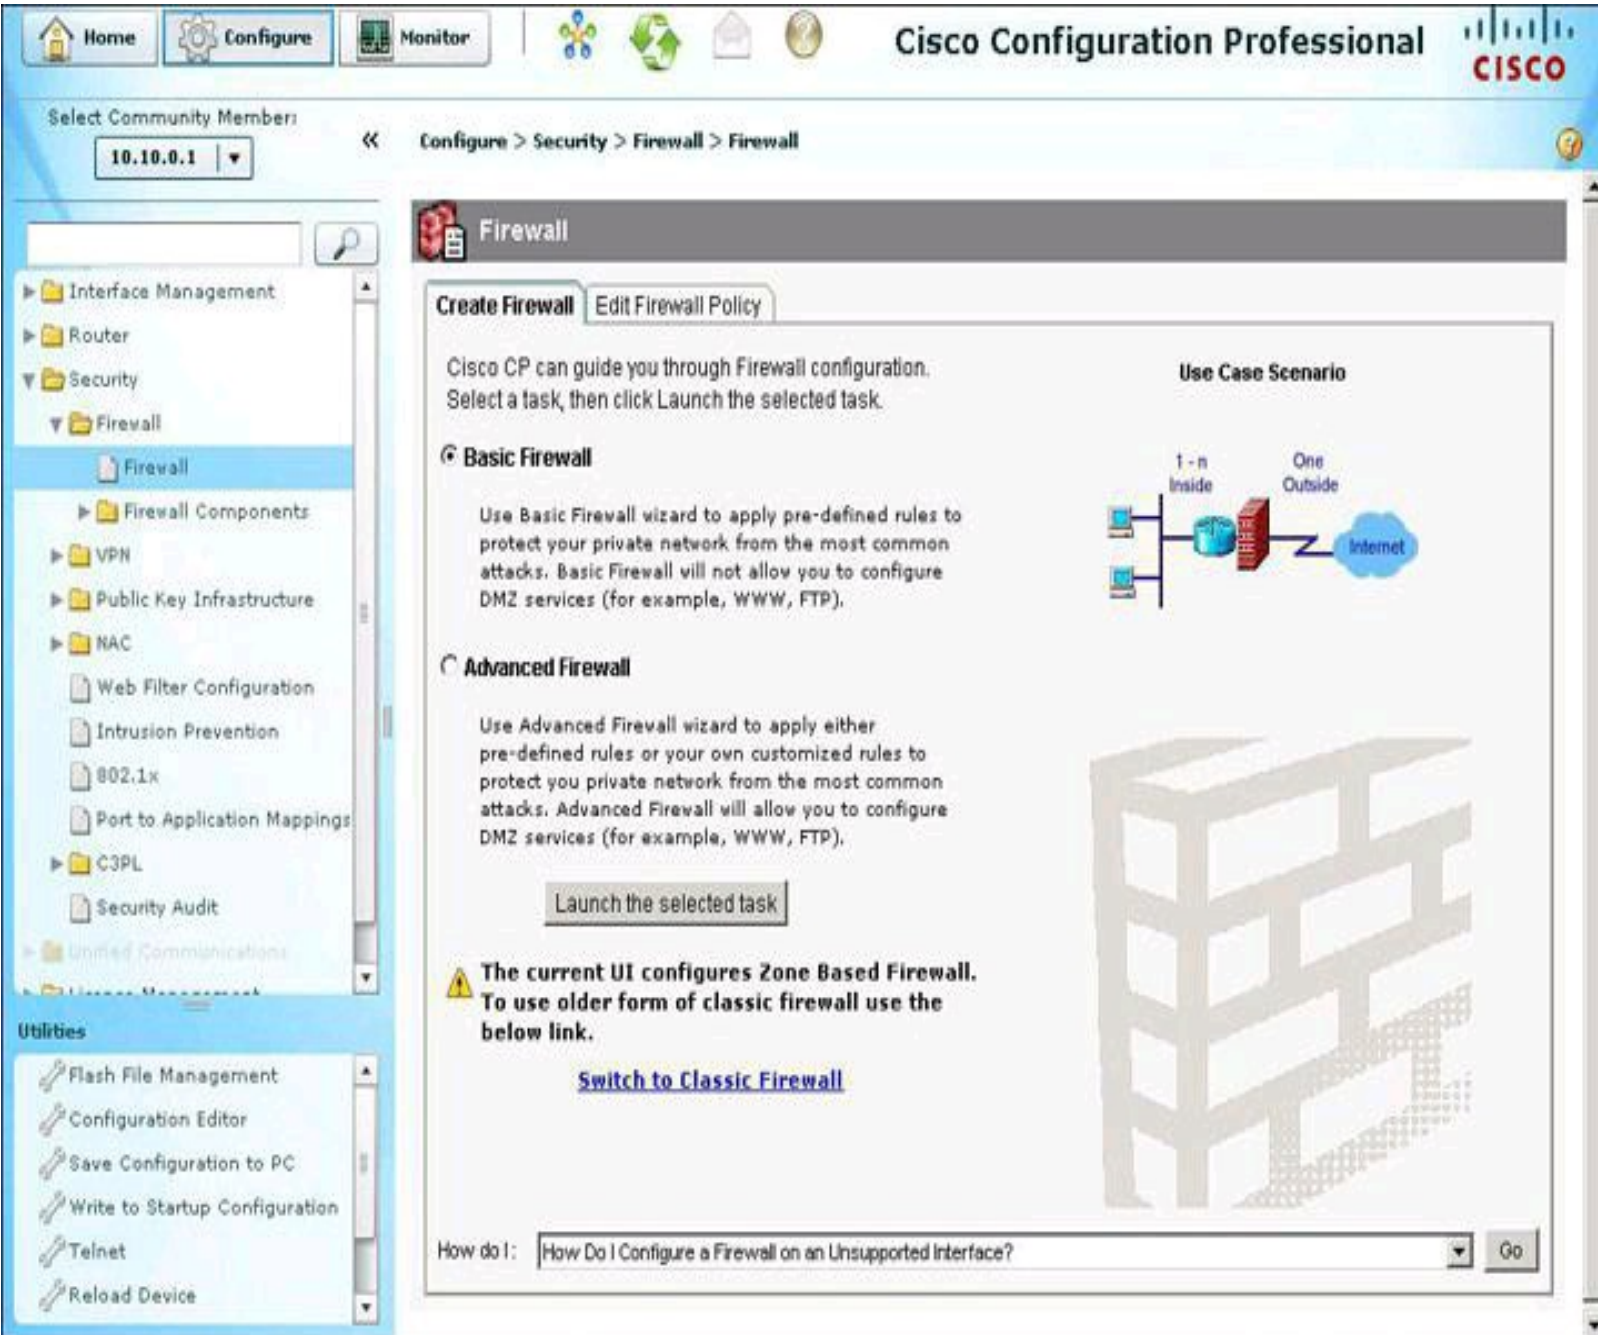
The CCP Firewall Wizard is available in two versions, Basic Firewall and Advanced Firewall. We will use the Basic Firewall wizard. These are the steps to configure a firewall using the Basic Firewall wizard:

- Step 1.** Start the Basic Firewall wizard.
- Step 2.** Select trusted and untrusted interfaces.
- Step 3.** Review and verify the resulting policies.
- Step 4.** (Optional) Enable logging.
- Step 5.** View firewall status and activity.
- Step 6.** (Optional) Modify basic policy objects.
- Step 7.** Verify CLI configuration.

The CCP Basic Firewall wizard will perform the specific operations related to zone-based policy firewalls. The operations include creating zones, creating zone pairs, creating C3PL components (class maps, policy maps, and service policies), and assembling the command hierarchy that implements the zone-based access control policy.

Step 1: Start the Basic Firewall Wizard

The first step is to access and run the Basic Firewall wizard, which you access by navigating to **Configure > Security > Firewall > Firewall**, clicking the **Create Firewall** tab, and then clicking the **Basic Firewall** radio button, as shown in [Figure 10-11](#). To run the wizard, click the **Launch the Selected Task** button. The Basic Firewall wizard asks you to identify the interfaces on your router, and then it uses CCP default access rules and inspection rules to create the firewall. (The Advanced Firewall wizard, on the other hand, shows you the default inspection rules and allows you to use them in the firewall or create your own inspection rules.) CCP will use a default access rule in the firewall.



Select Community Member: 10.10.0.1

Configure > Security > Firewall > Firewall

Firewall

Create Firewall | Edit Firewall Policy

Cisco CP can guide you through Firewall configuration. Select a task, then click Launch the selected task.

Basic Firewall

Use Basic Firewall wizard to apply pre-defined rules to protect your private network from the most common attacks. Basic Firewall will not allow you to configure DMZ services (for example, WWW, FTP).

Advanced Firewall

Use Advanced Firewall wizard to apply either pre-defined rules or your own customized rules to protect your private network from the most common attacks. Advanced Firewall will allow you to configure DMZ services (for example, WWW, FTP).

Launch the selected task

The current UI configures Zone Based Firewall. To use older form of classic firewall use the below link.

[Switch to Classic Firewall](#)

How do I: How Do I Configure a Firewall on an Unsupported Interface? Go

Figure 10-11. Starting the Basic Firewall Wizard

Note

The LAN and WAN configurations must be complete before you can configure a firewall.

Step 2: Select Trusted and Untrusted Interfaces

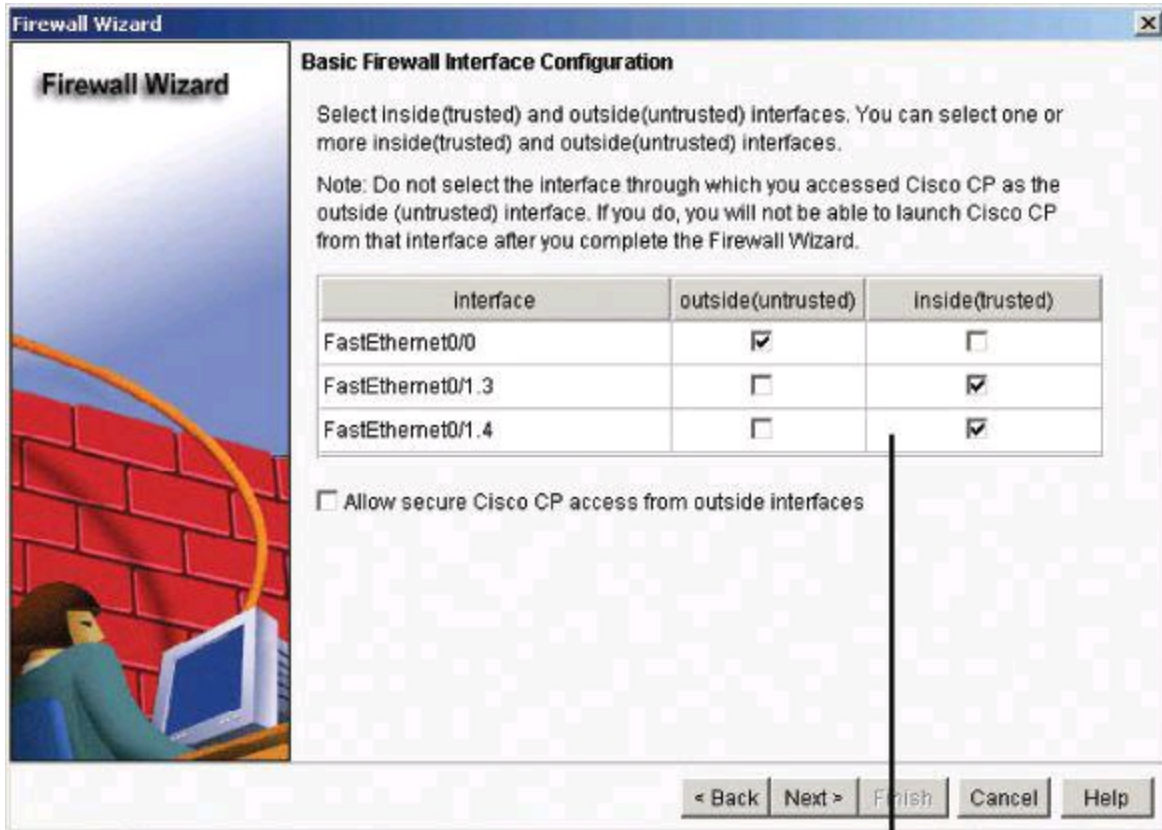
The first step of the Basic Firewall wizard is to identify the interfaces on the router so that the firewall will be applied to the correct interfaces, as shown in [Figure 10-12](#). The following interface categories are available:

- **Outside (untrusted) interface:** Select the router interface that is connected to the Internet or to your organization's WAN.

Note

Do not select the interface through which you accessed CCP as the outside (untrusted) interface. Doing so will cause you to lose your connection to CCP. Because it will be protected by a firewall, you will not be able to launch CCP from the outside (untrusted) interface after the Firewall Wizard completes.

- **Inside (trusted) interfaces:** Check the physical and logical interfaces connecting to the LAN. You can select multiple interfaces.



Interfaces are added to automatically created zones.

Figure 10-12. Selecting Trusted and Untrusted Interfaces

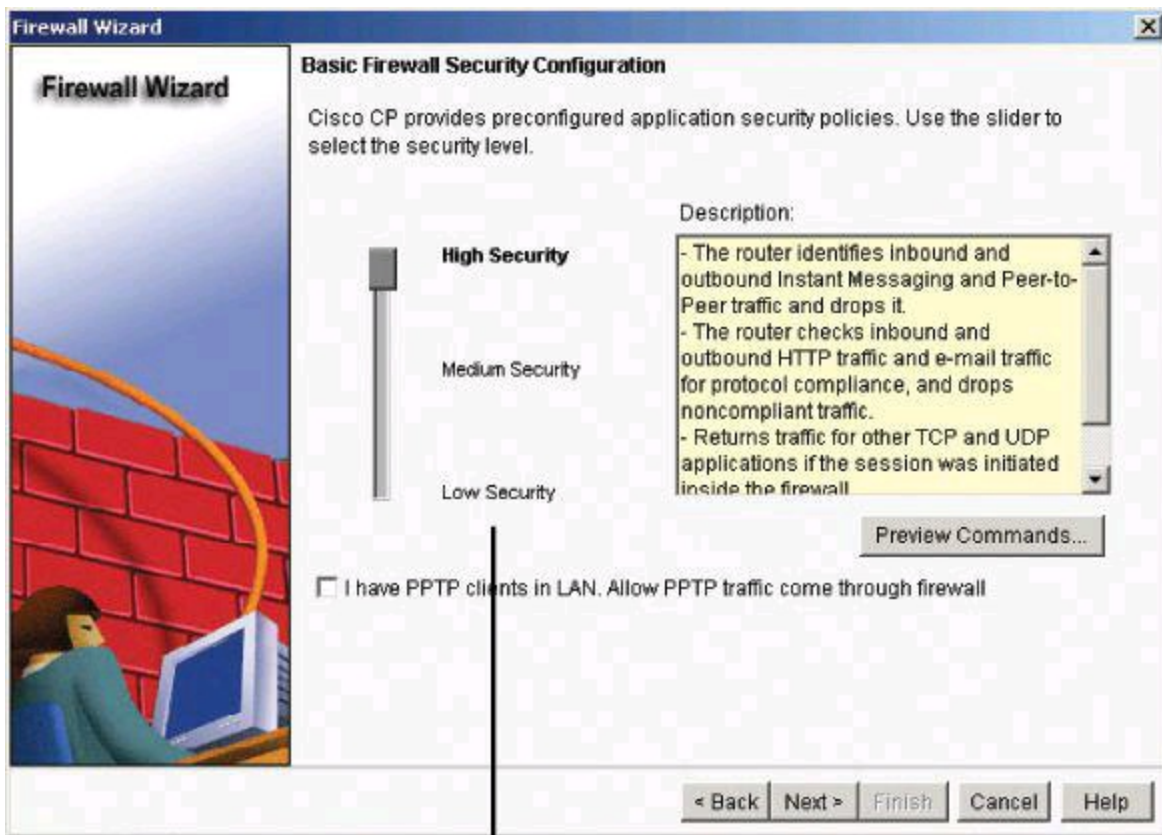
The first page of the wizard also has an Allow Secure Cisco CP Access from Outside Interfaces check box. Check this box if you want users outside the firewall to be able to access the router using CCP. The wizard will display a screen that allows you to specify a host IP address or a network address. The firewall will be modified to allow access to the address you specify. If you specify a network address, all hosts on that network will be allowed through the firewall.

Click **Next** when you are finished identifying interfaces.

The next screen of the wizard allows you to specify more lenient or strict firewall policies to match your requirements, as shown in [Figure 10-13](#). Three levels are available, implementing the following policies:

- High Security
 - The router identifies inbound and outbound instant messaging and peer-to-peer traffic and drops it.

- The router checks inbound and outbound HTTP traffic and email traffic for protocol compliance, and drops noncompliant traffic.
 - The router returns traffic for other TCP and UDP applications if the session was initiated inside the firewall.
 - Choose this option if you want to prevent use of these applications on the network.
- Medium Security
 - The router identifies inbound and outbound instant messaging and peer-to-peer traffic, and checks inbound and outbound HTTP traffic and email traffic for protocol compliance.
 - The router returns TCP and UDP traffic on sessions initiated inside the firewall.
 - Choose this option if you want to track use of these applications on the network.
- Low Security
 - The router does not identify application-specific traffic.
 - The router returns TCP and UDP traffic on sessions initiated inside the firewall.
 - Choose this option if you do not need to track use of these applications on the network.



Customizable policies result from different security levels.

Figure 10-13. Defining Security Levels for the Policy

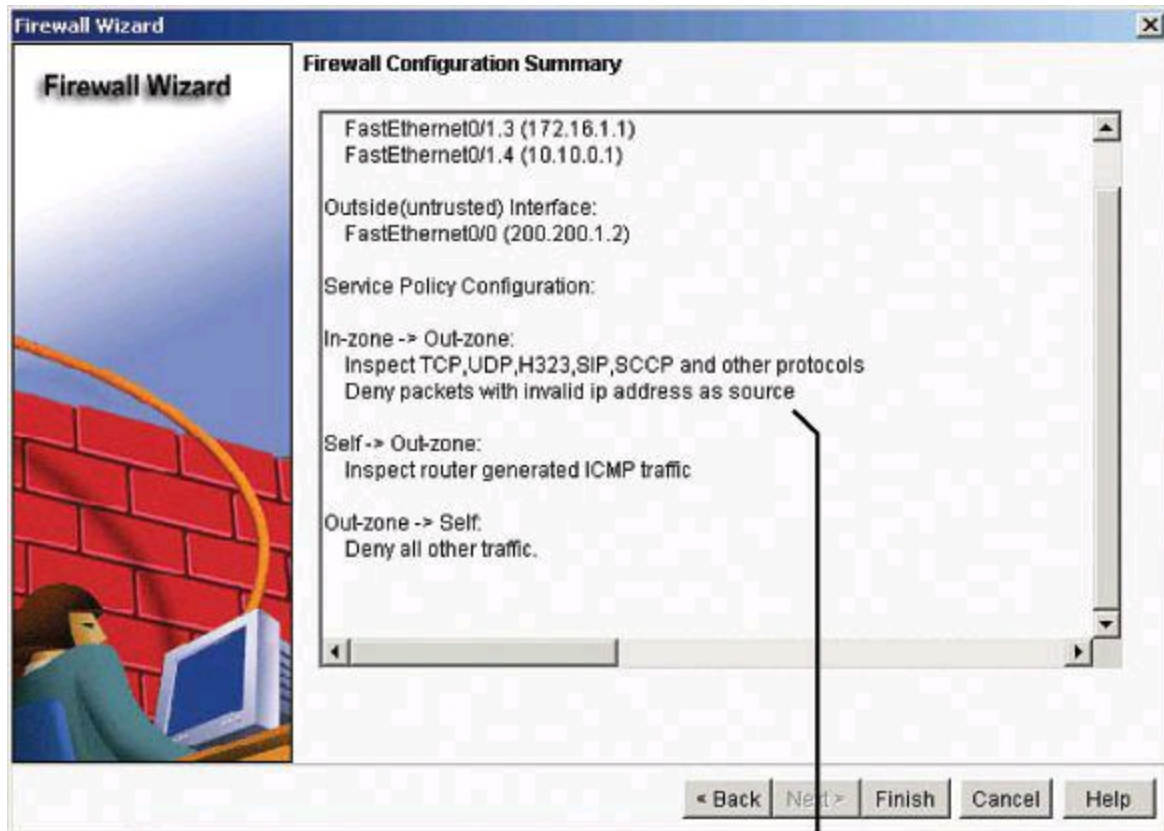
High Security Setting

The High Security setting enables many deep packet inspection features. Some 300

commands are then sent to the router to secure it. Use prudence with the High Security setting because it takes drastic steps to secure your firewall and the traffic going through it. Test amply in a lab environment before implementing it in a production environment. Be sure to preview the commands during your test, and to research them prior to implementing.

Step 3: Review and Verify the Resulting Policies

Click **Next**, and the wizard ends by displaying the configuration and settings, as shown in [Figure 10-14](#). Click **Finish** to continue.



Automatic policy is created, allowing all outbound traffic and blocking all inbound traffic.

Figure 10-14. Reviewing the Wizard Configuration Resulting from Selecting the Low Security Setting

Verifying and Tuning the Configuration

Firewall verification can be accomplished using several CCP tools. You can view the configured policies by navigating to **Configure > Security > Firewall > Firewall** and clicking the **Edit Firewall Policy** tab, as shown in [Figure 10-15](#). Use this tab to view the access and inspection rules in a context that displays the interfaces that the rules are associated with. You can also use it to modify the access and inspection rules that are displayed.

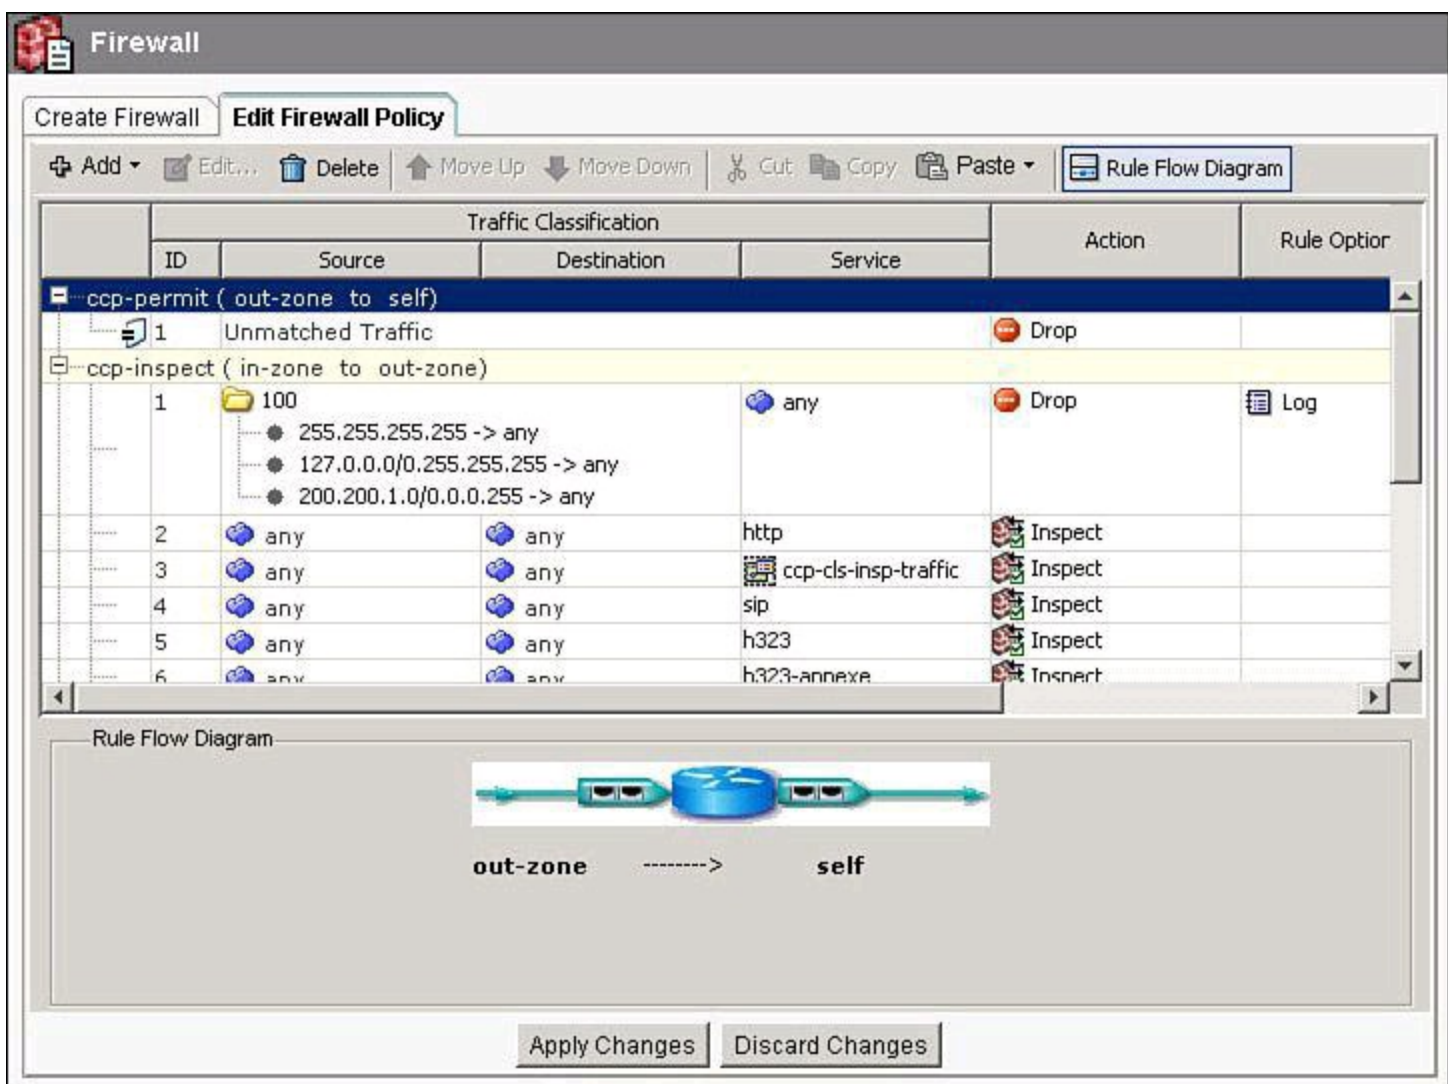


Figure 10-15. Verifying Firewall Policies Using CCP

Unmatched Traffic

Traffic that does not match any inspection rules which are created dynamically from allowed outbound traffic may be allowed inbound in a non-stateful manner, as long as the ACLs assigned to the interfaces permit the traffic bidirectionally.

Once the configuration is completed using CCP, you will want to turn on logging to monitor the status and activities of the firewall by viewing the logs. Depending on your findings, you might want to modify the zone-based firewall configuration.

Step 4: Enabling Logging

Activity on the firewall is monitored through the creation of log entries. If logging is enabled on the router, then whenever an access rule that is configured to generate log entries is invoked—for example, a connection is attempted from a denied IP address—a log entry is generated and can be viewed in monitor mode.

Before you attempt to view the firewall log, you must enable logging by navigating to **Configure > Router > Logging**. Then, follow these steps:

Step 1. On the **Additional Tasks > Logging** screen, shown in the background of [Figure 10-](#)

16, select **Logging to Buffer** and click **Edit** to open the Logging dialog box, shown in the foreground.

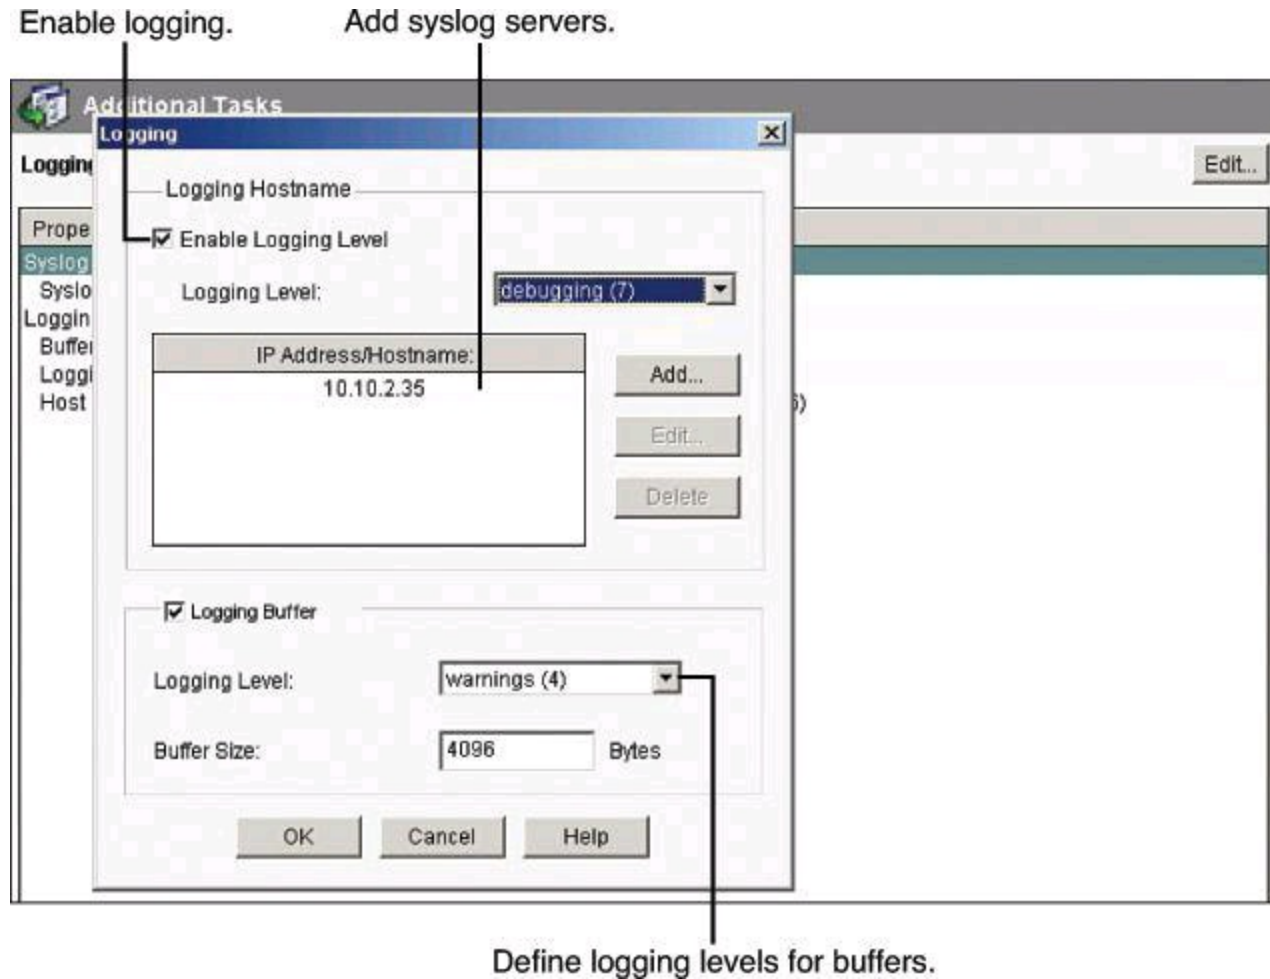


Figure 10-16. Enabling Logging to the Buffer

Step 2. In the Buffer Size field, enter the amount of router memory that you want to use for a logging buffer. The default value is 4096 bytes. A larger buffer will store more log entries, but you must balance your need for a larger logging buffer against potential router performance issues. Leave the other settings to their default values.

Step 3. Click **OK**.

Logging to the Console and Logging Dropped Packets

As explained in [Chapter 4](#), it is preferable to send logging results to vty sessions rather than to the console port. The console port performs at a much lower speed than vty lines. If you decide nevertheless to send logs to the console port, use a lower level than debugging. Use at most the informational logging level.

Earlier in this chapter, we saw that it is possible to log packets that are dropped based on the policy. Though we haven't looked at the CLI for ZBF, the command would look like this: Router(config-pmap-c)# **drop log**. This log action is available for packets that are dropped because they match a policy that *explicitly* says to drop the packets. Malformed or noncompliant packets dropped by the firewall, such as those with invalid UDP Header length, incompatible TCP options, and so forth, are not logged. To see packets dropped by the firewall, use the global command **ip inspect log drop-pkt**.

Step 5: Verifying Firewall Status and Activity

You can now view the firewall status and logs at **Monitor > Security > Firewall Status**. In the firewall statistics display, shown in [Figure 10-17](#), you can verify that your firewall is configured and view how many connection attempts have been denied. The table shows each router log entry that is generated by the firewall, including the time and the reason that the log entry was generated.

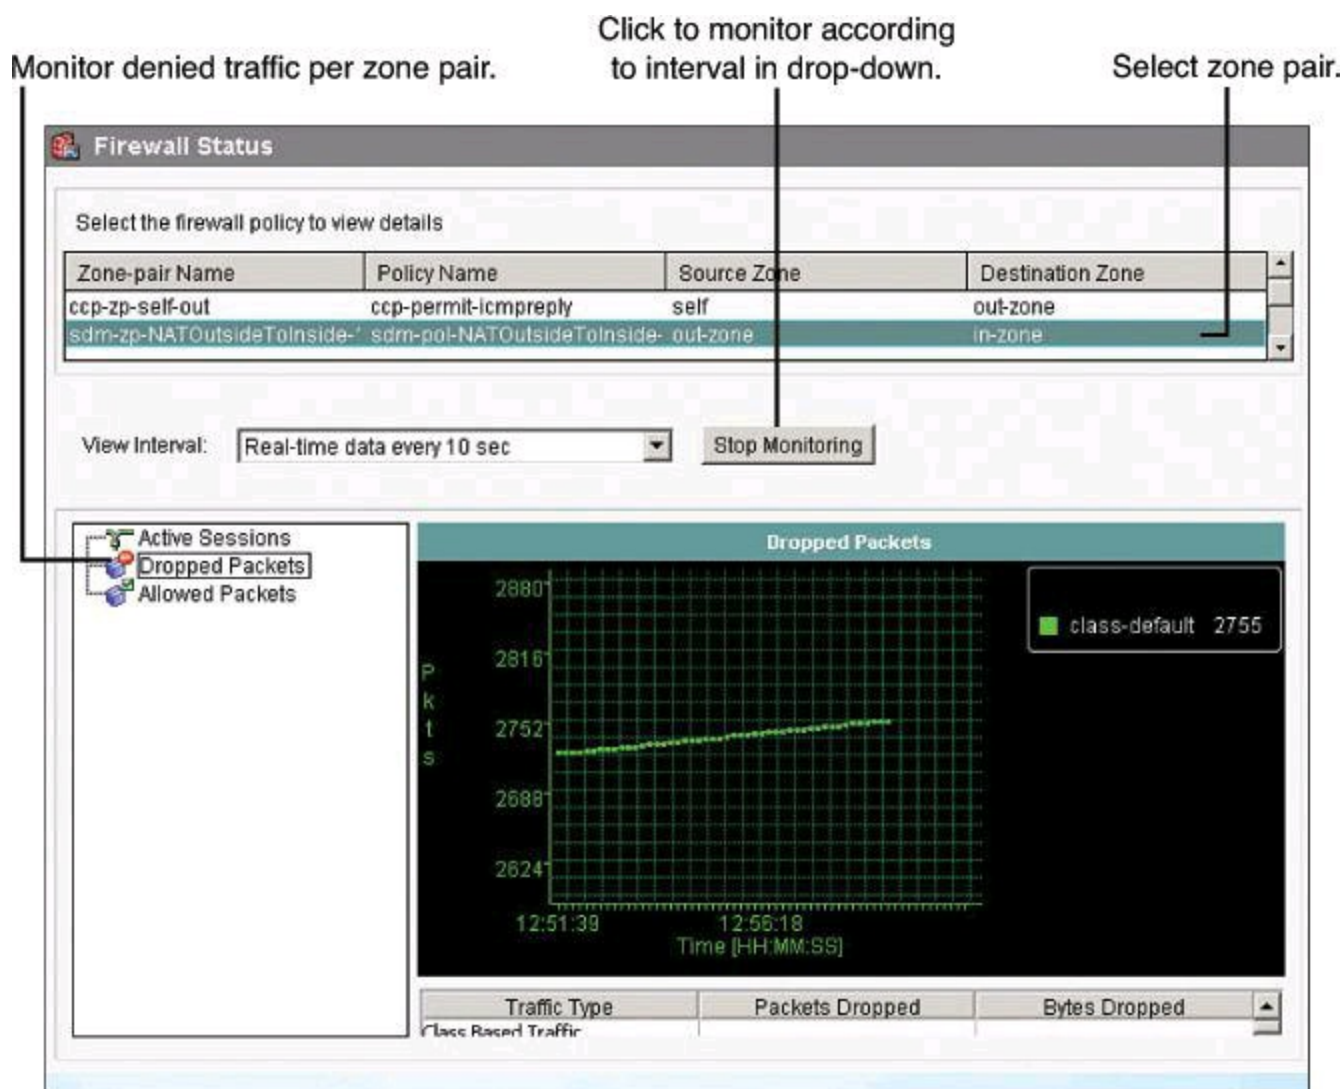


Figure 10-17. Verifying Firewall Status and Activity

Note

The default rule that is created by the Basic Firewall wizard enables logging for denied actions, and this log will display those messages. If logging is not enabled at the policy level, you will have to modify the policy maps to enable logging as an additional action.

Viewing Firewall Logs

Firewall logs can also be viewed by navigating to **Monitor > Router > Logging > Firewall Log**. The log entries are determined by log messages generated by the firewall. In order for the firewall to generate log entries, you must configure the following:

- Logging for the router. To obtain firewall logging messages, you must configure a logging level of debugging (7).

- Individual access rules to generate log messages when they are invoked.

The top-attacks table below the View drop-down menu displays the top attack entries. You can use the View drop-down menu to switch from a port-based report to an attacker-based report, which displays the IP addresses of the top attackers.

Note

Log entries are not refreshed automatically. You must click Update to refresh the view and display newly generated entries.

Step 6: Modifying Zone-Based Firewall Configuration Objects

The optional step of modifying zone-based policy firewall base objects uses the Configure > Security > Firewall > Firewall Components > Zones navigation option. This step allows you to add or change zones; add or modify zone pairs, as shown in the upper-right side of [Figure 10-18](#); or add or modify C3PL objects (class maps, policy maps, and service policies), as shown in the bottom portion of [Figure 10-18](#).

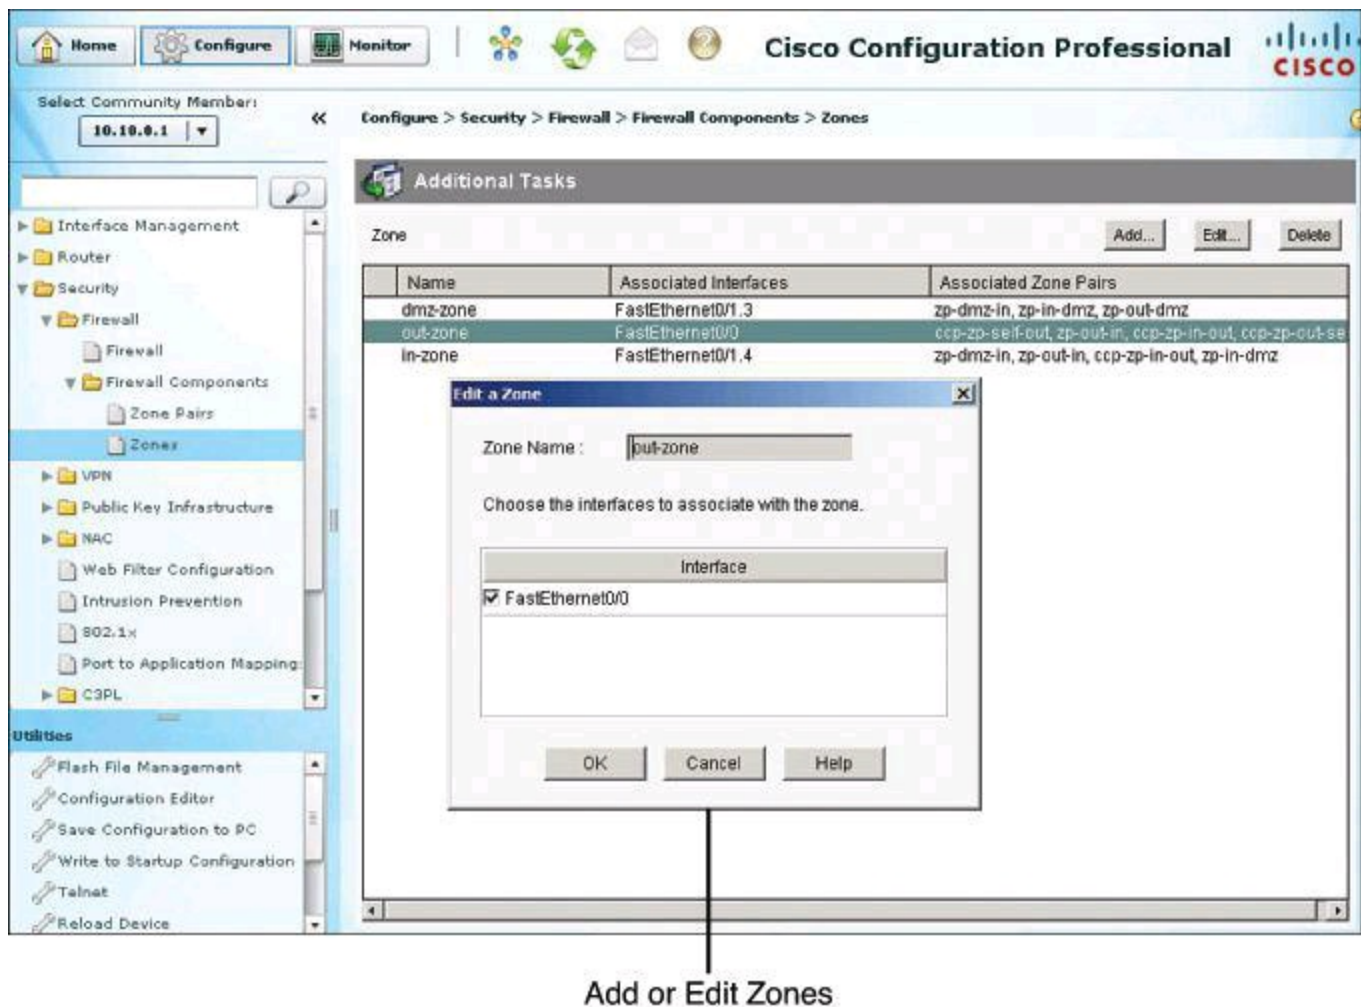


Figure 10-18. Modifying ZBF Configuration Objects

General Steps to Create ZBF

ZBF policies are built from the bottom up. You first define the class map, which specifies interesting traffic for the policy maps. You then define the policy maps, specifying which

actions to take on the traffic defined in the class maps. Then, the zone pairs reference both the policy maps and the zones. The following list shows the sequence of what is created and referenced:

1. ACL to identify the traffic
 2. Zones
 3. Class map
 4. Policy map (actions can be **inspect**, **pass**, and **drop**, and dropped traffic can also be logged)
 5. Zone pair (policy map + zones)
-

Step 7: Verifying the Configuration Using the CLI

[Example 10-1](#) shows an example of what the resulting CLI commands might look like after configuring a Cisco IOS Zone-Based Policy Firewall that uses two interfaces and the default inspection parameters. The breakdown is as follows:

First, the class map named OUTBOUND-PROTOCOLS is created. It identifies three protocols that are to be inspected: HTTP, SMTP, and FTP.

A policy map named ACCESS-POLICY is created. It applies stateful inspection to the protocols that are listed in the OUTBOUND-PROTOCOLS class map.

Two zones, named PRIVATE and INTERNET, are created. FastEthernet 0/0 is made a member of the PRIVATE zone and FastEthernet 0/1 is made a member of the INTERNET zone.

Finally, a zone pair named PRIV-TO-INTERNET is created with a source zone of PRIVATE, a destination zone of INTERNET, and the policy map ACCESS-POLICY applied to it.

Example 10-1. Commands of a Basic Cisco IOS Zone-Based Policy Firewall Configuration

[Click here to view code image](#)

```
class-map type inspect match-any OUTBOUND-PROTOCOLS
  match protocol http
  match protocol smtp
  match protocol ftp
!
policy-map type inspect ACCESS-POLICY
  class type inspect OUTBOUND-PROTOCOLS
  inspect
!
zone security PRIVATE
zone security INTERNET
!
interface fastethernet 0/0
  zone-member security PRIVATE
!
interface serial 0/0/0
```



```

zone-member security INTERNET
!
zone-pair security PRIV-TO-INTERNET source PRIVATE destination INTERNET
service-policy type inspect ACCESS-POLICY
!

```

Configuring NAT Services for Zone-Based Firewalls

Network Address Translation (NAT) can be added to zone-based firewall configurations by applying NAT rules at the interface level for interfaces that belong to zones. The configuration scenario that is shown in [Figure 10-19](#) illustrates an example of outbound dynamic Port Address Translation (PAT), where trusted hosts on the inside zone are translated dynamically, using PAT, to the IP address on the router's untrusted interfaces, in this example FastEthernet 0/1.

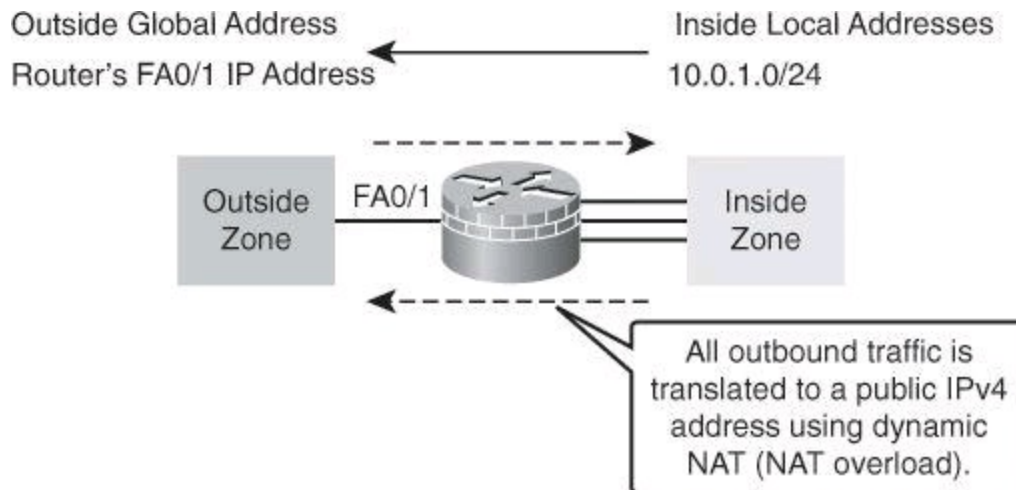


Figure 10-19. NAT with ZBF Configuration Scenario

There are three main steps to configure a NAT with Cisco IOS zone-based firewall:

Step 1. Run the Basic NAT wizard.

Step 2. Select NAT interfaces:

- Outside interface with global IP address
- Inside interface with original IP address

Step 3. Verify the configuration.

Step 1: Run the Basic NAT Wizard

The first step is to access and run the Basic NAT wizard by navigating to **Configure > Router > NAT** and clicking the **Create NAT Configuration** tab, as shown in [Figure 10-20](#). The options include the following:

- **Basic NAT:** Choose the Basic NAT wizard if you want to connect your network to the Internet (or the outside), and your network has hosts but no servers. Look at the sample diagram that appears to the right when you choose Basic NAT. If your network is made up only of PCs that require access to the Internet, choose **Basic NAT** and click **Launch**.
- **Advanced NAT:** Choose the Advanced NAT wizard if you want to connect your network to the Internet (or the outside), and your network has hosts and servers, and the servers must be accessible to outside hosts (hosts on the Internet). Look at the sample diagram that

appears to the right when you choose Advanced NAT.

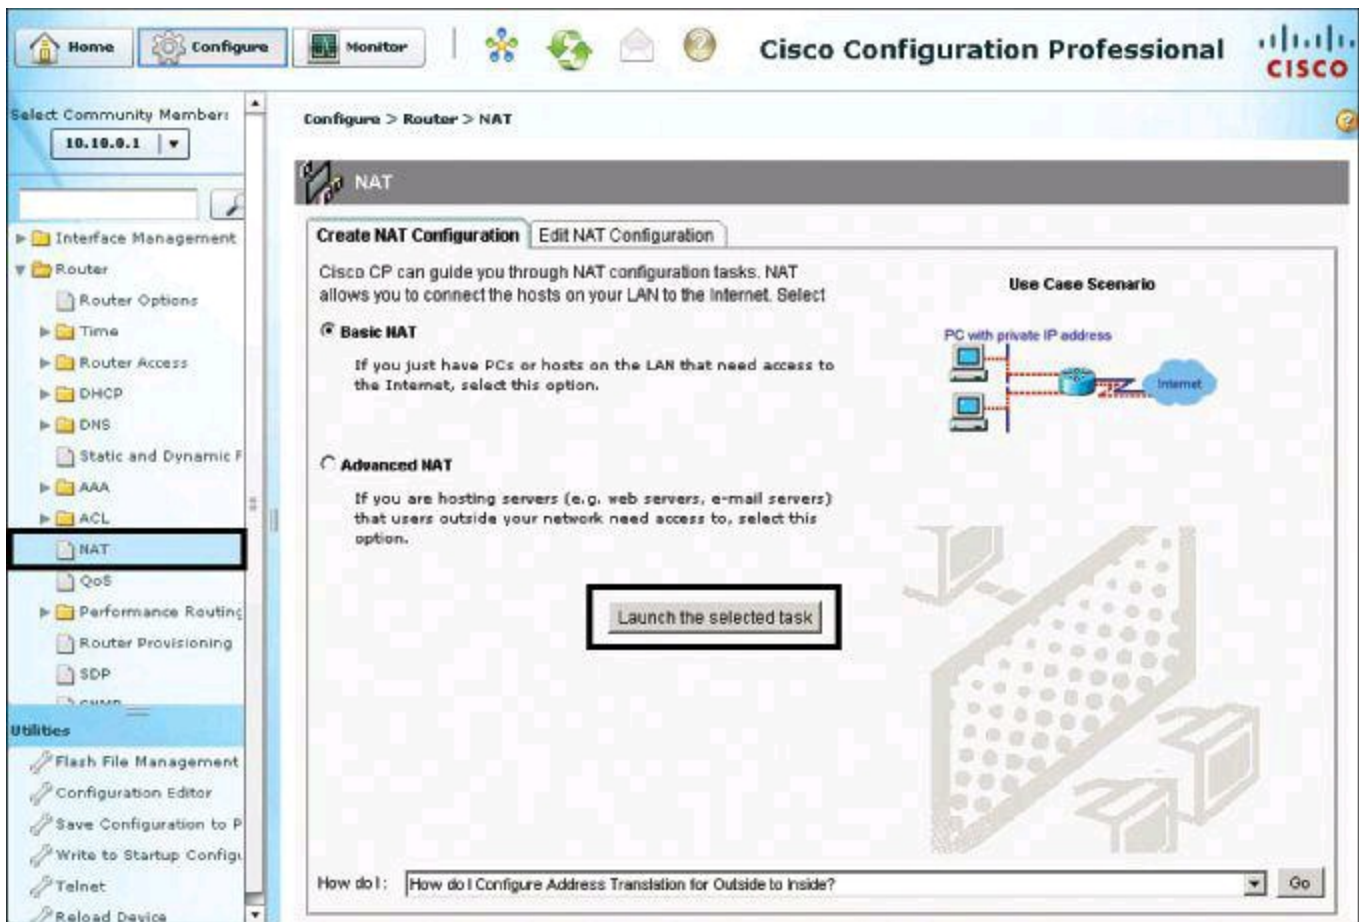


Figure 10-20. Starting the Basic NAT Wizard

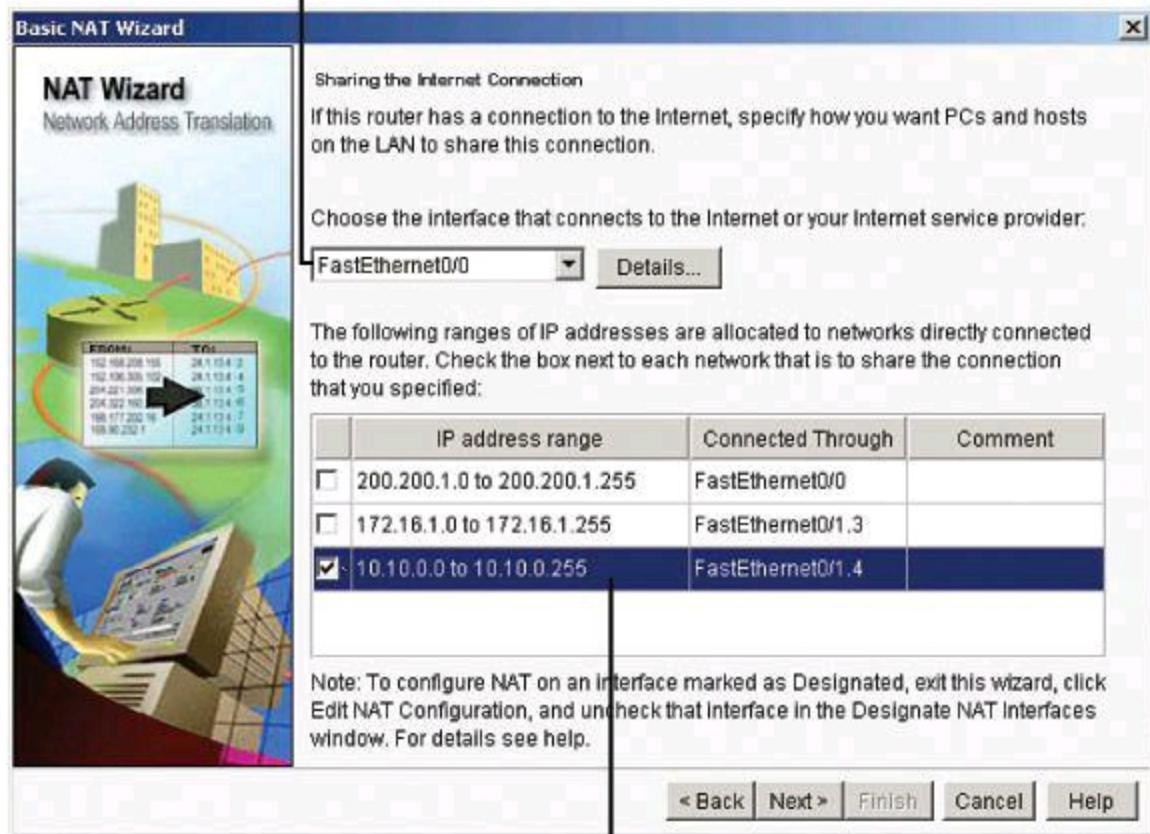
Click the **Launch the Selected Task** button to launch the wizard.

Step 2: Select NAT Inside and Outside Interfaces

The only configuration that is needed in the Basic NAT wizard is the definition of original versus translated interfaces:

- From the drop-down menu, shown in [Figure 10-21](#), choose the outside interface, which connects to the Internet or untrusted network. This is the router's WAN interface.

Select the Internet-facing interface from the list (outside).



Click to select the trusted networks to be translated.

Figure 10-21. Starting the Basic NAT Wizard

- From the list of available networks, shown with check boxes, choose which networks will share the WAN interface in the NAT configuration. To choose a network, check its check box in the list of available networks.

Note

Do not choose a network that is connected to the WAN interface set up in this NAT configuration. Remove that network from the NAT configuration by unchecking its check box.

The list of available networks shows the following information for each network:

- IP address allocated to the network
- Network LAN interface
- Comments entered about the network

To remove a network from the NAT configuration, uncheck its check box.

Note

If Cisco Configuration Professional detects a conflict between the NAT configuration and an existing virtual private network (VPN) configuration for the WAN interface, it will inform you with a dialog box after you click Next. The only configuration that is needed in

the Basic NAT wizard is the definition of original versus translated interfaces.

Click **Next**, and the wizard ends with a summary of the configuration, as shown in [Figure 10-22](#). Click **Finish**.

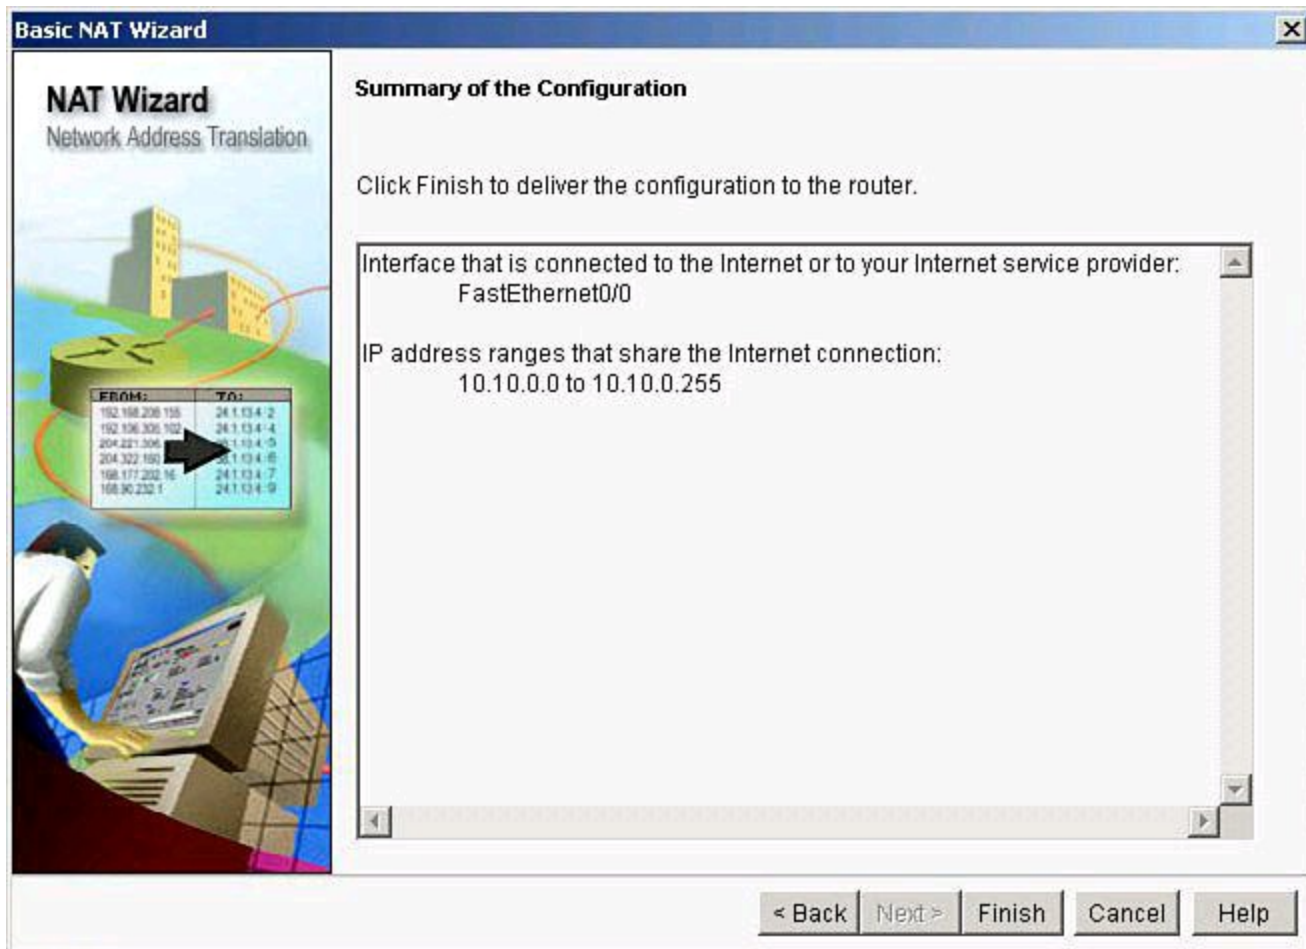


Figure 10-22. Finishing the Wizard

Step 3: Verify NAT with CCP and the CLI

Verification can be accomplished by looking at the various NAT rules, found by navigating to **Configure > Router > NAT** and clicking the **Edit NAT Configuration** tab, as shown in [Figure 10-23](#).

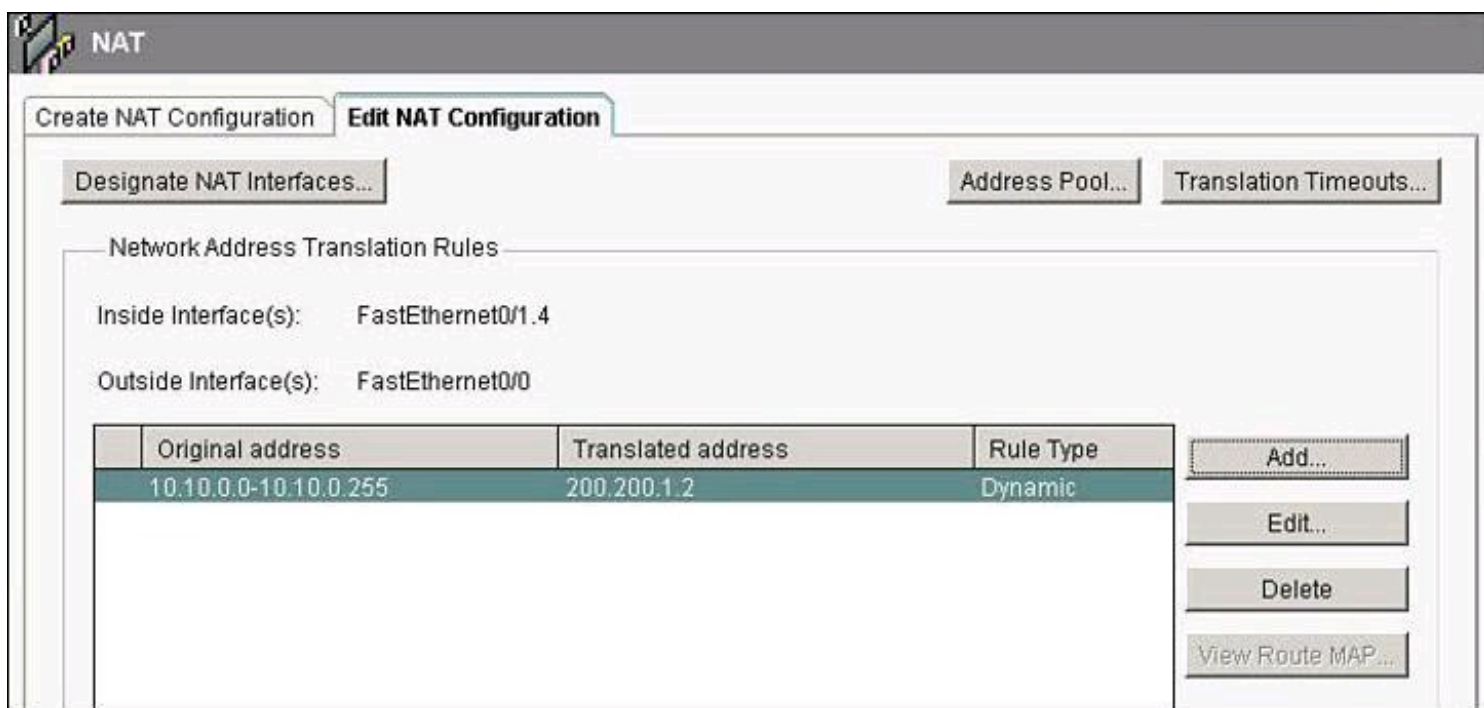


Figure 10-23. Verifying NAT Configuration with CCP

The resulting CLI configuration is shown in [Example 10-2](#). Notice the commands **ip nat outside** and **ip nat inside**, which define the location of original and translated IPv4 addresses. Also notice that the original addresses are matched using an ACL, list number 1 in [Example 10-2](#).

Example 10-2. NAT CLI Configuration

[Click here to view code image](#)

```

interface FastEthernet0/0
  ip address 200.200.1.2 255.255.255.0
  ip nat outside
  !
interface FastEthernet0/1.4
  ip address 10.10.0.1 255.255.255.0
  ip nat inside
  !
ip nat inside source list 1 interface FastEthernet0/0 overload
  !
access-list 1 permit 10.10.0.0 0.0.0.255
  !

```

The **show ip nat translations** command, shown in [Example 10-3](#), displays the current translation for live traffic, including parameters such as inside local, inside global, outside local, and outside global addresses, as well as local and global ports.

Example 10-3. Current Translation for Live Traffic

[Click here to view code image](#)

```

Router# show ip nat translations

```

Pro	Inside global	Inside local	Outside local	Outside
global				
TCP				
200.200.1.51:1050	10.10.10.20:1050	75.75.75.750:23	172.16.100.10:23	
TCP				
200.200.1.52:1776	10.10.10.10:1776	150.150.1.40:25	150.150.1.40:25	

ZBF: Work in Progress

Cisco keeps coming out with improvements to ZBF. One of the latest improvements, which came out with Cisco IOS Release 15, is intrazone support, which allows for zone configuration to include users both inside and outside a network. This allows for traffic inspection between users belonging to the same zone but different networks. Check Cisco.com regularly for updates to ZBF.

Cisco ASA Firewall

Cisco ASA 5505 Adaptive Security Appliance (ASA) implements a rich set of security technologies and can be effectively implemented as a perimeter firewall using several deployment modes. This section introduces Cisco ASA functionality, features, and underlying technologies. This section will demonstrate how to configure the Cisco ASA 5505 model for basic connectivity using Cisco Adaptive Security Device Manager (ASDM).

Cisco ASA 5500 Series are easy-to-deploy solutions that integrate a world-class firewall, Cisco Unified Communications (voice and video) security, Secure Sockets Layer (SSL) and IP Security (IPsec) VPNs, intrusion prevention systems (IPS), and content security services in a flexible, modular product family. The Cisco ASA 5500 Series provides intelligent threat defense and secure communications services that stop attacks before they affect business continuity. Designed to protect networks of all sizes, Cisco ASA 5500 Series security appliances enable organizations to lower their overall deployment and operations costs while delivering comprehensive multilayer security.

Cisco ASAs scale to meet a range of requirements and network sizes. At the time of publication, the Cisco ASA 5500 Series includes the Cisco ASA 5505, 5510, 5512-X, 5515-X, 5520, 5525-X, 5540, 5545-X, 5550, and 5555-X, the four different footprints of the 5585-X (each one offering a different footprint of the Security Services Processor, or SSP), and the ASA Security Services Modules.

Stateful Packet Filtering and Application Awareness

The Cisco ASA security appliance is fundamentally a stateful packet filter with application inspection and control, with a rich set of additional integrated software and hardware features that enable you to expand its functionality beyond those fundamental filtering mechanisms.

The heart of the Cisco ASA is an application-aware stateful packet inspection algorithm, which controls flows between networks that are controlled by the security appliance. The connection-oriented stateful packet inspection algorithm design creates session flows based on source and destination addresses and ports. The information is kept in a state table, as shown in [Figure 10-24](#). The stateful packet inspection algorithm randomizes TCP sequence numbers, port numbers, and additional TCP flags before completion of the connection. This function is always in operation,

monitoring return packets to ensure that they are valid, and it transparently allows any additional dynamic transport layer sessions of the same application session. This fundamental function of the security appliance allows you to minimize host and application exposure and only allow the minimal traffic on OSI network and transport layers that is required to support your applications.

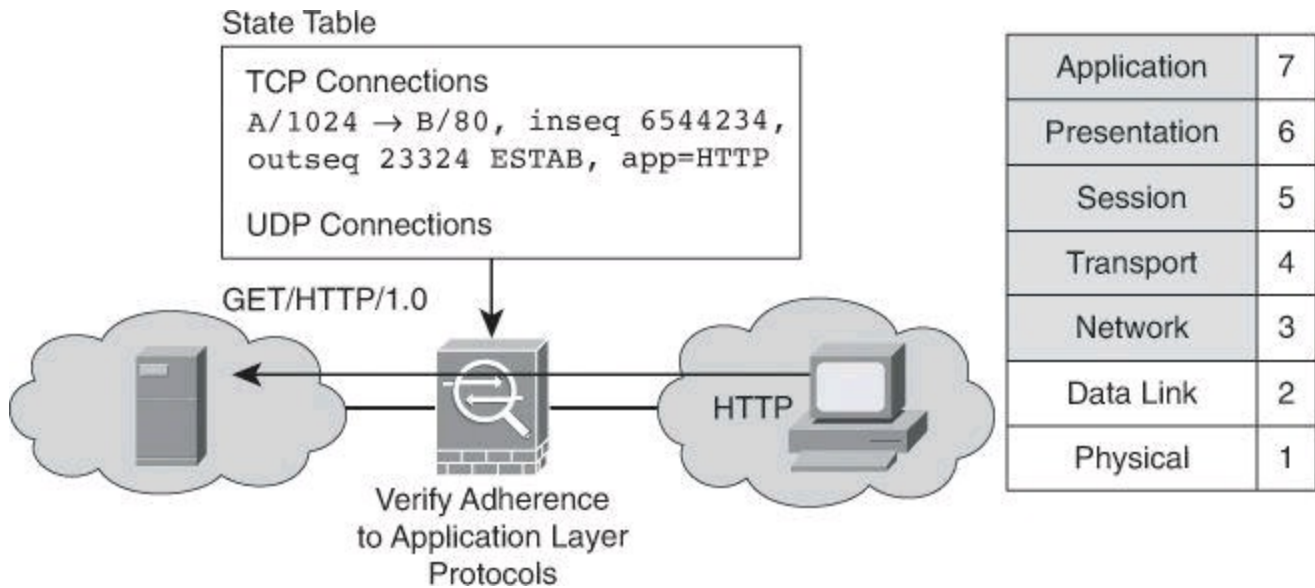


Figure 10-24. State Table Created for All Inspected Traffic

Interface Names

You will also notice in [Figure 10-24](#) that interfaces receive names. With the Cisco ASA, to pass traffic, an interface needs to have a name (configured with the **nameif** command), and if you wish to route IP traffic, that interface also needs an IP address, which we will cover later. The **nameif inside** command assigns automatically the security level 100. This topic will be discussed later in this section.

Network Services Offered by the Cisco ASA 5500 Series

The Cisco ASA 5500 series offers services in addition to firewalling. Some of those additional services, examined in this section, are NAT, DHCP server, and routing.

Network Address Translation

The Cisco ASA security appliance includes support for many different varieties of NAT.

Cisco ASA Pre-8.3 to 8.3 NAT

Cisco has significantly changed the way NAT is done, starting with Cisco ASA version 8.3. If you have invested time and effort learning NAT pre-8.3, read the document “ASA Pre-8.3 to 8.3 NAT Configuration Examples,” found at <https://supportforums.cisco.com/docs/DOC-9129>. There is also a good video to help you make the transition, “Cisco ASA Version 8.3 and 8.4 NAT Configuration Example,” available at <https://supportforums.cisco.com/docs/DOC-12324>. The terminology has changed with NAT 8.3. As an example, if you run Packet Tracer (a troubleshooting tool covered later in this chapter), you might notice the term UN-NAT, which refers to a packet

using its real IP address and port (typically the private IP address). Another significant change with version 8.3 is that NAT exemption, introduced with Cisco ASA 7.0, is no longer used. For detailed information about Cisco ASA 8.3 NAT, refer to *CCNP Security FIREWALL 642-618 Official Cert Guide* (Cisco Press, 2012).

Another substantial difference with 8.3 is that when creating an ACL, you must always use the real IP address, whereas in pre-8.3, the IP address had to be either the real IP address or the NATted IP address, depending on where the ACL was applied.

Note that for the purpose of passing the IINSv2 certification exam, knowing the material presented in this section suffices.

These different varieties of NAT allow for flexible deployment of NAT services:

- **Inside and outside NAT:** These terms are typically used in reference to the physical location of the hosts that require and use NAT services. If the host resides on the inside network and its address is translated for outbound traffic flows, inside NAT is used. If the host resides on the outside network and its address is translated for inbound traffic flows, outside NAT is used. The terminology of inside and outside NAT has its origins in how the interfaces were named by the Cisco ASA predecessor, the PIX Firewall; this nomenclature is still used today: inside interface and outside interface. The inside interface, as its name implies, is the secure side of the firewall. The outside interface points to the less secure side of the firewall, usually the Internet.
- **Dynamic NAT and PAT:** Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool may include fewer addresses than the real group, in which case source ports are used to map multiple local addresses to one or a few global addresses. When a host that you want to translate accesses the destination network, the Cisco ASA assigns to the host an IP address from the mapped pool. The translation is added only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, although the connection is allowed by an ACL, and the Cisco ASA rejects any attempt to connect to a real host address directly. If all the real addresses are translated to one single IP address instead of a pool of multiple addresses, it is then referred to as Dynamic PAT. On a Cisco ASA firewall, the default dynamic NAT timeout is 3 hours. That is, after 3 hours of no activity seen for a translated address, the entry is removed from the translation table. This timer is adjustable. Note that on the Cisco IOS router, the default idle timeout for translation is 24 hours.
- **Static NAT and PAT:** Static NAT creates a fixed translation of real addresses to mapped addresses. With dynamic NAT and PAT, each host uses a different address or port for each subsequent translation. Because the mapped address is the same for each consecutive connection with static NAT, and a persistent translation rule exists, static NAT allows hosts on the destination network to initiate traffic to a translated host, if an ACL exists that allows it.

The main difference between using dynamic NAT and using a range of addresses for

static NAT is that static NAT allows a remote host to initiate a connection to a translated host (if an ACL exists that allows it), while dynamic NAT does not. You also need an equal number of mapped addresses as real addresses with static NAT.

Static PAT is the same as static NAT, except that it lets you specify the protocol (TCP or UDP) and port for the real and mapped addresses.

- **Policy NAT:** Policy NAT lets you identify real addresses for address translation by specifying the source and destination addresses. You can also optionally specify the source and destination ports. Regular NAT can consider only the source addresses and cannot consider the destination. For example, with policy NAT, you can translate the real address to mapped address A when it accesses server A, and translate the real address to mapped address B when it accesses server B.
- **NAT exemption:** NAT exemption exempts addresses from translation and allows both real and remote hosts to originate connections. NAT exemption lets you specify the real and destination addresses when determining the real traffic to exempt (similar to policy NAT), so you have greater control using NAT exemption than dynamic identity NAT. However, unlike policy NAT, NAT exemption does not consider the ports.

Cisco ASA Connection Table

In addition to the translation table kept by the Cisco ASA, which you can look at with the **show xlate** command, the Cisco ASA maintains a connection table. You can see it with the **show conn** command. A TCP connection has a 1-hour idle timeout. A UDP connection has a 2-minutes idle timeout. These default timeouts are adjustable.

Additional Network Services

The Cisco ASA can provide a DHCP server or DHCP relay services to DHCP clients attached to Cisco ASA interfaces. The DHCP server provides network configuration parameters directly to DHCP clients. DHCP relay passes DHCP requests received on one interface to a DHCP server located behind a different interface. DHCP relay takes a DHCP broadcast and forwards it to the DHCP server located on a different network as a unicast. You can configure a DHCP server on each interface of the Cisco ASA. Each interface can have its own pool of addresses to draw from.

The Cisco ASA security appliance supports Routing Information Protocol (RIP) (versions 1 and 2), Open Shortest Path First (OSPF), and Enhanced Interior Gateway Routing Protocol (EIGRP) dynamic routing protocols to integrate into existing routing infrastructures. Where dynamic routing is not available, the Cisco ASA security appliance can use static routing instead.

Cisco ASA Security Technologies

Many other features and technologies are available in Cisco ASA. The appliances implement a rich set of firewall and security features. Some of these features are enabled by default, while others can be enabled to implement multiple levels of access and threat control and protection. In addition to the basic security features presented in this section, the following functions are available in Cisco ASAs:

- Stateful inspection and application level controls

- ACL packet filtering
- Object groups
- Application Inspection and Control (AIC)
- User-based access control (cut-through proxy)
- Identity firewall
- Session auditing
- Threat control and containment
 - IPS via Cisco ASA Advanced Inspection and Prevention Security Services Module (AIP-SSM) and Advanced Inspection and Prevention Security Services Card (AIP-SSC)
 - Botnet traffic filtering
 - Category-based URL filtering
 - Threat detection (basic, advanced, scanning)
- Network integration
 - Virtualization
 - Security modules
 - IPv6 and multicast support
 - NAT and DHCP services
 - Site-to-site and remote-access IPsec and SSL VPNs
 - Transparent firewall mode
 - IP routing
 - High-availability failover

The majority of these advanced functions and features are beyond the scope of this book and are covered in Cisco Press's *CCNP Security FIREWALL 642-618 Official Cert Guide*.

Cisco ASA Configuration Fundamentals

In Cisco ASA, default and custom access rules are based on interface security levels. Security levels define the level of trustworthiness of an interface, and help implement a layered, defense-in-depth approach to security. The higher the level, the more trusted the interface. The numbers themselves are not really important, but the relationship between the numbers is. In that sense, traffic flows are defined as inbound or outbound like this:

- **Inbound traffic:** Travels from a less trusted interface to a more trusted interface; that is, from a lower security level to a higher security level
- **Outbound traffic:** Travels from a more trusted interface to a less trusted interface; that is, from a higher security level to a lower security level

Each interface must have a security level from 0 (lowest) to 100 (highest). For example, as shown in [Figure 10-24](#), you should assign your most secure network, such as the inside host network, to level 100, while the outside network connected to the Internet can be level 0. Other networks, such as DMZs, can be in between. You can assign interfaces to the same security level, but additional

configuration will be needed for traffic to traverse interfaces of the same security level.

Default Security Policy

Security levels define default behavior and processing of traffic flows. Merely by assigning numbers to interfaces, this policy applies the following:

- Outbound traffic is allowed and inspected by default.
- Returning traffic is allowed because of stateful packet inspection.
- Inbound traffic is denied by default.
- When you access the Cisco ASA (for management), you have to access the nearest interface.

Remembering the Default Security Policy

Remember this rhyme: “High to low, good to go. Low to high, must die.” In other words, traffic entering by an interface with security level 100 and leaving the interface with security level 0 is allowed and the returning traffic can come in. However, traffic originating from an interface with security level 50 destined for an interface with security level 100 would by default be denied.

The security level controls the following behavior:

- **Network access:** By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an ACL to the interface. If you enable communication for same security level interfaces, there is an implicit permit for interfaces to access other interfaces on the same security level or lower.
- **Inspection engines:** Some application inspection engines are dependent on the security level. For same security level interfaces, inspection engines apply to traffic in either direction.
 - **NetBIOS inspection engine:** Applied only for outbound connections.
 - **SQL*Net inspection engine:** If a control connection for the SQL*Net (formerly OraServ) port exists between a pair of hosts, then only an inbound data connection is permitted through the Cisco ASA.
- **Filtering:** As an example, HTTP(S) and FTP filtering applies to outbound connections (from a higher level to a lower level). If you enable communication for same security level interfaces, you can filter traffic in either direction.

[Figure 10-25](#) illustrates an example of security levels in action.

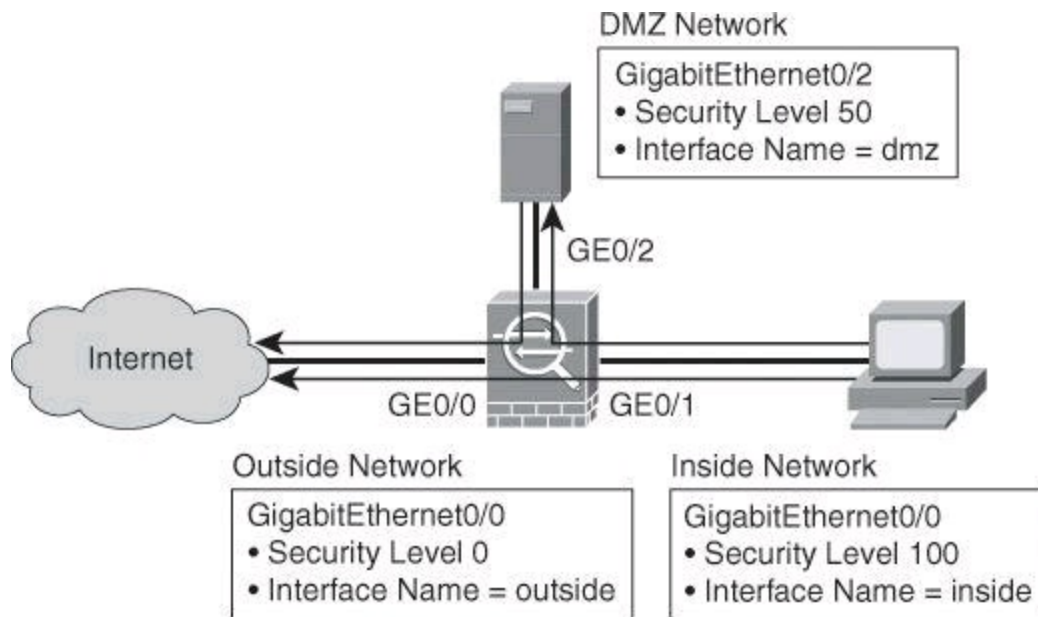


Figure 10-25. Cisco ASA Security Levels and Interface Names

The traffic flows from the inside network to the outside network are permitted by default, because they go from more trusted to less trusted interfaces. In a real network, this means that internal users on the inside interface can access resources on the DMZ freely. They can also initiate connections to the Internet with no restrictions and without the need for an additional policy or additional commands.

By the same token, traffic that is sourced on the outside network, originating from the Internet and going into either the DMZ or the inside network, is denied by default. Return traffic, however, originating on the inside network and returning via the outside interface, would be allowed.

These default traffic flows match the requirements of real-life networks. It is worth noting that firewalls should not be the only security countermeasure in your security architecture. A sound security policy, along with effective countermeasures, should be put in place. For instance, a restrictive policy for outbound traffic should be implemented to overwrite the default permissive behavior based on security levels, to prevent malicious attackers from initiating attacks from the inside network and abusing the default policies based on those security levels.

Managing the Cisco ASA Using the CLI

Cisco ASAs contain a command set that is based on Cisco IOS Software. The appliance provides five configuration modes, similar to Cisco IOS devices:

- ROM monitor mode
- User EXEC mode
- Privileged EXEC mode
- Global configuration mode
- Specific configuration modes

Upon first accessing a Cisco ASA, you are presented with a standard prompt, as shown in [Figure 10-26](#). Multiple configuration and monitoring modes are available at this point:

- **EXEC mode:** Allows view-only capabilities on a restricted group of settings.
- **Privileged EXEC mode:** Expands the set of available commands and settings in EXEC

mode to a set with full administrative privileges. It also allows for maintenance actions in addition to the read-only and monitoring capabilities.

- **Global configuration mode:** Allows changes to the system configuration.
- **Component-specific configuration modes:** Allow changes to specific components, such as interfaces.

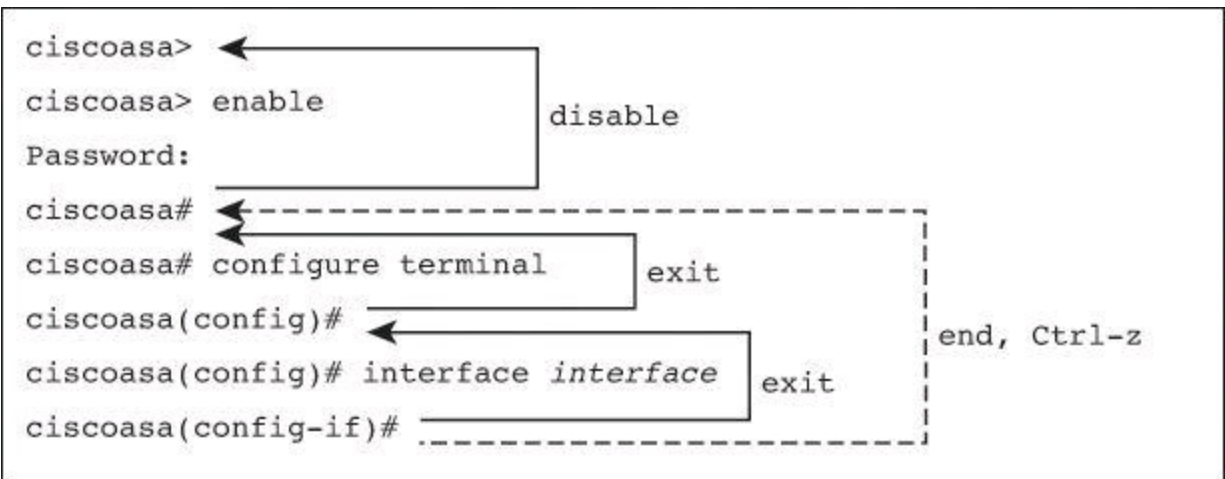


Figure 10-26. Cisco ASA Prompts

You can transition between modes using these commands and keystrokes:

- Use the **enable** command to access the privileged EXEC mode.
- Use **exit** or **disable** to return to the user EXEC mode.
- Use **configure terminal** to access the global configuration mode.
- Use **exit** to exit from global configuration mode, or to exit from a component-specific configuration mode into the previous level.
- Use Ctrl-Z or **end** to exit any configuration mode, any number of levels deep, and return to privileged EXEC mode.

Cisco ASA 5505

The Cisco ASA 5505 is designed for small offices, home offices, and enterprise teleworker environments. As shown in [Figure 10-27](#), the appliance showcases an eight-port built-in Layer 2 switch that provides flexible VLAN configuration. Two of the ports are Power over Ethernet capable, and all of the ports can be assigned to specific VLANs to implement outside, inside, and DMZ subnets. Layer 2 ports are assigned to VLANs, and VLAN interfaces are created for Layer 3 connectivity. VLAN interfaces are used to assign IP addresses, security levels, and other Layer 3 components.




Figure 10-27. Cisco ASA 5505

There are two types of ports and interfaces to configure:

- **Physical switch ports:** The eight Fast Ethernet switch ports of the Cisco ASA forward traffic at Layer 2, using the switching function in hardware. You can configure these interfaces either in access mode or as trunk ports.
- **Logical VLAN interfaces:** Logical VLAN interfaces forward traffic between VLANs at Layer 3, using the configured security policy to apply firewall and VPN services. VLAN interfaces let you divide your Cisco ASA 5505 into separate VLANs, such as home, business, and Internet VLANs.

To segregate the switch ports into separate VLANs, you assign each switch port to a VLAN interface. Switch ports on the same VLAN can communicate with one another using hardware switching. However, when a switch port on VLAN 1 wants to communicate with a switch port on VLAN 2, for instance, the Cisco ASA applies the security policy to the traffic and routes or bridges between the two VLANs.

Other components shown [Figure 10-27](#) include the following:

- **Power connector:** Used for attaching the power cord
- **SSC slot:** Used to install the optional Cisco AIP SSC-5 network IPS module
- **Console port:** Used to connect a computer to the Cisco ASA 5505 using an RJ-45 cable for console operations
- **Reset button:** Reserved for future use
- **USB ports:** Reserved for future use

Cisco ASDM

Cisco ASDM is to the Cisco ASA what Switch Database Management (SDM) or CCP is to Cisco IOS routers, and more. It is a configuration tool that is designed to facilitate the setup, configuration, monitoring, and troubleshooting of Cisco ASAs. The application hides the complexity of commands from administrators, and allows streamlined configurations without requiring extensive knowledge of the security appliance CLI. Cisco ASDM can be used to monitor and configure multiple security appliances that run the same ASDM version.

Cisco ASDM features include the following:

- Runs on a variety of platforms

- Implemented in Java to provide robust, real-time monitoring
- Works with SSL, as shown in [Figure 10-28](#), to ensure secure communication with the security appliance




Figure 10-28. Cisco ASDM

- Comes preloaded in flash memory on new Cisco ASA security appliances running version 7.0 or later
- Can manage multiple security appliances from a single ASDM interface, one at a time
- Cisco ASDM sessions:
 - Five Cisco ASDM sessions per unit (single mode) or context (multimode)
 - 32 sessions per unit in multimode
- Supported on all Cisco ASAs

Cisco ASDM Demo Mode

Similar to the CCP Demo Mode version mentioned earlier in the chapter, Cisco ASDM also has a Demo Mode version. For a tutorial on how to install and use the ASDM Demo Mode, check out “ASDM Demo Mode Tour,” also by Doug McKillip, the URL for which is provided in the “References” section at the end of this chapter.

Preparing the Cisco ASA 5505 for ASDM

Cisco ASDM access requires some minimal configuration so that you can communicate over the network with a management interface.

With a factory default configuration, you can connect to Cisco ASDM using the following interface and network settings:

- The management interface depends on your model:
 - **Cisco ASA 5505:** The switch port to which you connect to Cisco ASDM can be any port, except for Ethernet 0/0.
 - **Cisco ASA 5510 and later:** The interface to which you connect to Cisco ASDM is Management 0/0.
- The default management address is 192.168.1.1.
- The clients that are allowed to access Cisco ASDM must be on the 192.168.1.0/24 network.

Cisco ASA 5510 and Later and Cisco ASDM

With Cisco ASA 5510 and later, the firewall looks for an interface with the name

Management for the setup to be allowed to run. The **nameif management** command doesn't need to be applied to the actual Management 0/0 interface for setup to run. It's not a significant issue, but interesting to know.

If you do not have a factory default configuration, or want to change the firewall or context mode, the following steps are required.

To run Cisco ASDM on a Cisco ASA 5505, perform the following steps:

Step 1. Assign the name “inside” to one interface.

Step 2. Run the interactive setup dialog using the **setup** command.

Step 3. (Optional) Specify which Cisco ASDM image file to use.

Cisco ASDM Features and Menus

Once initialized and connected, the Cisco ASDM Home page, shown in [Figure 10-29](#), lets you view important information about your Cisco ASA. Status information in the Home page is updated every 10 seconds. This page usually has two tabs, Device Dashboard and Firewall Dashboard. If you have a hardware module such as the CSC-SSM (Content Security and Control Security Service Module) installed in your Cisco ASA, the module-specific tab also appears in the Home page. The additional tab displays status information about the software on the specific module. If you have IPS software installed in your Cisco ASA, the Intrusion Prevention tab also appears in the Home page. The additional tab displays status information about the IPS software.

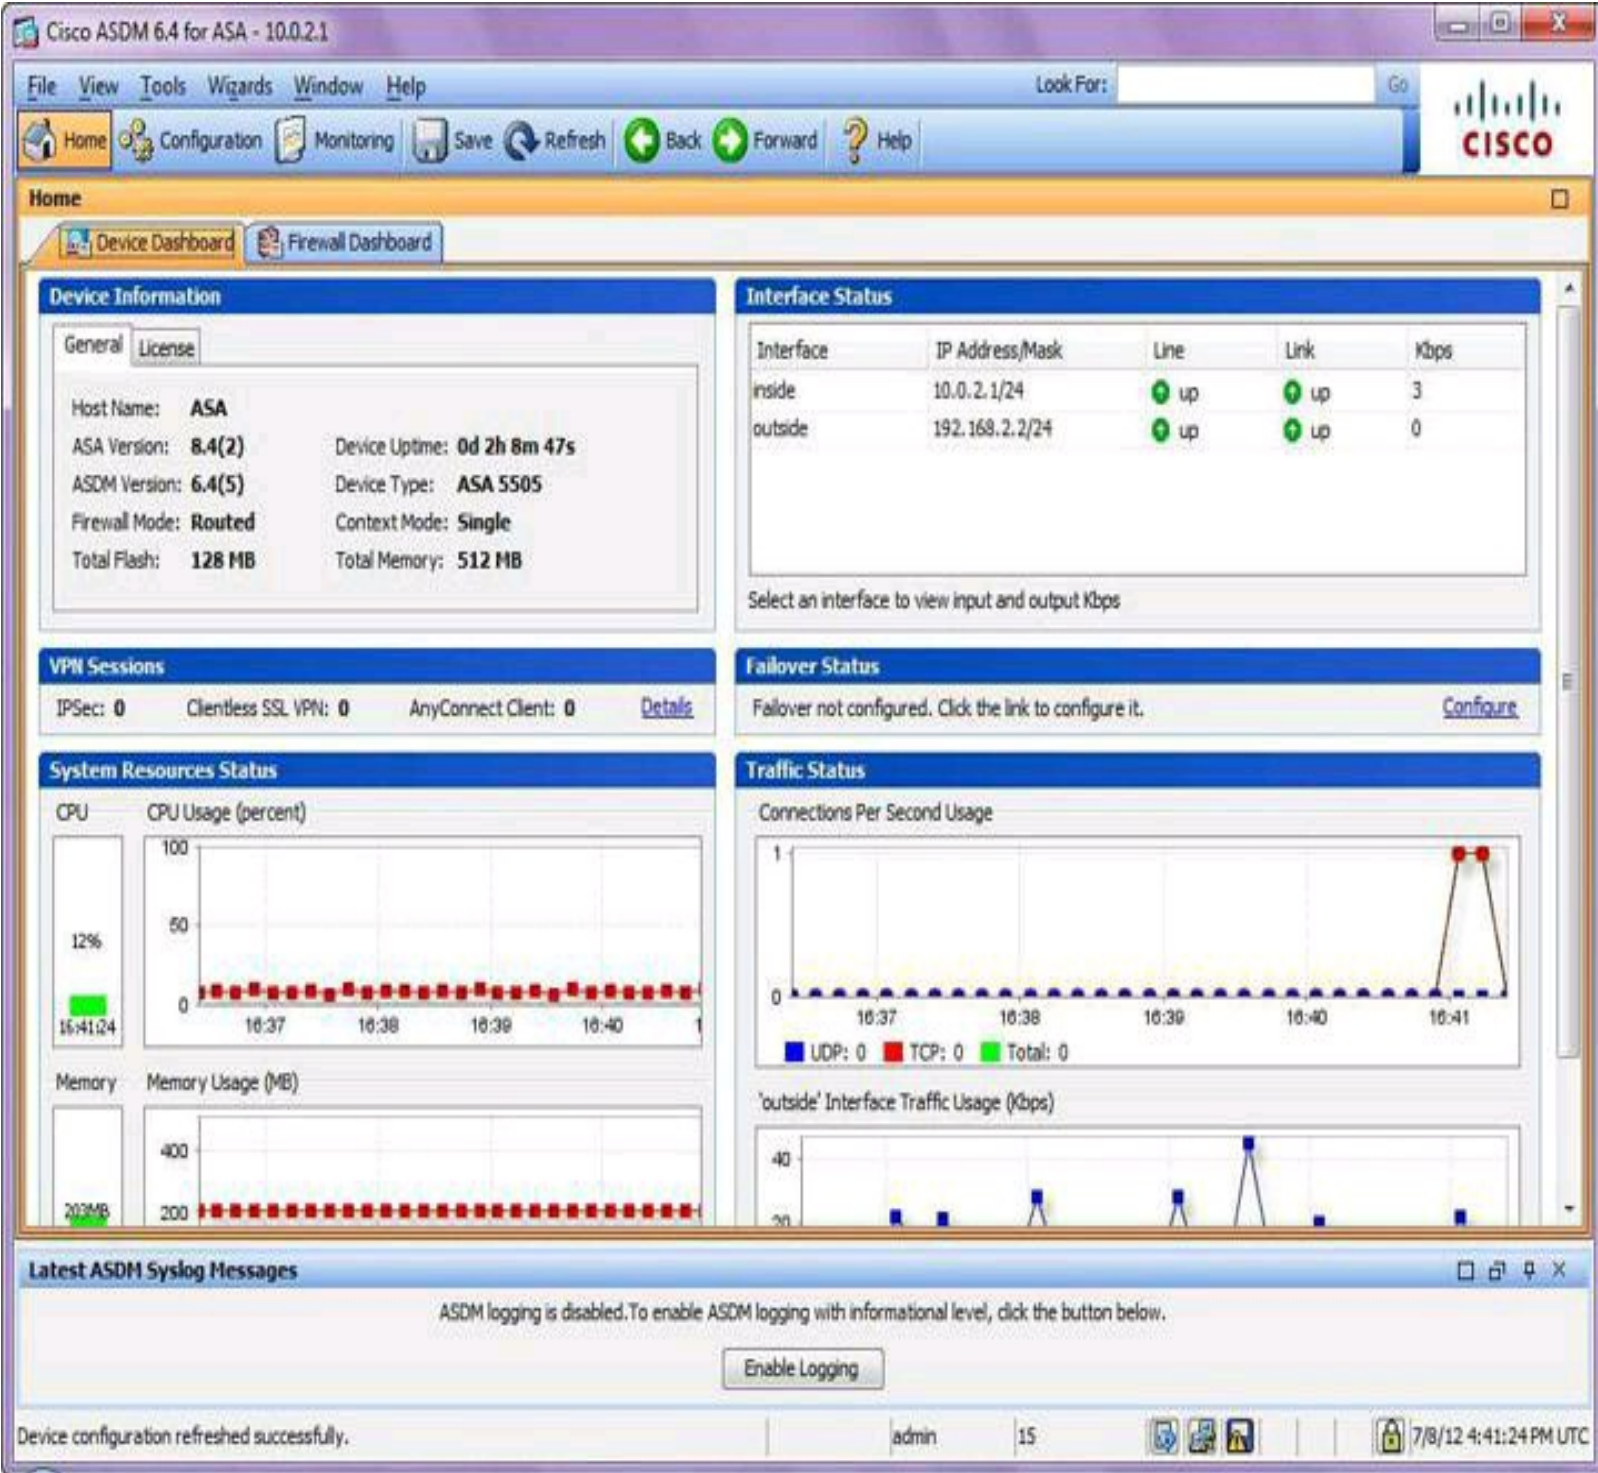


Figure 10-29. Cisco ASDM Home Pane

The Device Dashboard tab lets you view at a glance important information about your Cisco ASA, such as the status of your interfaces, the version you are running, licensing information, and performance.

The Cisco ASDM user interface is designed to provide easy access to the many features that the Cisco ASA supports. It includes the following elements, shown in [Figure 10-30](#):

- A menu bar that provides quick access to files, tools, wizards, and help. Many menu items also have keyboard shortcuts.
- A toolbar that enables you to navigate Cisco ASDM. From the toolbar, you can access the Home, Configuration, and Monitoring pages. You can also get help and navigate between

pages.

- A dockable left navigation pane to move through the Configuration and Monitoring pages. You can click one of the buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it. To access the Configuration and Monitoring pages, you can do one of the following:
 - Click links on the left side of the application window in the left navigation pane. The content pane then displays the path in the title bar of the selected pane.
 - If you know the exact path, you can type it directly into the title bar of the content pane on the right side of the application window, without clicking any links in the left navigation pane.
- A maximize and restore button in the right corner of the content pane that lets you hide and show the left navigation pane.
- A dockable Device List pane with a list of devices that you can access through Cisco ASDM. You can click one of the three buttons in the header to maximize or restore this pane, make it a floating pane that you can move, hide it, or close it.
- A status bar at the bottom of the application window that shows the time, connection status, user, memory status, running configuration status, privilege level, and SSL status.
- The left navigation pane, as mentioned earlier, that shows various objects that you can use in the rules tables when you create access rules, NAT rules, authentication, authorization, and accounting (AAA) rules, filter rules, and service rules. The tab titles within the pane change according to the feature that you are viewing. In addition, the Cisco ASDM Assistant appears in this pane.
- Above the left navigation pane, but not shown in [Figure 10-30](#), is an optional Device List useful when toggling between two or more firewall configurations.

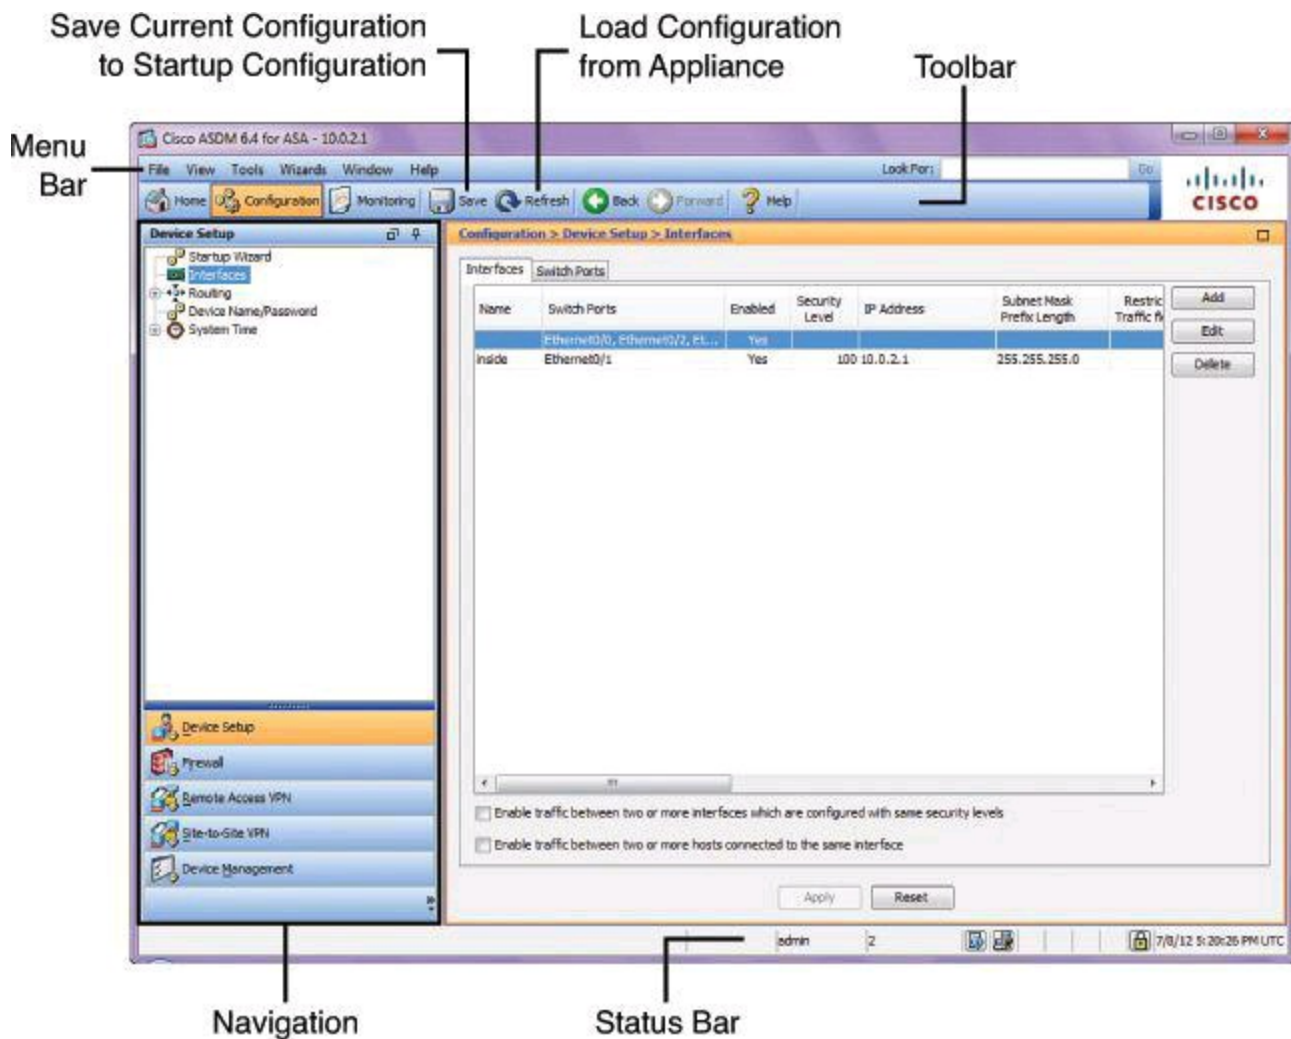


Figure 10-30. Cisco ASDM Interface Elements

When deploying Cisco ASAs, your access policy is made up of one or more access rules per interface or globally for all interfaces. An access rule permits or denies traffic based on the protocol, a source and destination IP address or network, and, optionally, the source and destination ports. Note that global ACLs are only available starting with Cisco ASA version 8.3.

The order of rules is important. When the Cisco ASA decides whether to forward or drop a packet, the Cisco ASA tests the packet against each rule in the order in which the rules are listed, as shown in [Figure 10-31](#). After a match is found, no more rules are checked. For example, if you create an access rule at the beginning that explicitly permits all traffic for an interface, no further rules are ever checked.

Action	Source Address	Destination Address	Service
DENY	10.1.1.0 255.255.255.0	192.168.1.0 255.255.255.0	Any
PERMIT	10.1.1.0 255.255.255.0	172.16.0.0 255.255.255.0.0	SSH
PERMIT	10.1.1.0 255.255.255.0	ANY	HTTP
DENY	ANY	ANY	ANY

Figure 10-31. Example of Stateful Packet Filtering on Cisco ASA

There is an implicit deny-all rule at the end of each interface access rule list. Therefore, if a connection does not match any of the rules, by default, it will be denied.

Note

Interface access rules apply to through traffic—that is, traffic entering the appliance on one interface and then leaving the appliance on the same interface or another interface. Traffic that is directed to the IP addresses or interfaces of the appliance is not subject to these rules, and requires an additional set of rules to implement access control.

Access rules apply based on the security levels described earlier in this chapter. In that sense, these defaults apply if you do not apply an interface access ruleset to an interface:

- All outbound (that is, from higher to lower security level interfaces) connections for hosts on that interface are permitted.
- All inbound (that is, from lower to higher security level interfaces) connections for hosts on that interface are denied.

You need to specifically allow connectivity between interfaces with the same security levels, as well as traffic in and out of the same interface. All such connectivity is then still controlled by the access rules of the relevant interface.

The Cisco ASA supports two types of ACLs:

- **Ingress:** Ingress ACLs apply to traffic as it enters an interface. An ingress ACL can bind an ACL to a specific interface or apply a global rule on all interfaces.
 - **Egress:** Egress ACLs apply to traffic as it exits an interface. An egress ACL is useful, for example, if you want to allow only certain hosts on the inside networks to access a web server on the outside network. Rather than creating multiple ingress ACLs to restrict access, you can create a single egress ACL that allows only the specified hosts. The egress ACL prevents any other hosts from reaching the outside network.
-

Note

The terms *ingress* and *egress* refer to the application of an ACL on an interface, either to traffic entering the Cisco ASA on an interface or traffic exiting the Cisco ASA on an interface. These terms do not refer to the movement of traffic from a lower security interface to a higher security interface, commonly known as *inbound*, or from a higher to lower interface, commonly known as *outbound*.

Access rules can be created and maintained using the Access Rules pane in Cisco ASDM, which you can reach by navigating to **Configuration > Firewall > Access Rules**, as shown in [Figure 10-32](#). Note that the screenshot shown in [Figure 10-32](#) was captured on a Cisco ASA 5505 running code 8.4(2), which explains the presence of the extra column titled User.

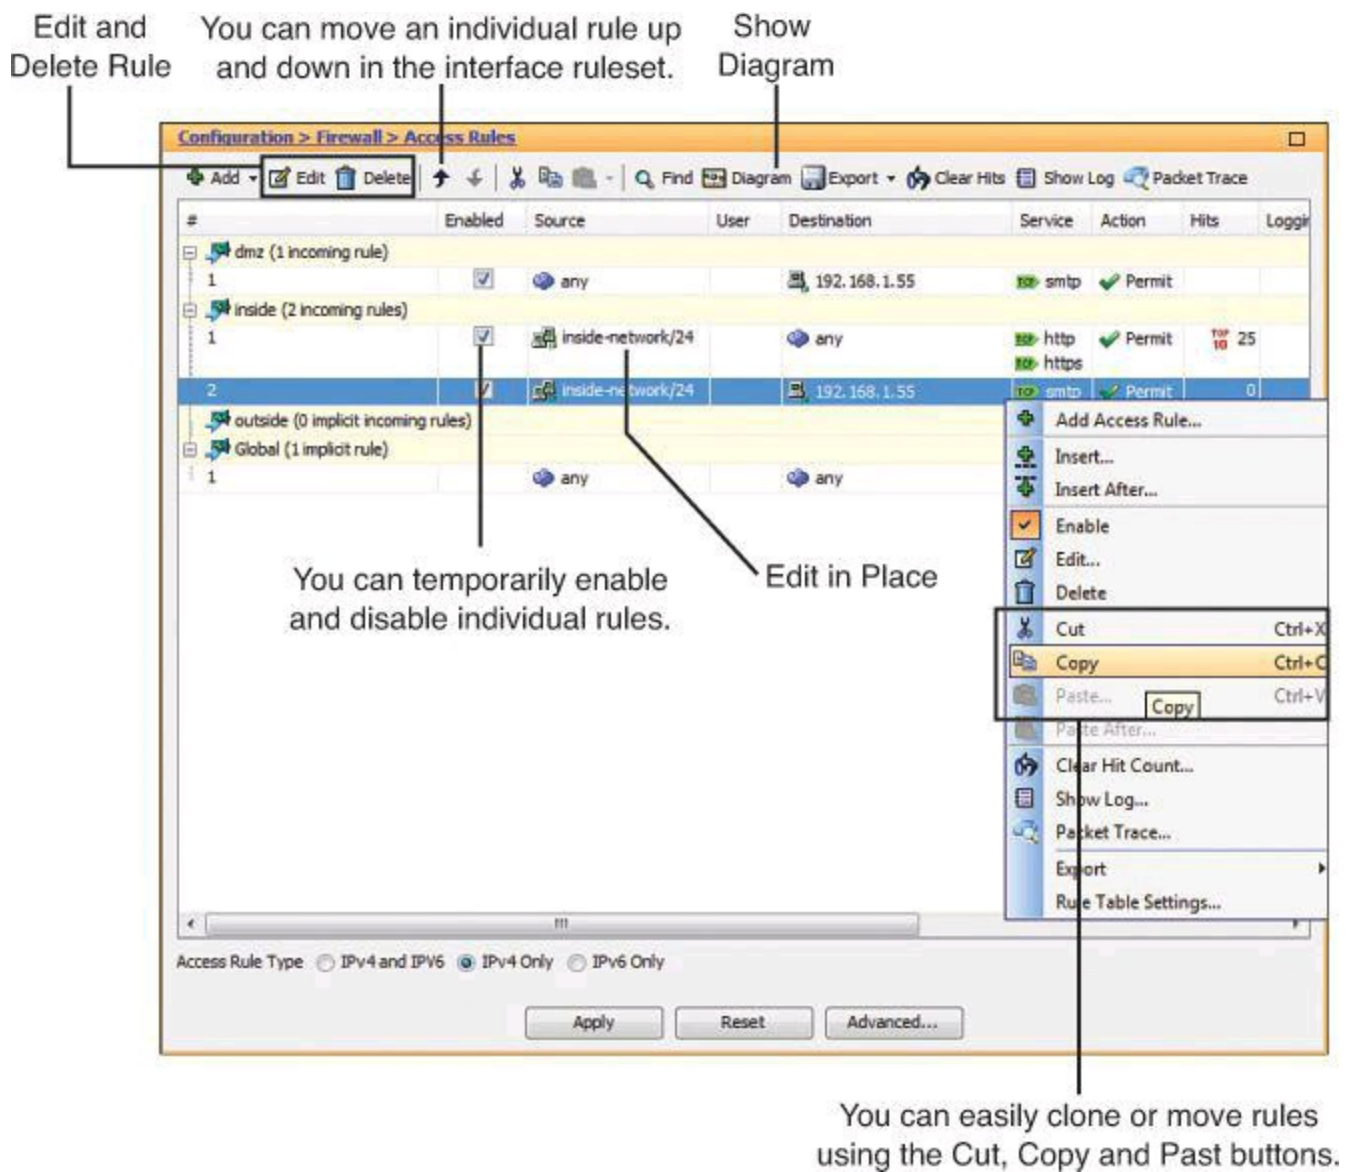


Figure 10-32. Cisco ASA Interface Access Rules Configured with Cisco ASDM

Using the options that are shown [Figure 10-32](#), you can add, delete, or edit rules either by using the options in the top menu or by right-clicking a particular rule and choosing from the context menu. You can also temporarily disable and enable rules.

Cisco Modular Policy Framework

For application layer inspection, as well as other advanced options, the Cisco Modular Policy Framework (MPF) is also available on Cisco ASAs. Cisco MPF on ASAs is similar to C3PL on Cisco IOS routers. Cisco MPF uses these three configuration objects to define modular, object-oriented, hierarchical policies:

- **Class maps:** Define a match criterion to identify qualifying traffic
- **Policy maps:** Associate actions to the match criteria
- **Service policies:** Attach the policy map to an interface, or globally to all interfaces of the appliance

Cisco MPF, illustrated in [Figure 10-33](#), allows granular classification of traffic flows, in order to apply different advanced policies to different flows. These features, among others, use Cisco MPF:

- **Hardware modules:** Traffic is redirected granularly from the Cisco ASA to the modules using Cisco MPF.
- **Advanced Inspection and Control:** Traffic can be classified at Layers 5 through 7 for application layer inspection.
- Rate limiting and quality of service (QoS) features.

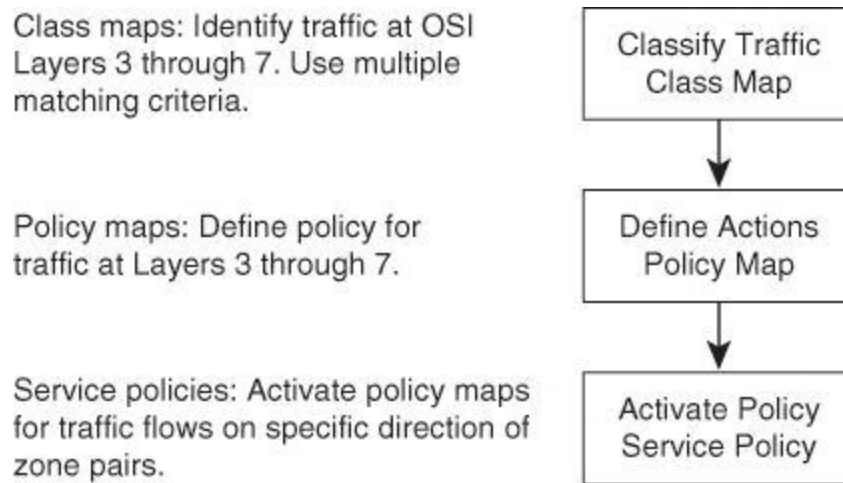


Figure 10-33. Cisco ASA and Modular Policy Framework

Class Map: Identifying Traffic on Which a Policy Will Be Enforced

Cisco ASA class maps allow a richer set of match criteria, as compared to their counterparts in Cisco IOS Zone-Based Policy Firewalls. The list includes these criteria for matching information at Layers 3 and 4:

- Access list
- Any packet
- Default inspection traffic
- IP differentiated services code point (DSCP)
- IP flow
- TCP and UDP ports
- IP precedence
- Real-Time Transport Protocol (RTP) port numbers
- VPN tunnel group

Layers 5 through 7 class maps are also available, and provide an even richer set of criteria for application-specific parameters. For instance, you can match HTTP URLs and request methods.

Policy Map: Configuring the Action That Will Be Applied to the Traffic

Cisco ASA policy maps also allow a more rich set of actions, as compared to their router counterparts for Cisco IOS Zone-Based Policy Firewalls. The actions include not only firewall-related functions, such as application inspection and advanced connection settings, but also QoS actions, such as traffic prioritization, and traffic rerouting to hardware modules integrated into the security appliance. The list of possible actions includes

- Sending traffic to the Advanced Inspection and Prevention Security Services Module

(AIP-SSM) or the Trend Micro InterScan for Cisco CSC-SSM

- Sending NetFlow information
- Prioritizing, policing, or shaping traffic (QoS)
- Configuring advanced connection settings (such as maximum number of simultaneous embryonic connections per client)
- Application inspection (such as preparing for dynamic port for FTP sessions)

When creating a policy map, you need to provide it with a name. You also need to refer to class maps, where traffic is identified.

Service Policy: Activating the Policy

The service policy, defined in the policy map configuration, will be applied to an interface or globally. You can apply only one service policy per interface. However, a service policy might contain multiple policy maps, and thus different actions to be applied on different flows of traffic.

Note

Interface service policies take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.

Cisco ASA Modular Policy Framework: Simple Example

[Figure 10-34](#) provides an example of MPF at work. Three map classes have been created:

- **Internet:** Traffic that has port 80 listed in its TCP header will be classified as **Internet** traffic.
- **Engineers:** Traffic with IP subnet 10.66.0.0 will be classified as **Engineers** traffic.
- **Voice:** Traffic with Layer 3 DSCP bits set to decimal value 46 will be recognized as VoIP traffic and will be classified as **Voice** traffic.

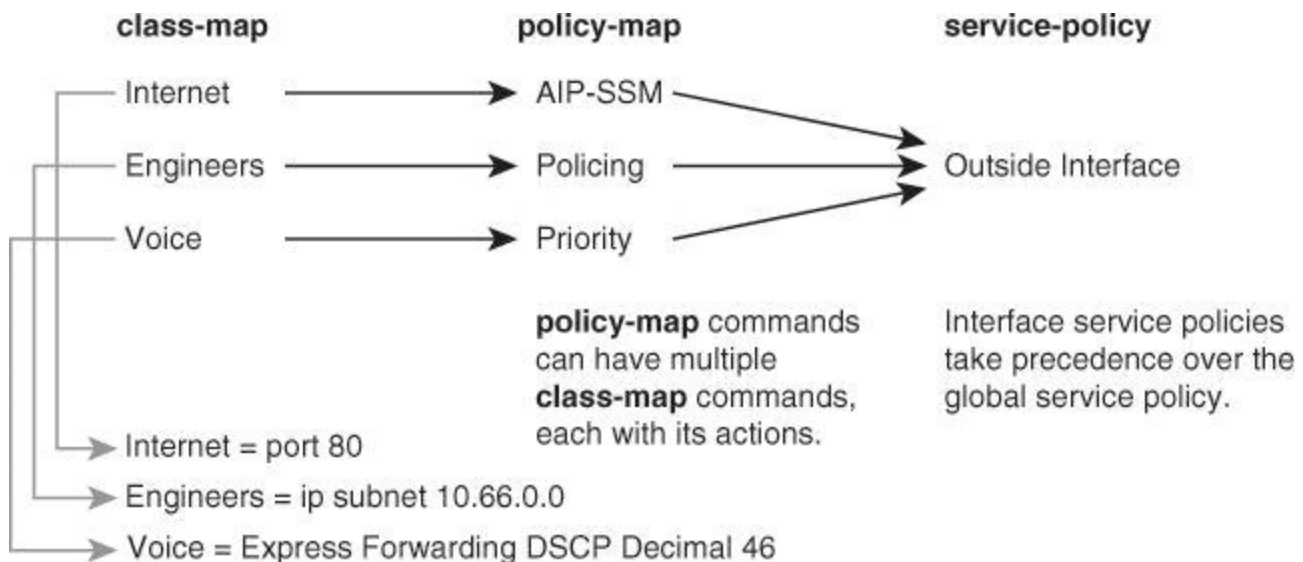


Figure 10-34. Example of How Modular Policy Framework Can Be Used on Cisco ASA

The traffic identified as **Internet** will be passed to the IPS module inserted in the Cisco ASA firewall. This module is the Advanced Inspection and Prevention Security Services Module (AIP-SSM). Traffic identified as **Engineers** will be throttled using rate limitation. Traffic identified as **Voice** will be prioritized on the way out.

With the example in [Figure 10-34](#), the outside interface is made aware with the **service-policy** command that it is responsible for enforcing the three **policy-map** commands, which themselves list the actions to be applied on the different flows of traffic created by the **class-map** commands.

First Packet Dictates the Policy Used

The policy is assigned upon receipt of the first packet. As an example, in a TCP communication, the policy is “chosen” upon the SYN packet transmission. All subsequent packets from the same session will use the action that was stipulated when the policy was selected at the beginning of the connection. Any changes made to that policy during the session lifetime do not affect that specific connection. Only new connections will be affected. Therefore, if you make a change to a policy, it is best to clear all current connections that could be affected by the policy, to ensure that the new policy will be used.

Basic Outbound Access Control on Cisco ASA Using Cisco ASDM

Cisco ASDM can be used to initialize Cisco ASAs for basic operations. The Cisco ASDM Startup Wizard is available to configure interface settings, basic access control, NAT, DHCP, basic management features, and others.

[Figure 10-35](#) describes a configuration scenario for the use of the Cisco ASDM Startup Wizard. This network, using a Cisco ASA 5505 security appliance, requires outbound access for clients on the 10.0.2.0/24 subnet. No public servers are required, and inbound access that is originated at the untrusted network should be blocked. The Cisco ASA will act as DHCP server for the 10.0.2.0/24 subnet, and will perform dynamic PAT (NAT overload) on outbound traffic using the outside interface IP address of the appliance.

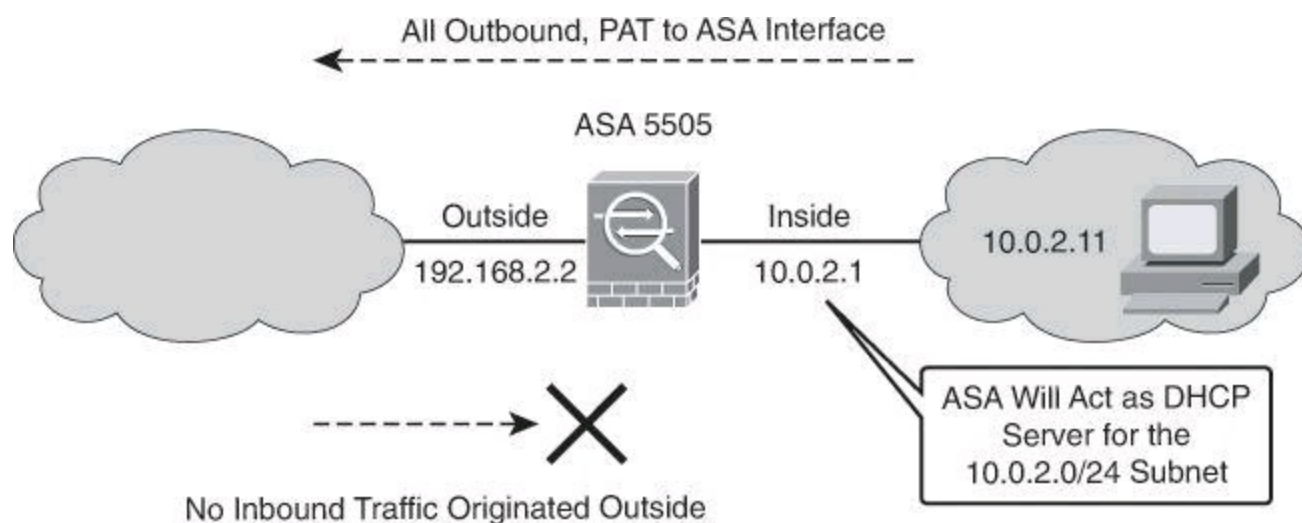


Figure 10-35. Configuration Scenario for Outbound Traffic Control on Cisco ASA

Scenario Configuration Steps Using Cisco ASDM

The configuration steps are as follows:

Step 1. Prepare Cisco ASA for Cisco ASDM access via the CLI.

Step 2. Run the interactive setup dialog in Cisco ASDM.

Step 3. Verify the configuration.

Step 4. Verify firewall activity using the Packet Tracer tool.

The preparation steps are configured using the CLI. Once the appliance is ready, Cisco ASDM is used to invoke the Startup Wizard and continue the basic configuration. Verification is shown for both the CLI and ASDM, and the Packet Tracer tool is used for verification and troubleshooting.

Step 1: Prepare Cisco ASA for Cisco ASDM Access Using the CLI

In preparing a Cisco ASA 5505 security appliance for Cisco ASDM access, the first step is to enable an interface and configure it with a logical name. The **nameif** command, mentioned earlier in a sidebar in this chapter, assigns a name to an interface when in interface configuration mode. Security levels are automatically assigned when you issue the **nameif** command. The inside interface acquires security level 100, while other interfaces acquire level 0 by default. This default security level can be changed with the **security-level** command.

Notice that in the case of the Cisco ASA 5505 model, security levels and IP addresses are assigned to VLAN interfaces, as shown in [Example 10-4](#). Physical interfaces in this model are Layer 2 ports, and VLAN interfaces are needed to provide Layer 3 connectivity and configure the appliance with IP addresses, security levels, ACLs, and other components.

Note

These guidelines and configuration examples apply to the Cisco ASA 5505 model. Other models, such as 5510, 5520, 5540, and 5580, are subject to different rules and requirements. As examples, with the Cisco ASA 5510, an interface would be referred to as Ethernet 0/0; with the Cisco ASA 5520, it would be referred as GigabitEthernet 0/0.

Example 10-4. Configuring an Interface on Cisco ASA 5505

[Click here to view code image](#)

```
asa5505(config)# interface vlan21
asa5505(config-if)# nameif inside
```

The next task is to enter the **setup** command to run the setup initialization dialog, as shown in [Example 10-5](#). The **setup** command configures the following:

- Device settings, such as the firewall mode, clock settings, enable password, and DNS information.
- Interface settings, such as IP addresses and subnet masks.
- Cisco ASDM settings, specifically access restrictions in the form of IP addresses and subnets that are allowed to access the appliance via Cisco ASDM. In [Example 10-5](#), this

address is 10.0.2.11.

Example 10-5. Running the Cisco ASA Setup

[Click here to view code image](#)

```
!  
asa5505(config)# interface vlan21  
asa5505(config-if)# nameif inside  
INFO: Security level for "inside" set to 100 by default.  
ciscoasa(config)# setup  
Pre-configure Firewall now through interactive prompts [yes]? <Enter>  
Firewall Mode  
[Routed]: <Enter>  
Enable password [<use current password>]: cisco123  
Allow password recovery [yes]? <Enter>  
Clock (UTC):  
  Year [2012]: <Enter>  
  Month [Aug]: <Enter>  
  Day [26]: <Enter>  
  Time [10:21:49]: 15:34:00  
Inside IP address [0.0.0.0]: 10.0.2.1  
Inside network mask [255.255.255.255]: 255.255.255.0  
Host name [ciscoasa]: ASA Domain name: cisco.com  
IP address of host running Device Manager: 10.0.2.11  
Use this configuration and write to flash? Y  
!
```

Routed Versus Transparent

In the setup dialog of [Example 10-5](#), you might have noticed **Firewall Mode [Routed]**. The Cisco IOS Firewall, like the Cisco ASA firewall, can operate in routed mode or transparent mode. The default mode is routed mode, where the firewall is a routed hop and acts as a default gateway. In transparent mode, the firewall can perform its duties at Layer 2. Transparent firewalls are beyond the scope of this book.

Optionally, you should define the location of Cisco ASDM files in the flash memory of the appliance, as shown in [Example 10-6](#). This is especially important when using multiple versions of Cisco ASDM.

Example 10-6. Specifying the Cisco ASDM File to Use

[Click here to view code image](#)

```
!  
asa5505(config)# asdm image disk0:/asdm-641.bin  
!
```

Note

Cisco ASA's internal flash is referred to as **disk0**, while its external flash (compact flash) is referred to as **disk1**. We refer to Cisco PIX Firewall flash, even when running code 8.x, as **flash0**.

[Figure 10-36](#) illustrates and explains the CLI commands that result from the Setup Wizard. Notice how the Cisco ASA 5505 security appliance uses a VLAN interfaces to define the inside interface. The VLAN for VLAN interface 21, in this example, is assigned to the physical interface Ethernet 0/1, which is physically connected to the inside network.

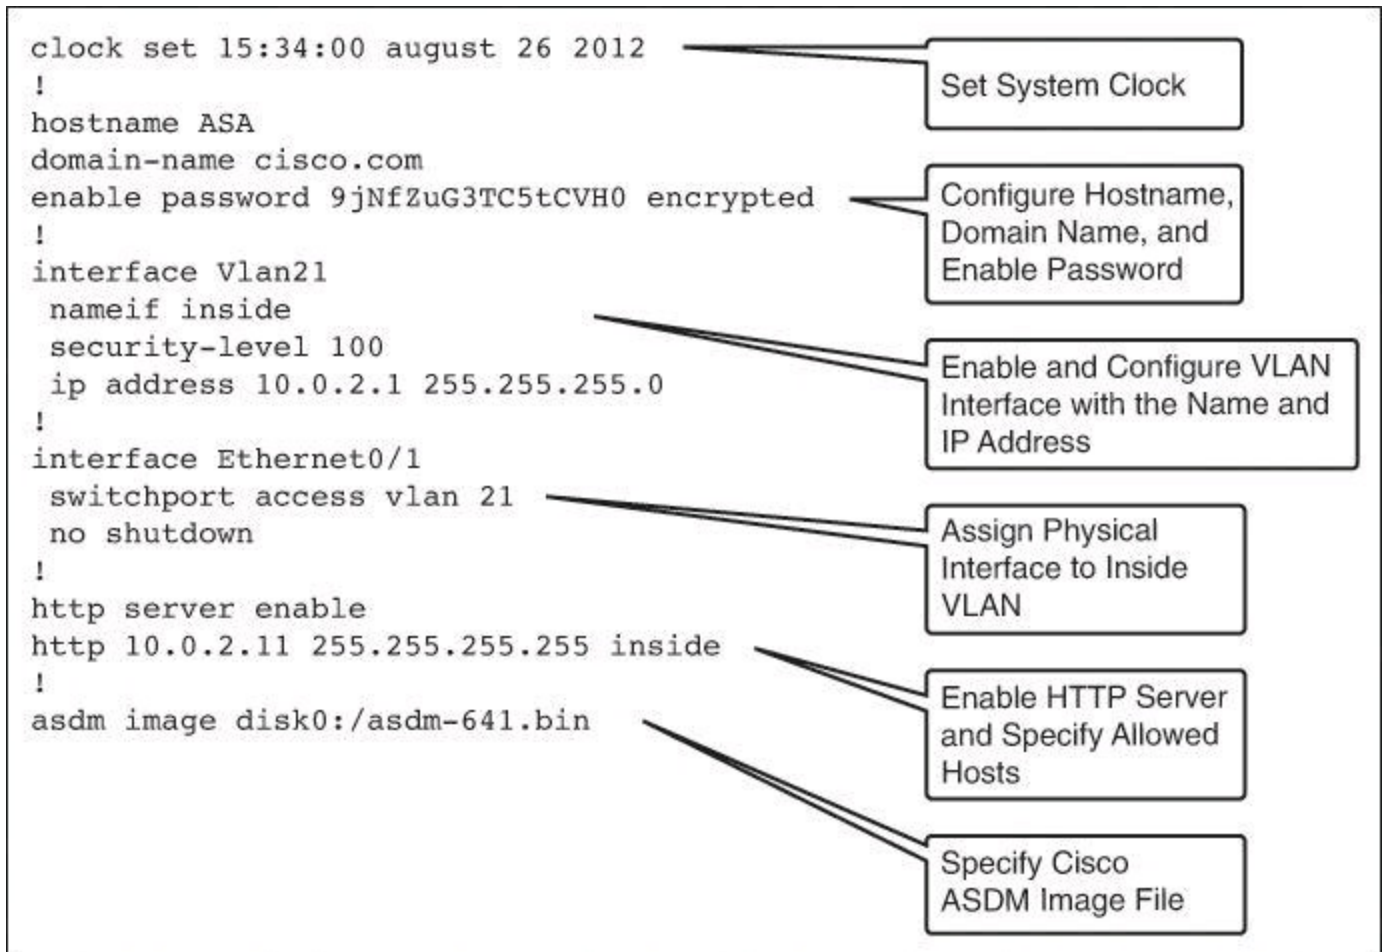


Figure 10-36. Verifying and Understanding the Setup Configuration of a Cisco ASA

Step 2: Run the Startup Wizard from Cisco ASDM

You can start Cisco ASDM using either of two methods, shown in [Figure 10-37](#):

- **ASDM-IDM Launcher (Windows only)**: The Launcher is an application (downloaded from the Cisco ASA using a web browser) that you can use to connect to any Cisco ASA IP address. You do not need to re-download the Launcher if you want to connect to other Cisco ASAs. The Launcher also lets you run a virtual Cisco ASDM in Demo Mode using files that are downloaded locally.
- **Java Web Start**: For each Cisco ASA that you manage, you need to connect with a web browser and then save or launch the Java Web Start application. You can optionally save the application to your PC. However, you need separate applications for each Cisco ASA

IP address.

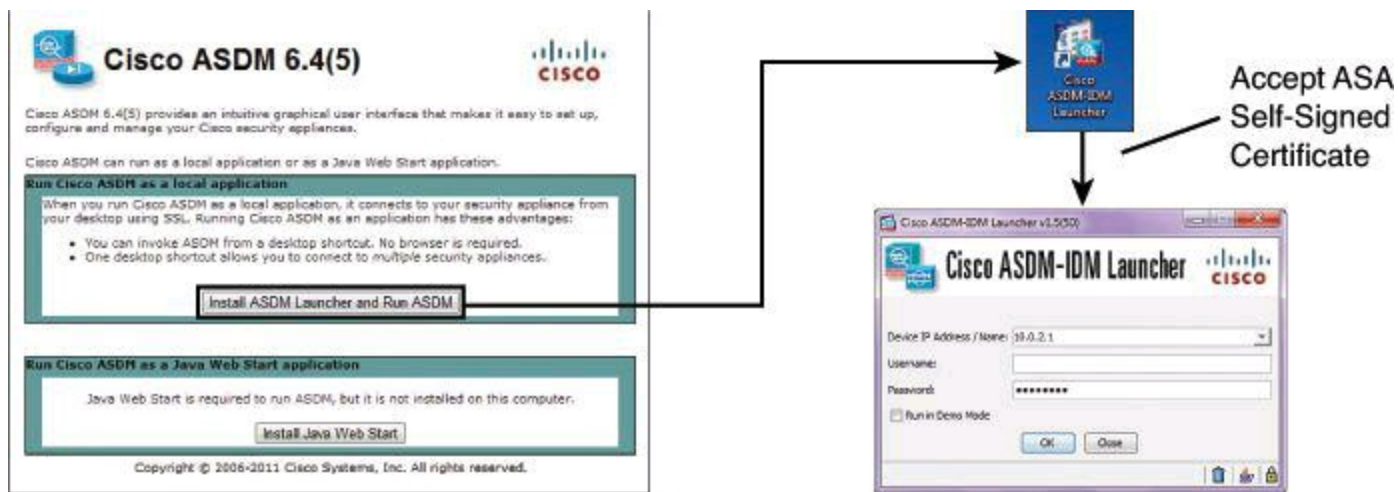


Figure 10-37. Launching Cisco ASDM

Once started, the security appliance will prompt you for credentials for administrative access.

Note

ASDM uses SSL for communications between the management station and the security appliance. By default, a self-signed certificate is generated for use by Cisco ASA. Your browser will prompt you for confirmation that you want to accept and use this digital certificate.

The next task is to run the Startup Wizard. To access this feature in the main Cisco ASDM application window, choose one of the following:

- Navigate to **Wizards > Startup Wizard**.
- Navigate to **Configuration > Device Setup > Startup Wizard**, and then click **Launch Startup Wizard**.

Step 1 of the Startup Wizard is to define the starting point in terms of configuration, as shown in [Figure 10-38](#). The options are as follows:

- To change the existing configuration, click the Modify Existing Configuration radio button.
- To set the configuration to the factory default values, click the Reset Configuration to Factory Defaults radio button.
 - To configure the IP address and subnet mask of the Management 0/0 (Cisco ASA 5510 and later) or VLAN 1 (Cisco ASA 5505) interface to be different from the default value (192.168.1.1), check the **Configure the IP Address of the Management Interface** check box.
 - Choose the first option and click **Next** if, like in our case, you wish to modify the current configuration.



Figure 10-38. Starting the Cisco ASDM Startup Wizard and the First Screen

Note

If you reset the configuration to factory defaults, you cannot undo these changes by clicking Cancel or by closing this screen.

Step 2 of the Startup Wizard defines the basic configuration of the appliance, as shown in [Figure 10-39](#). The following components can be configured:

- **Configure the Device for Teleworker Usage:** Check this check box to specify a group of configuration settings for a remote worker. This is available only for Cisco ASA 5505 security appliances.
- **Cisco ASA Host Name:** The hostname appears in the command-line prompt, and if you establish sessions to multiple devices, the hostname helps you keep track of where you enter commands. The hostname is also used in system messages.
- **Domain Name:** The Cisco ASA appends the domain name as a suffix to unqualified names. For example, if you set the domain name to “example.com,” and specify a syslog server by the unqualified name of “jupiter,” then the security appliance qualifies the name to “jupiter.example.com.”
- **Change Privileged Mode (Enable) Password:** Checking this check box enables you to change the configured enable (privilege 15) password. The enable password lets you access privileged EXEC mode after you log in. Also, this password is used to access Cisco ASDM as the default user, which is blank. The default user shows as “enable_15” in the User Accounts pane. (If you configure user authentication for enable access, then each user has their own password, and this enable password is not used. In addition, you can configure authentication for HTTP/ASDM access.)

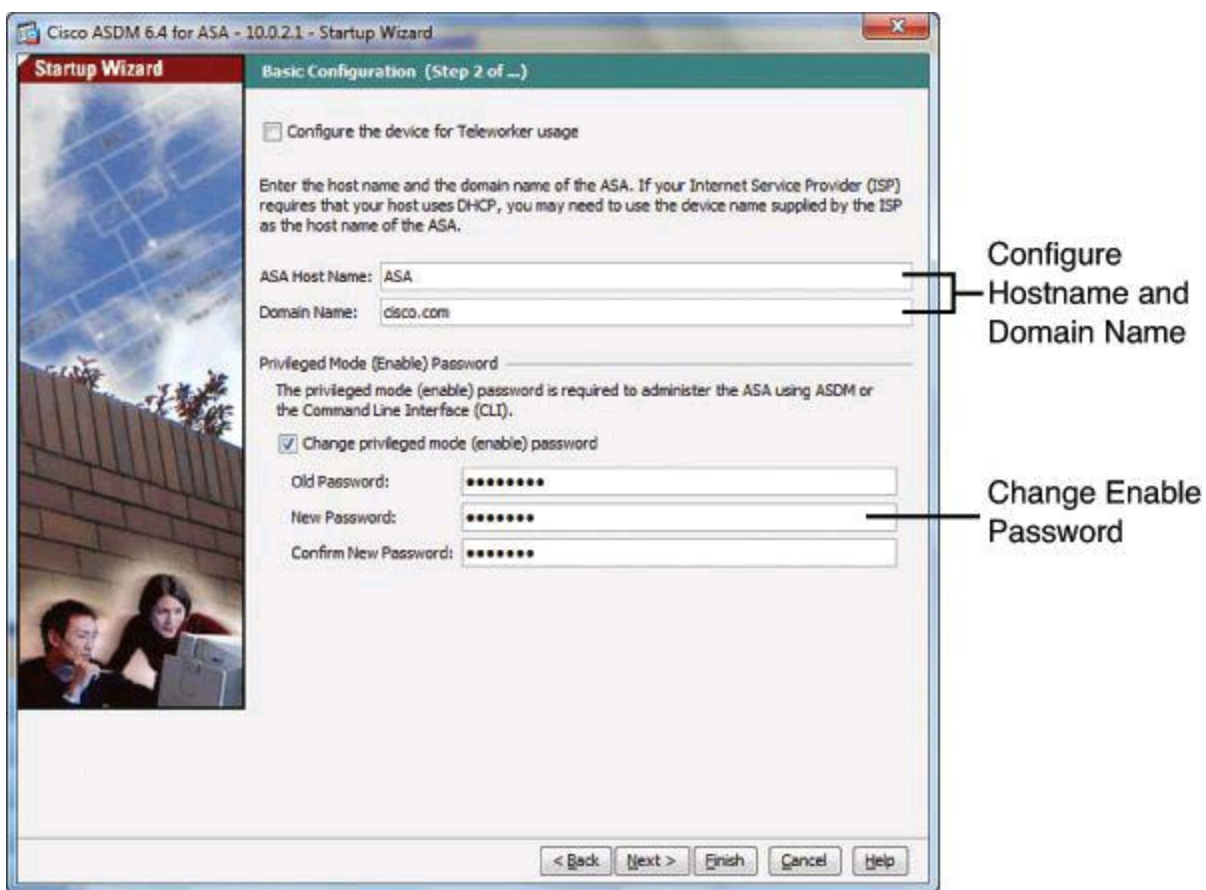


Figure 10-39. Second Screen of the Cisco ASDM Startup Wizard

Multiple screens allow the configuration of the appliance interfaces. In Step 3 of the wizard, (running the Plus code instead of the Base code), you group the eight Fast Ethernet switch ports on the Cisco ASA 5505 into three VLANs. These VLANs function as separate Layer 3 networks. You can then choose or create the VLANs that define your network—one for each interface: Outside, Inside, or DMZ. You have the option to not define a DMZ by not creating a DMZ VLAN.

Security levels are configured automatically when you use the wizard, but you can change them later using Cisco ASDM options. The outside interface acquires a security level of 0, the DMZ interface acquires a security level of 50, and the inside interface acquires a security level of 100.

In [Figure 10-40](#), VLAN 21 is allocated to the inside interface, while VLAN 22 is allocated to the outside interface. No DMZ VLAN is created.

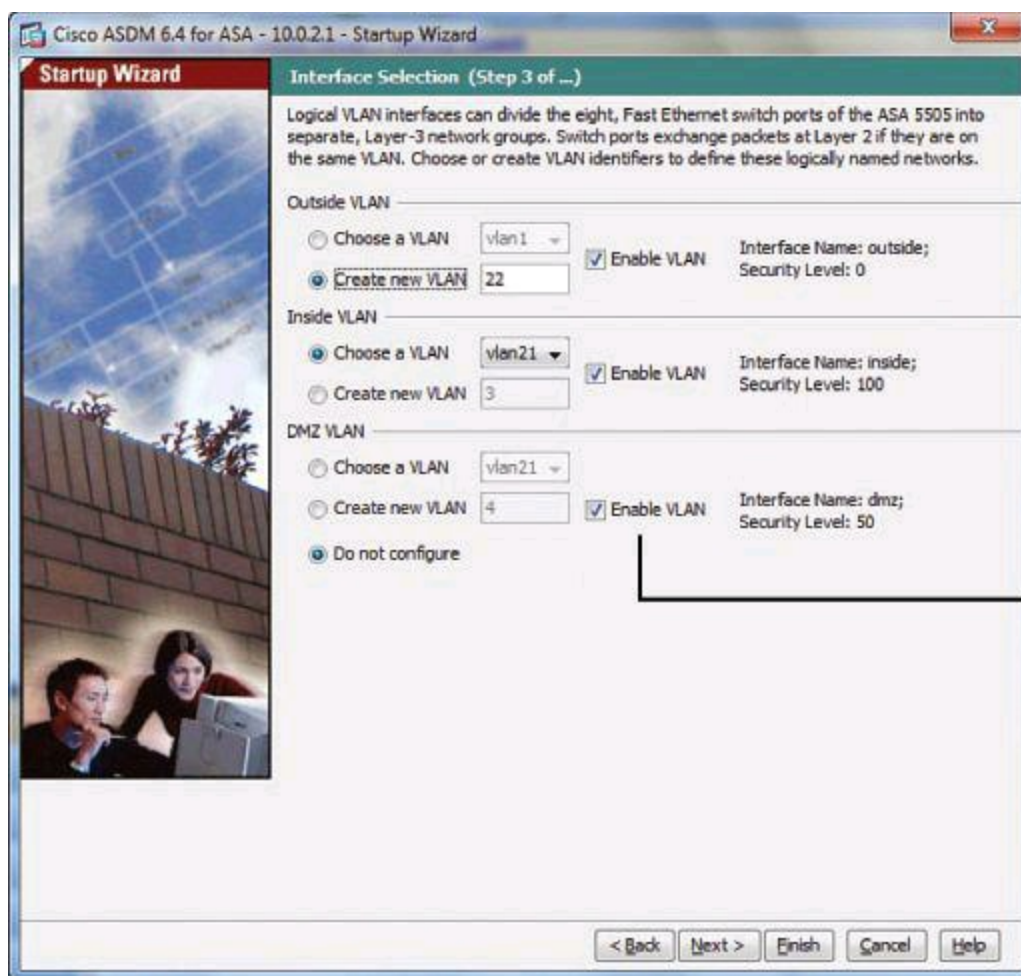
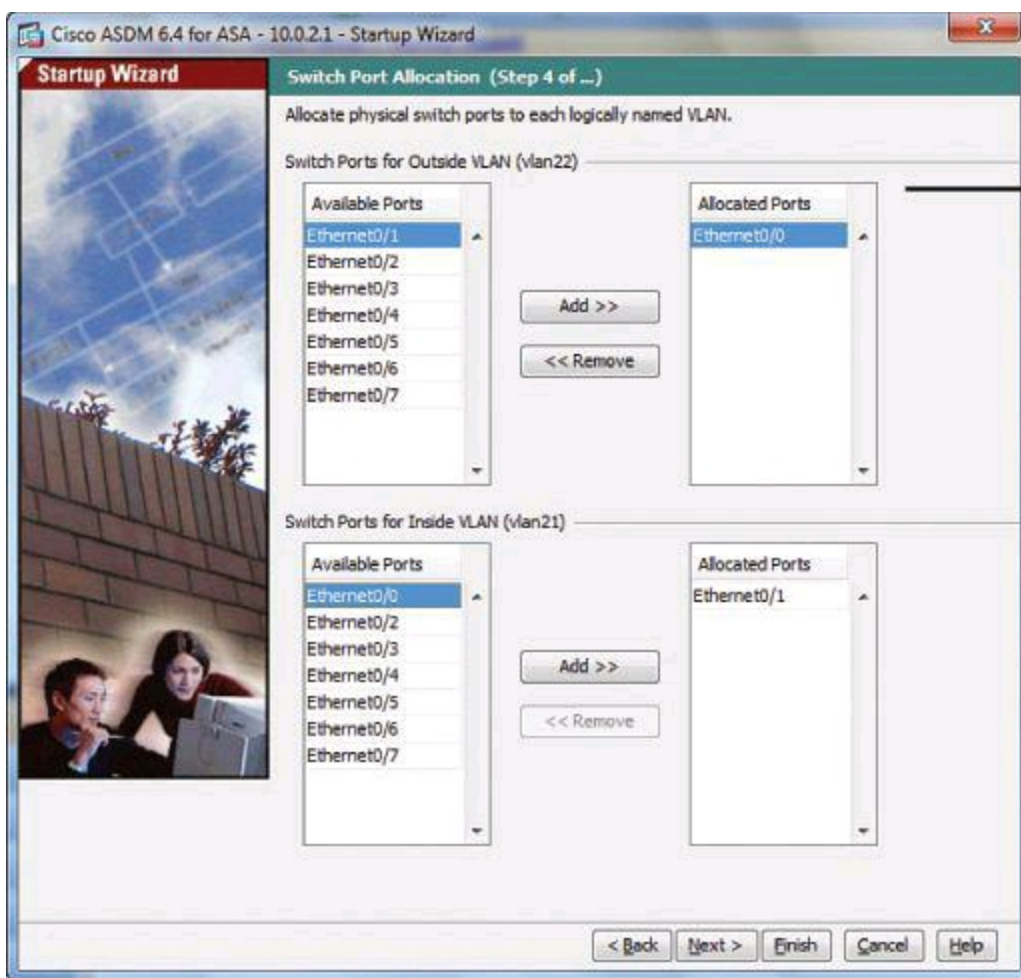


Figure 10-40. Interface Selection from Cisco ASDM Startup Wizard

Step 4 of the wizard is switch port allocation, which lets you allocate switch ports to outside, inside, or DMZ interfaces. By default, all switch ports are assigned to VLAN 1 (inside).

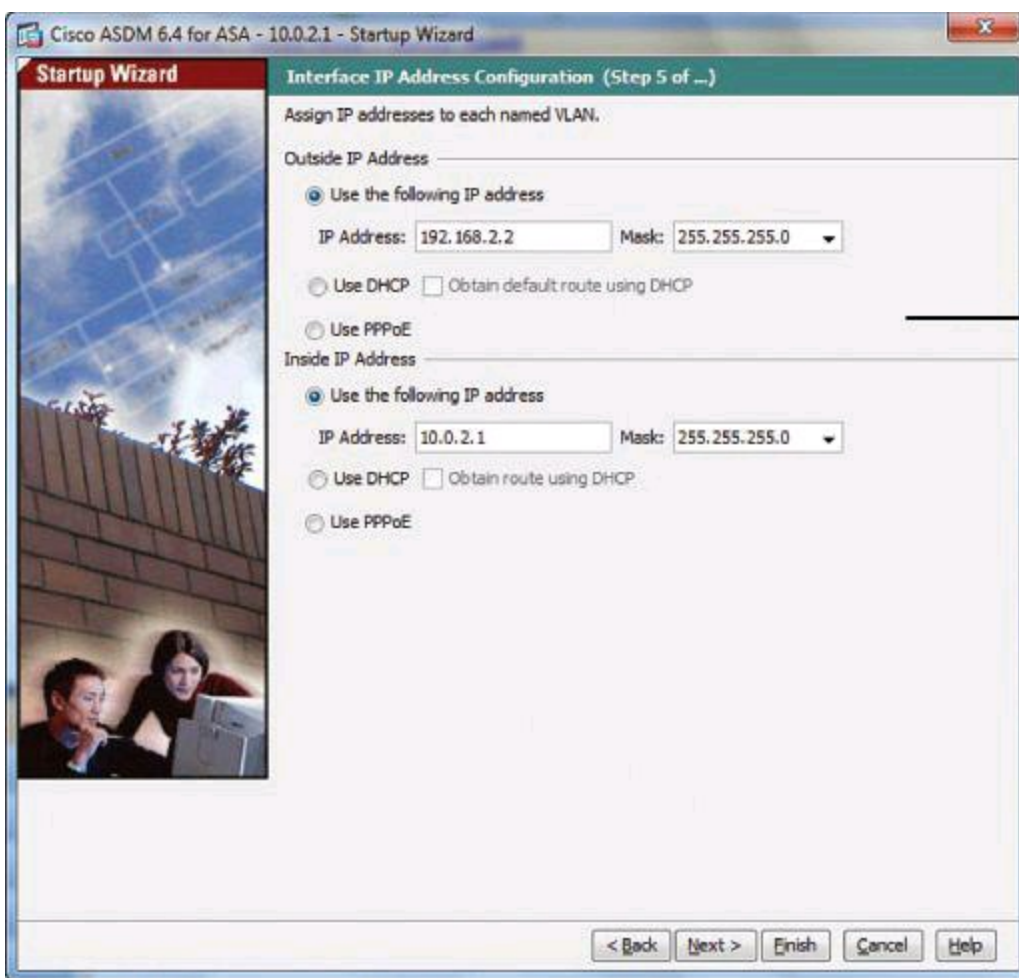
[Figure 10-41](#) shows the example for a Cisco ASA 5505 security appliance, which includes an eight-port switch. In this example, the Ethernet 0/1 interface is allocated to the inside VLAN, while the Ethernet 0/0 interface is allocated to the outside VLAN.



Allocate switch ports to VLANs.

Figure 10-41. Switch Port Selection from Cisco ASDM Startup Wizard (Cisco ASA 5505 Only)

Step 5 of the wizard lets you configure the IP address for each VLAN interface, as shown in [Figure 10-42](#). Notice the options to configure Cisco ASA interfaces to obtain their IP address, as well as a default route, automatically from the network via DHCP.



Assign IP addresses and masks per logical interface.

Figure 10-42. Interface IP Address Configuration from Cisco ASDM Startup Wizard

Step 6 of the wizard enables you to configure network services, either DHCP or PPPoE, depending on what selections you made in Step 5. [Figure 10-43](#) illustrates the options to configure the Cisco ASA as a DHCP server. The options allow the inside interface to deliver DHCP services to DHCP clients reachable through that interface.

Enable DHCP services if required.

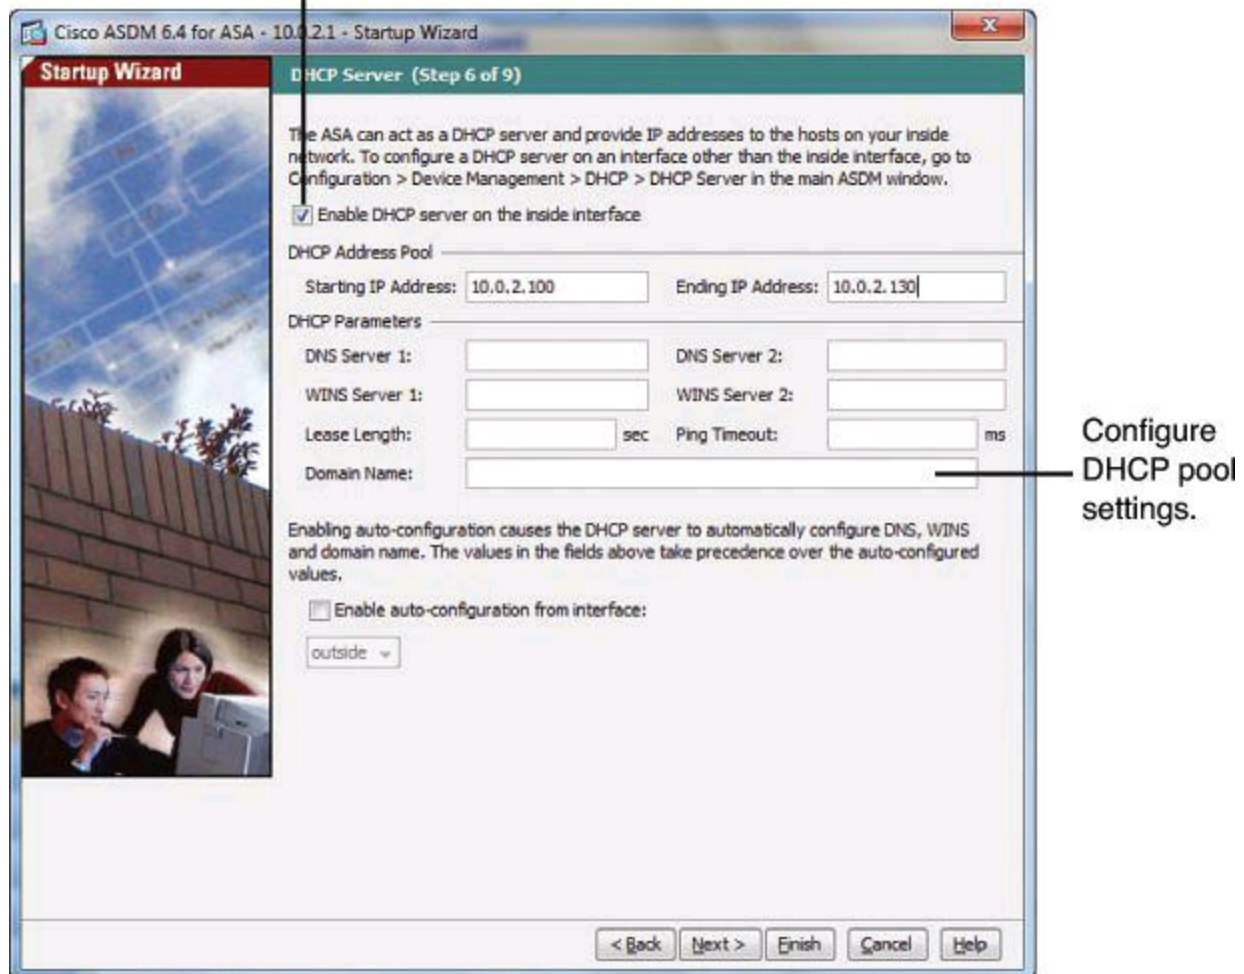


Figure 10-43. DHCP Server Configuration from Cisco ASDM Startup Wizard

In our scenario, the inside interface provides DHCP servers using the range of addresses 10.0.2.100 to 10.0.2.130.

Step 7 of the wizard is to define NAT and PAT settings, if required. You can enable dynamic NAT or dynamic PAT. You have the option to define additional Cisco ASA objects such as IP addresses and IP subnets, in order to define the PAT IP addresses or the ranges of source IP addresses to match. You configure these objects in additional windows that open from the main wizard window when you choose the different options.

As shown in [Figure 10-44](#), you can configure PAT and define the IP address on the outside interface to overload for all traffic heading outbound when originated on the inside interface.

Enable PAT.

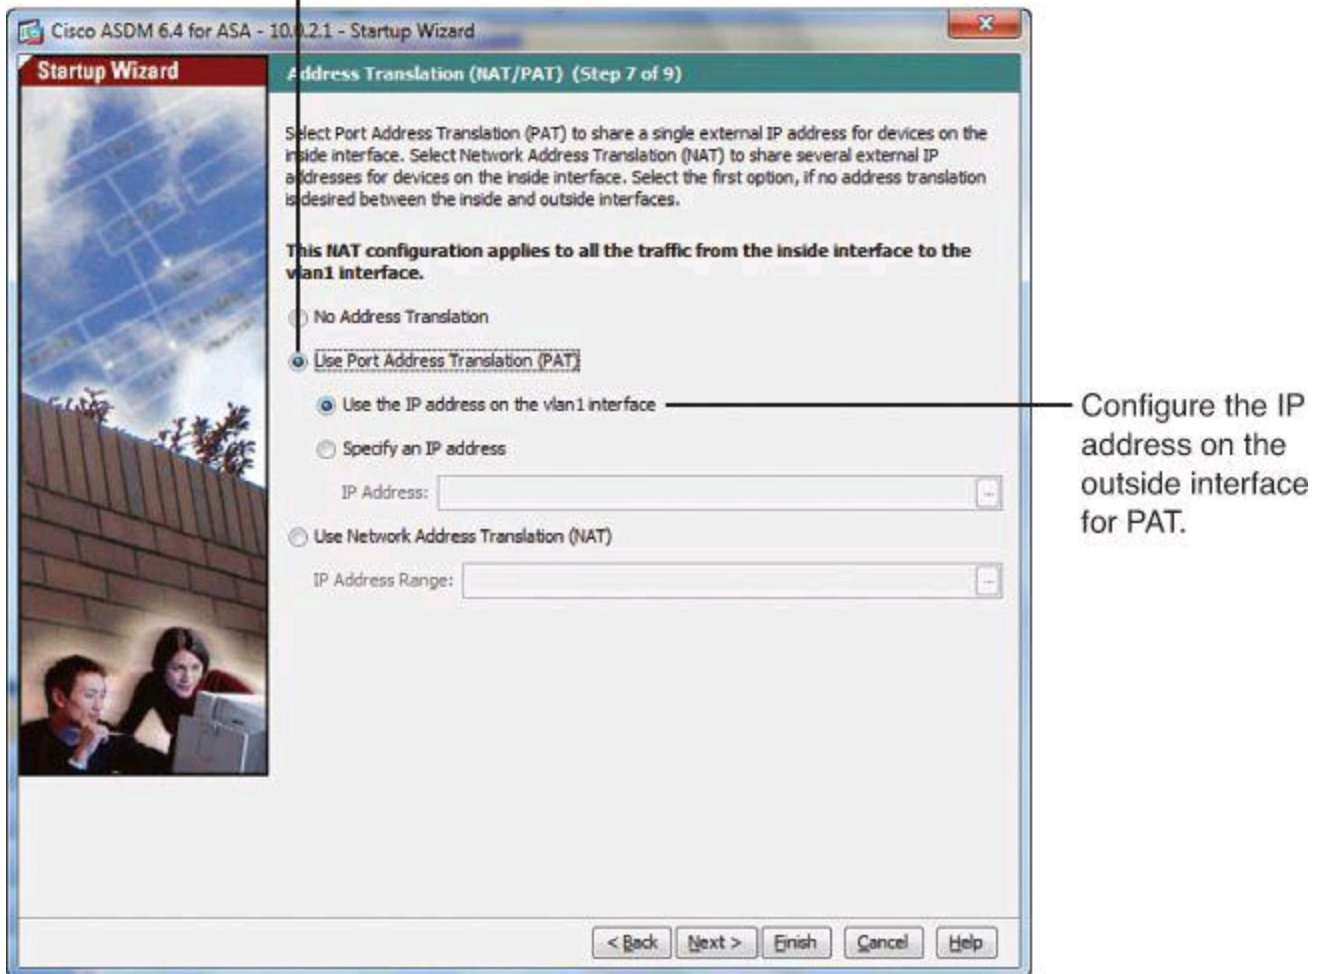
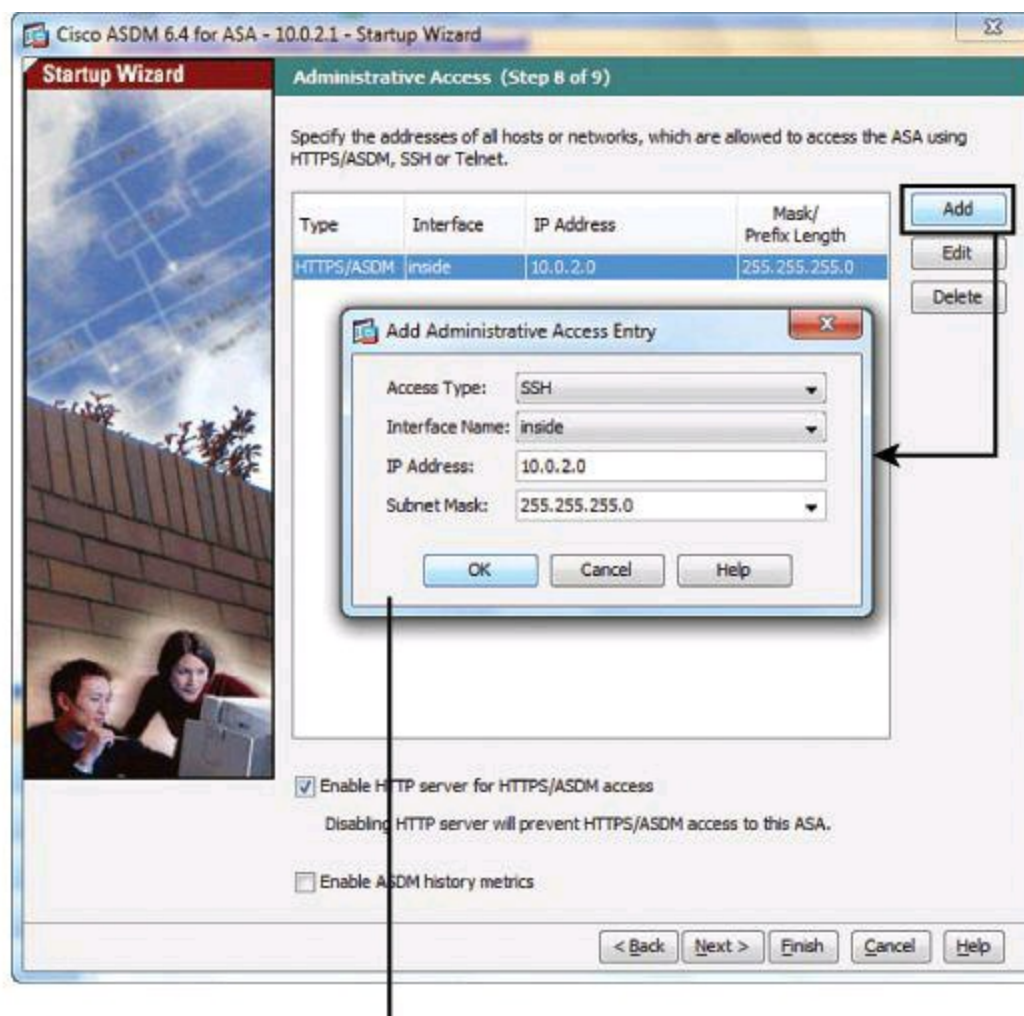


Figure 10-44. NAT and PAT Configuration from Cisco ASDM Startup Wizard

Step 8 of the wizard enables you to configure administrative access. The following options are available:

- Configure ASDM, Telnet, or SSH access. Clicking the Add button enables you to configure the IP address or range of addresses that are allowed to use the protocol to manage the Cisco ASA, as well as the interface where they are allowed.
- To enable a secure connection to an HTTP server to access Cisco ASDM, check the **Enable HTTP Server for HTTPS/ASDM Access** check box.
- To allow Cisco ASDM to collect and display statistics, check the **Enable ASDM History Metrics** check box.

In [Figure 10-45](#), SSH services are defined on the inside interface to allow the whole subnet 10.0.2.0/24 to have SSH access to the appliance.

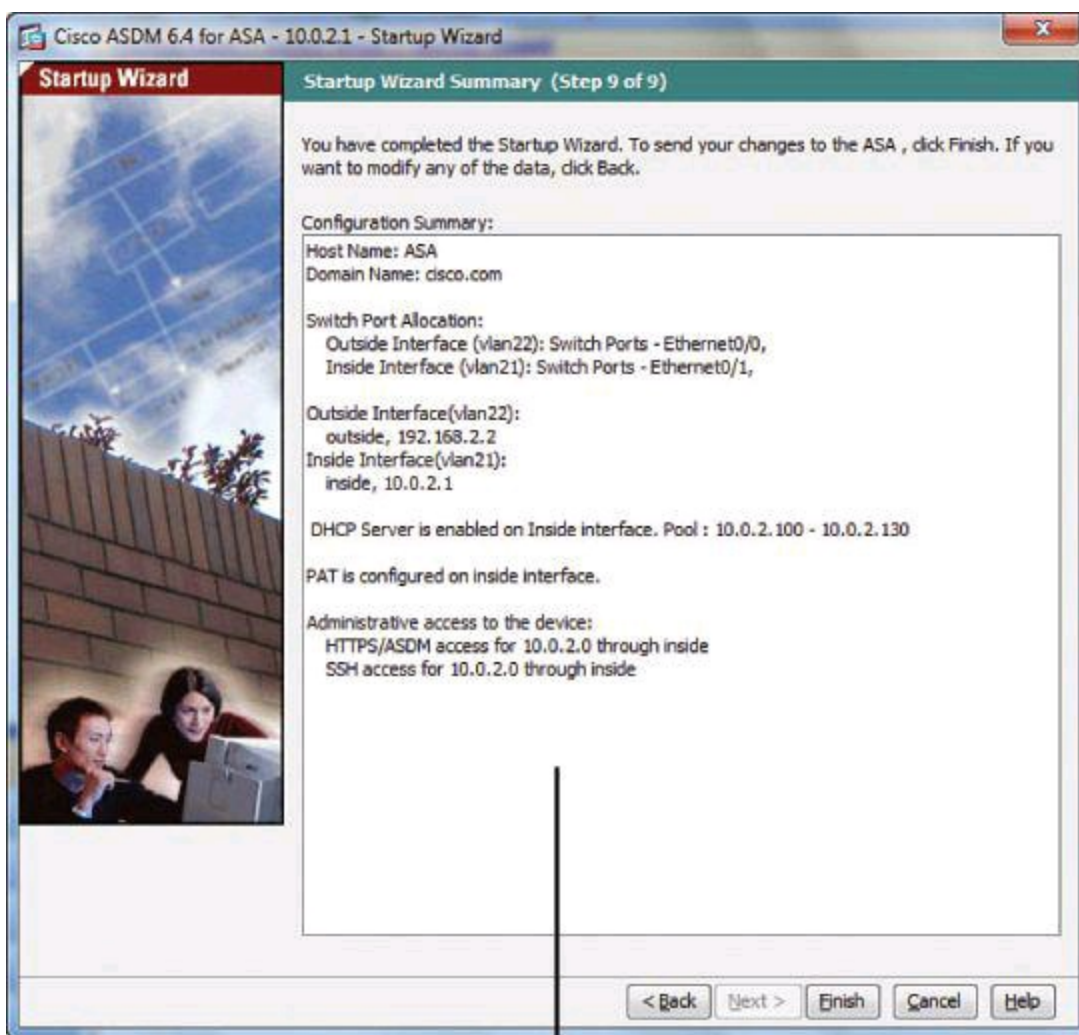


Configure access control for HTTP, SSH, or Telnet access.

Figure 10-45. Administrative Access Configuration from Cisco ASDM Startup Wizard

The final wizard screen, shown in [Figure 10-46](#), summarizes all of the configuration settings that you have made for the ASA. The following are your options:

- To change any of the settings in previous screens, click Back.
- Choose one of the following:
 - If you ran the Startup Wizard directly from a browser, when you click Finish, the configuration settings that you created through the wizard are sent to the Cisco ASA and saved in flash memory automatically.
 - If you ran the Startup Wizard from within Cisco ASDM, you must explicitly save the configuration in flash memory by choosing **File > Save Running Configuration to Flash**.



Review and click Finish.

Figure 10-46. Cisco ASDM Startup Wizard Summary

Step 3: Verify the Configuration Created by the Cisco ASDM Startup Wizard

Verification can be accomplished using the Cisco ASDM Device Dashboard tab from the Home page. This screen displays interface status and statistics, as well as overall information about device health.

The access rules and NAT rules that result from the Startup Wizard can be verified by navigating to the Cisco ASDM options shown in [Figure 10-47](#) and [Figure 10-48](#). Note that the screenshots shown in these figures were captured on a Cisco ASA 5505 running code 8.4(2). This explains the extra column titled User, in [Figure 10-47](#), that you might not be seeing on your Cisco ASA.

#	Enabled	Source	User	Destination	Service	Action	Hits	Logging
inside (1 implicit incoming rule)								
1		any		Any less secure networks	IP ip	Permit		
inside IPv6 (1 implicit incoming rule)								
1		any		Any less secure networks	IP ip	Permit		
outside (0 implicit incoming rules)								
outside IPv6 (0 implicit incoming rules)								
Global (1 implicit rule)								
1		any		any	IP ip	Deny		
Global IPv6 (1 implicit rule)								
1		any		any	IP ip	Deny		

Figure 10-47. Verifying Access Rules Created by Cisco ASDM Startup Wizard

#	Match Criteria: Original Packet					Action: Translated Packet			Options	De
	Source Intf	Dest Intf	Source	Destination	Service	Source	Destination	Service		
"Network Object" NAT (No rules)										
1	inside	outside	any	any	any	outside (P)	-- Original --	-- Original --		

Figure 10-48. Verifying NAT Rules Created by Cisco ASDM Startup Wizard

The last required step to enable connectivity through the appliance is to define a default route for outbound traffic. This task can be accomplished by navigating to **Configuration > Device Setup > Routing > Static Routes** and clicking **Add** to create a static route. The options include the outgoing interface, the destination prefix for the route, and the next hop.

In [Figure 10-49](#), a default route is created using the **any** keyword on the Network field. The Network field indicates the destination prefix of the route, and using **any** is similar to using 0.0.0.0 as a destination. The outgoing interface for traffic following this route is the outside interface.

Configure static route specifying outgoing interface and next hop.

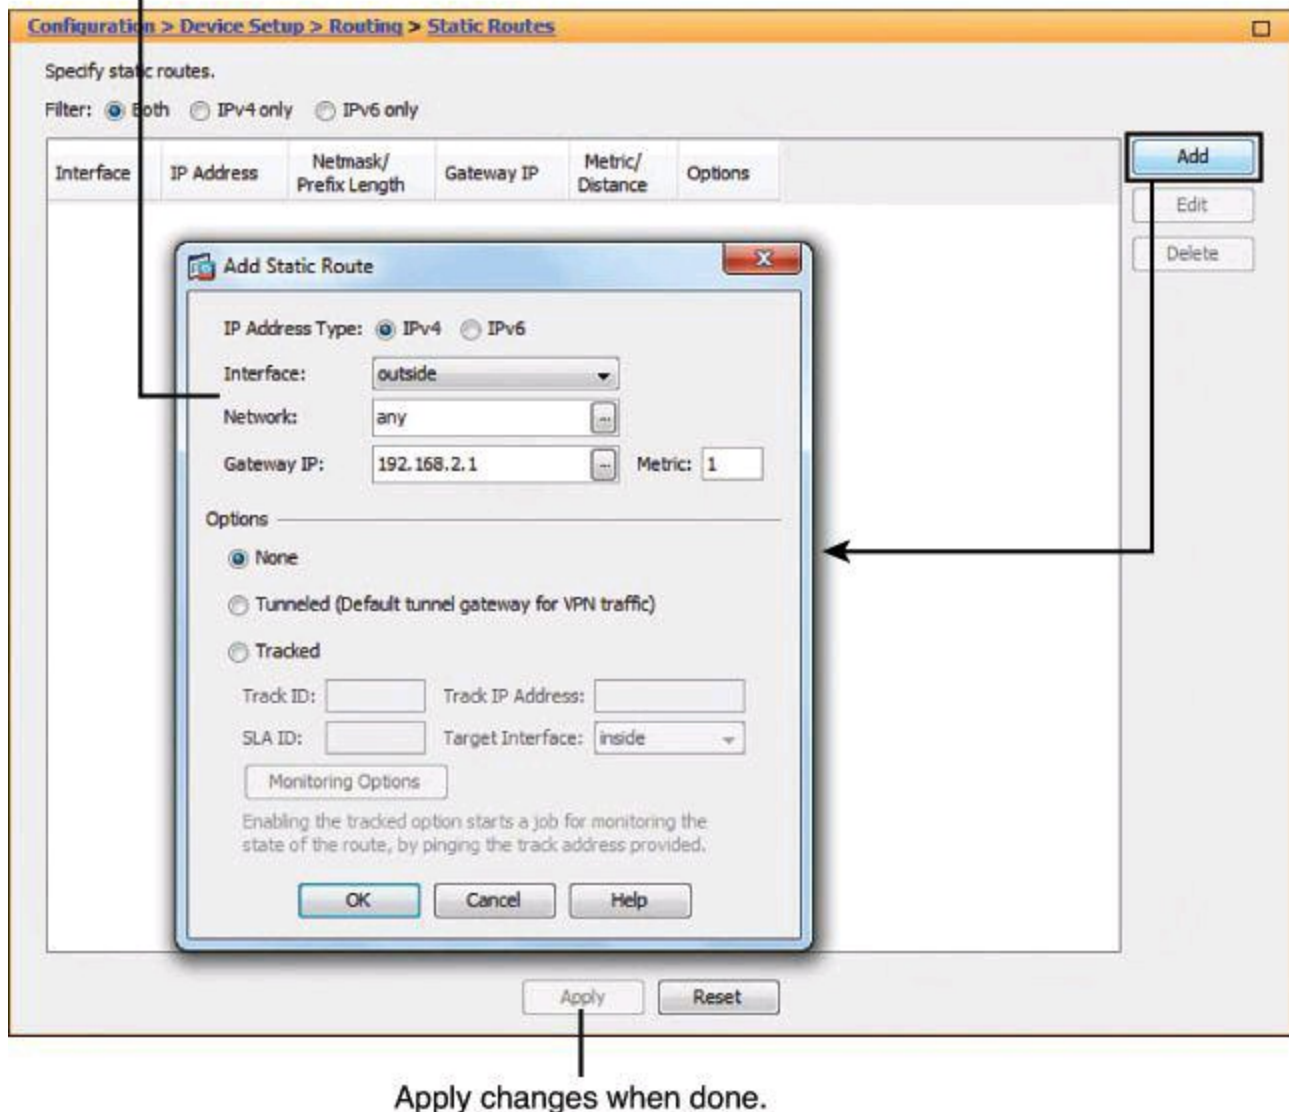


Figure 10-49. Adding a Static Route Using Cisco ASDM

Notice that this is a change made outside of the wizard, so it must be synchronized from Cisco ASDM to the appliance by clicking **Apply** at the bottom of the Static Routes window.

Step 4: Verify Firewall Activity Using the Packet Tracer Tool

The Packet Tracer tool provides packet tracing for packet sniffing and network fault isolation, as well as detailed information about the packets and how they are processed by the Cisco ASA. If a configuration command did not cause the packet to drop, Packet Tracer provides information about the cause in an easily readable manner.

In addition, you can use the Packet Tracer tool to trace the lifespan of a packet through the Cisco ASA to see whether the packet is operating correctly. This tool lets you do the following:

- Debug all packet drops in a production network
- Verify that the configuration is working as intended
- Show all rules applicable to a packet, along with the CLI commands that caused the rule addition
- Show a timeline of packet changes in a data path

- Inject tracer packets into the data path

To open the Packet Tracer, perform the following steps, shown in [Figure 10-50](#):

Step 1. In the main Cisco ASDM application window, navigate to **Tools > Packet Tracer**.

Step 2. The Cisco ASDM Packet Tracer dialog box opens.

Step 3. Choose the source interface for the packet trace from the Interface drop-down list.

Step 4. Specify the protocol type for the packet trace. Available protocol types include TCP, UDP, ICMP, and IP.

Step 5. Enter the source address for the packet trace in the Source IP Address field.

Step 6. Choose the source port for the packet trace from the drop-down list.

Step 7. Enter the destination IP address for the packet trace in the Destination IP Address field.

Step 8. Choose the destination port for the packet trace from the drop-down list.

Step 9. Click Start to trace the packet.

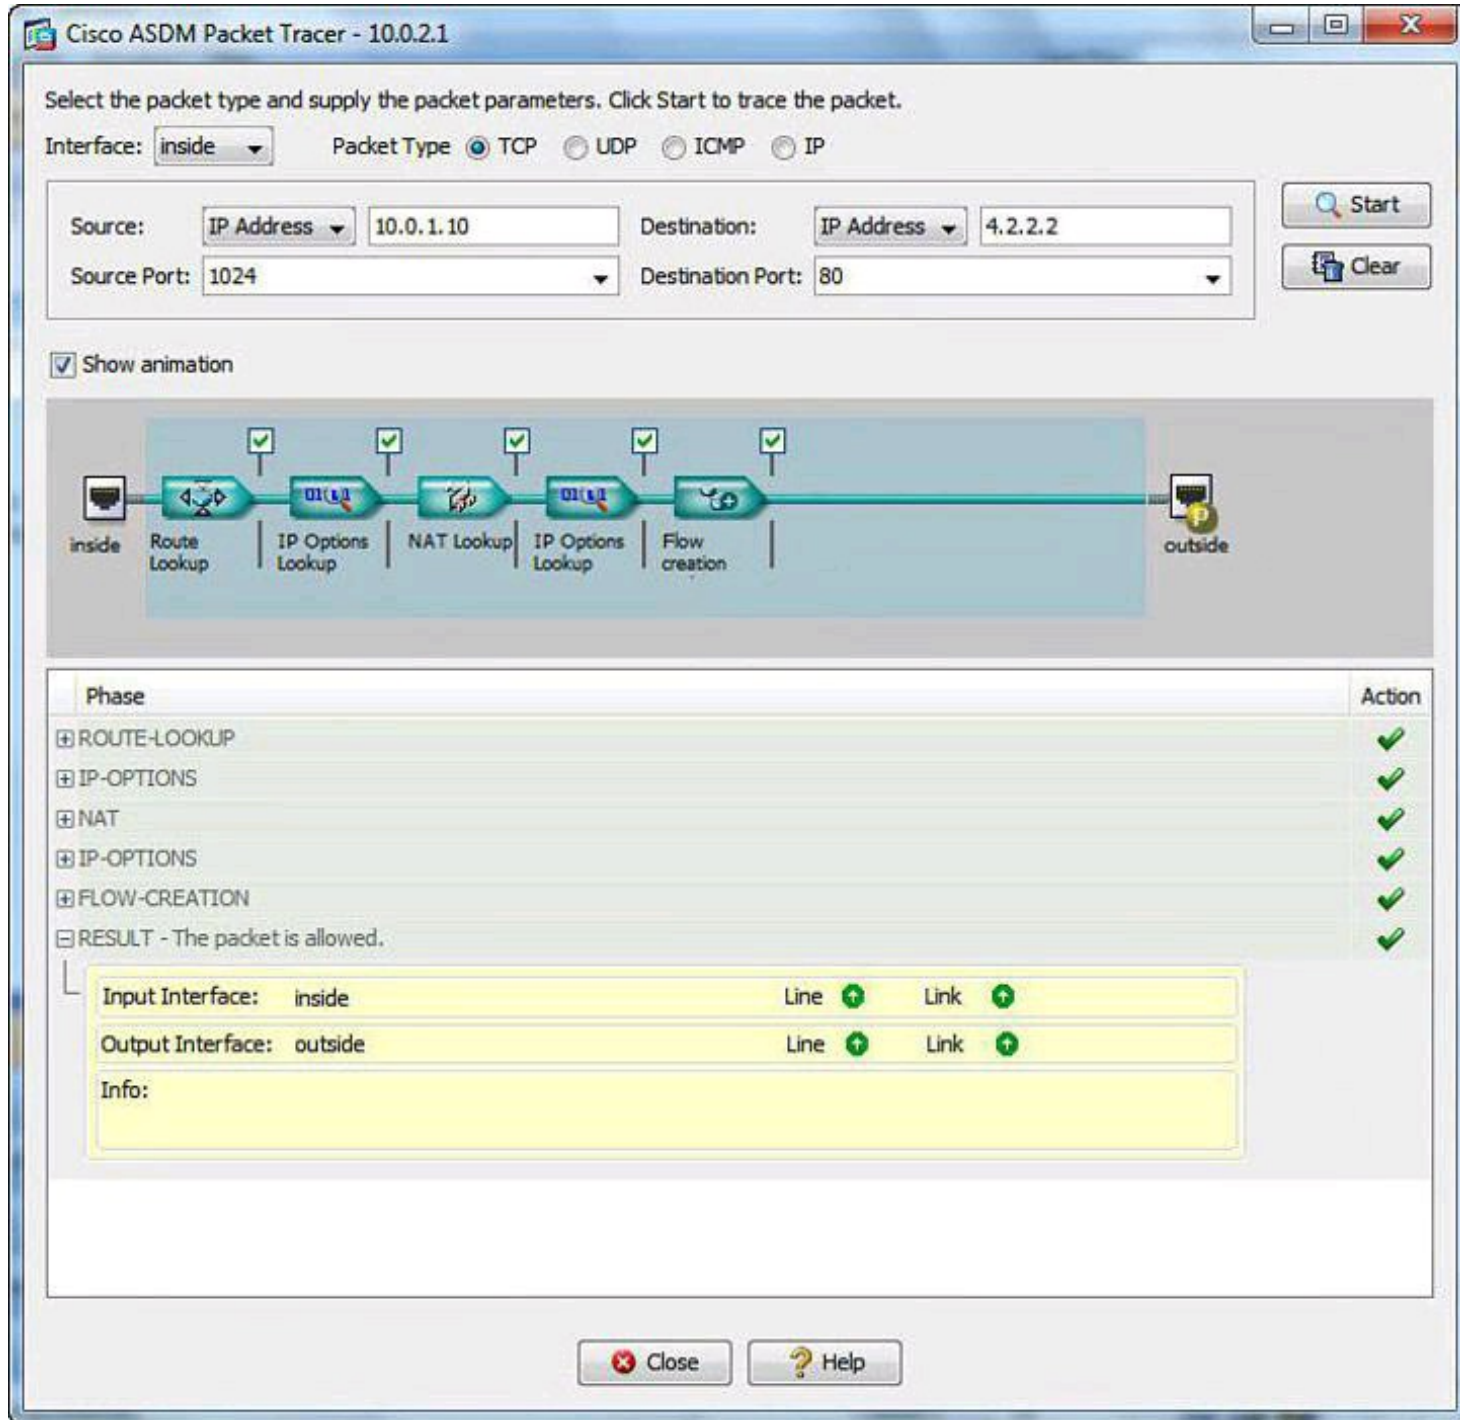


Figure 10-50. Packet Tracer Tool Using Cisco ASDM

The information display area shows detailed messages about the packet trace.

Note

To display a graphical representation of the packet trace, check the Show Animation check box.

Summary

In this chapter, you learned that a firewall is a set of rules designed to enforce an access control policy between two networks. You learned how to create firewall rules to implement your security policies. Cisco provides a range of firewall products that help you implement your security policies

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions.](#)”

1. Which of the following is not an action of zone-based policy firewalls?
 - a. drop
 - b. reset
 - c. pass
 - d. inspect
2. When using zone-based policy, for traffic to flow among all the interfaces in a router, all of the interfaces must be which of the following?
 - a. Configured with the CCP Wizard
 - b. A member of a security zone
 - c. Placed in a virtual template
 - d. A member of a dialer interface
3. What is the role of class maps?
 - a. They group interfaces to which security policies will be applied.
 - b. They allow you to specify a unidirectional firewall policy between two security zones.
 - c. They specify the actions to be taken when traffic matches the criteria.
 - d. They identify traffic and traffic parameters that a Cisco IOS Zone-Based Policy Firewall selects for policy application.
4. What programming mechanism is used to provide zone-base firewall services in a Cisco IOS router?
 - a. Adaptive Security Algorithm
 - b. Access Control Lists
 - c. Cisco Common Classification Policy Language
 - d. Modular Policy Framework
5. In CLI, the **zone-pair** command is used to associate together which of the following? (Choose two.)
 - a. zones
 - b. class-map
 - c. access-list
 - d. service-policy
 - e. interface
 - f. class-type

6. An interface belongs to a zone for which a zone pair with the self zone was created. No policies have yet been created. Which action will be applied when traffic originating on the interface that belongs to the zone is destined to the router itself?

- a. inspect**
- b. pass**
- c. drop**
- d. log**
- e. drop and log**

7. Multiple service policies can be applied to the same zone pair.

- a. True**
- b. False**

8. You want to apply a restrictive policy to traffic between interfaces that belong to the same zone in a zone-based policy firewall. Which of the following options can be used?

- a. Define the zone as the source and destination of a zone pair and create a policy for the zone pair**
- b. Create a new zone and change interface zone membership**
- c. Create a zone-based policy that applies ACLs to interfaces of the zone**
- d. Use ACLs on the interfaces on the zone**

9. Which Cisco Configuration Professional feature can be used to create zone-based firewalls?

- a. Advanced Firewall wizard only**
- b. Basic Firewall wizard only**
- c. Zone-based Firewall wizard**
- d. Both Advanced Firewall and Basic Firewall wizards**

10. Your Cisco ASA has multiple interfaces configured. There are two DMZ interfaces—DMZ1 with security level 51 and DMZ2 with security level 52. Which of the following is to be expected from this environment?

- a. All traffic between the two DMZs requires an explicit policy.**
- b. Traffic originating at DMZ2 and destined to DMZ1 is allowed.**
- c. Traffic originating at DMZ1 and destined to DMZ2 is allowed.**
- d. None of the above.**

11. Which condition would require a special configuration on Cisco ASA 8.3?

- a. Connectivity for through traffic**
- b. ARP traffic for devices on the same interfaces**
- c. Outbound connectivity**
- d. Connectivity between interfaces of the same security level**

12. What action is not available on Cisco ASA when using Modular Policy Framework?

- a. Sending SNMP traps
- b. Connection settings
- c. NetFlow-related actions
- d. QoS policies

13. Which function is not available using the Packet Tracer tool?

- a. Traffic simulation
- b. Testing ICMP traffic
- c. Capturing packets for further analysis
- d. Showing the rules applied to a packet flow

14. What is the significance of the **nameif** command with Cisco ASAs?

- a. An interface needs to have a name to pass traffic.
- b. Naming interfaces is useful for troubleshooting.
- c. The interface name appears in CDP packets.
- d. IP name resolution is based on name of interfaces.

15. Which statement is true regarding the default behavior of the Cisco ASA?

- a. Traffic is allowed to flow from a lower security level interface to a higher security level interface.
- b. Traffic is allowed to flow from a higher security level interface to a lower security level interface.
- c. Traffic is allowed to flow between interfaces of the same security level.
- d. Traffic is allowed to flow in and then out of the same interface.

Chapter 11. Intrusion Prevention Systems

This chapter describes the functions and operations of intrusion detection systems (IDS) and intrusion prevention systems (IPS). This chapter will introduce you to

- The fundamentals of intrusion prevention, comparing IDS and IPS
- The building blocks of IPS, introducing the underlying technologies and deployment options
- The use of signatures in intrusion prevention, highlighting the benefits and drawbacks
- The need for IPS alarm monitoring, evaluating the options for event managers
- Analyzing the design considerations in deploying IPS

IPS Fundamentals

Intrusion detection system (IDS) and intrusion prevention system (IPS) solutions form an integral part of a robust network defense solution. Maintaining secure network services is a key requirement of a profitable IP-based business. To show how these systems work, this chapter uses Cisco products and technologies as examples.

Introducing IDS and IPS

The evolution of threats and business models requires a new approach to intrusion prevention:

- **Targeted, mutating, stealth threats are increasingly difficult to detect.** Intrusion prevention needs to be distributed and embedded end to end along the “path of intrusion.” Shared threat and security intelligence is a key feature.
- **Attackers have insidious motivations and exploit high-impact targets, often for financial benefit or economic and political reasons.** The impact on business continuity is immense. More than ever, intrusion prevention needs to be accurate, providing a timely and precise response only to threats relevant to the infrastructure, application, and business environments of the organization.
- **Attackers are taking advantage of new ways of communication.** Borderless networks provide a launching pad for high-impact intrusion from mobile users to consumer devices (smartphones, PDAs, tablet devices, entertainment devices, book readers, and so on), exploiting data wherever it may be stored (user’s personal device, enterprise data centers, cloud service providers, and so on). The “army of intrusion sensors” requires flexible deployment options, from appliances to hardware modules on existing network elements to software-based and virtualized form factors.

Traditionally, an IDS and an IPS work together to provide a network security solution. An IDS captures packets in real time, processes them, and can respond to threats, but it works on *copies* of data traffic to detect suspicious activity by using signatures. This is called *promiscuous mode*. In the process of detecting malicious traffic, an IDS allows some malicious traffic to pass before the IDS can respond to protect the network. An IDS analyzes a copy of the monitored traffic rather than the actual forwarded packet. The advantage of operating on a copy of the traffic is that the IDS does not affect the packet flow of the forwarded traffic. The disadvantage of operating on a copy of the traffic

is that the IDS cannot stop malicious traffic from single-packet attacks from reaching the target system before the IDS can apply a response to stop the attack. An IDS often requires assistance from other networking devices, such as routers and firewalls, to respond to an attack.

An IPS works inline in the data stream to provide protection from malicious attacks in real time. This is called *inline mode*. Unlike an IDS, an IPS does not allow packets to enter the trusted side of the network. An IPS monitors traffic at Layer 3 and Layer 4 to ensure that the traffic's headers, states, and so on are those specified in the protocol suite. However, the IPS sensor analyzes the payload of the packets at Layer 2 to Layer 7 for more sophisticated embedded attacks that might include malicious data. This deeper analysis lets the IPS identify, stop, and block attacks that would normally pass through a traditional firewall device. When a packet comes in through an interface on an IPS, that packet is not sent to the outbound or trusted interface until the packet has been determined to be clean. An IPS builds upon previous IDS technology; Cisco IPS platforms use a blend of detection technologies, including profile-based intrusion detection, signature-based intrusion detection, and protocol analysis intrusion detection.

The key to differentiating an IDS from an IPS is that an IPS responds immediately and does not allow any malicious traffic to pass, whereas an IDS allows malicious traffic to pass before it can respond. IDSs typically are IPSs that are configured to work on a copy of the traffic (promiscuous mode) instead of working on traffic going through it (inline mode). Because IDS devices are usually IPS devices configured differently, in this chapter we often use the term IPS, which, in our generic conversation on the topic, will also encompass IDS, unless indicated otherwise.

IDS:



- Analyzes copies of the traffic stream
- Does not slow network traffic
- Allows some malicious traffic into the network

IPS:

- Works inline in real time to monitor Layer 2 through Layer 7 traffic and content
- Needs to be able to handle network traffic
- Prevents malicious traffic from entering the network

IDS and IPS technologies share several characteristics:

- IDS and IPS technologies are deployed as sensors. An IDS or an IPS sensor can be any of the following devices:
 - A router configured with Cisco IOS IPS Software
 - An appliance specifically designed to provide dedicated IDS or IPS services
 - A network module installed in a Cisco adaptive security appliance, switch, or router
- IDS and IPS technologies typically monitor for malicious activities in two spots:
 - **Network:** To detect attacks against the network, including attacks against hosts

and devices, using network IDS and network IPS.

- **Hosts:** To detect attacks launched from or against target machines, using host IPS (HIPS). Host-based attacks are detected by reading security event logs, checking for changes to critical system files, and checking system registries for malicious entries.
 - IDS and IPS technologies use signatures to detect patterns of misuse in network traffic. A signature is a set of rules that an IDS or IPS uses to detect typical intrusive activity. Signatures are usually chosen from a broad cross section of intrusion detection signatures, and can detect severe breaches of security, common network attacks, and information gathering.
 - IDS and IPS technologies look for the following general patterns of misuse:
 - **Atomic pattern:** In an atomic pattern, an attempt is made to access a specific port on a specific host, and malicious content is contained in a single packet. An IDS is particularly vulnerable to an atomic attack because until the IDS finds the attack, malicious single packets are being allowed into the network. An IPS prevents these packets from entering at all.
 - **Composite pattern:** A composite pattern, also referred to as *compound*, is a sequence of operations distributed across multiple hosts over an arbitrary period of time.
-

Inline Sensors

Sensors, even inline, might not be completely successful at dropping packets of an attack. An attack might be underway, if only partially, before an inline sensor even starts to drop packets matching a composite pattern signature. The drop action is much more effective for atomic signatures because the sensor makes a single-packet match.

[Figure 11-1](#) shows a sensor deployed in IDS mode and a sensor deployed in IPS mode.

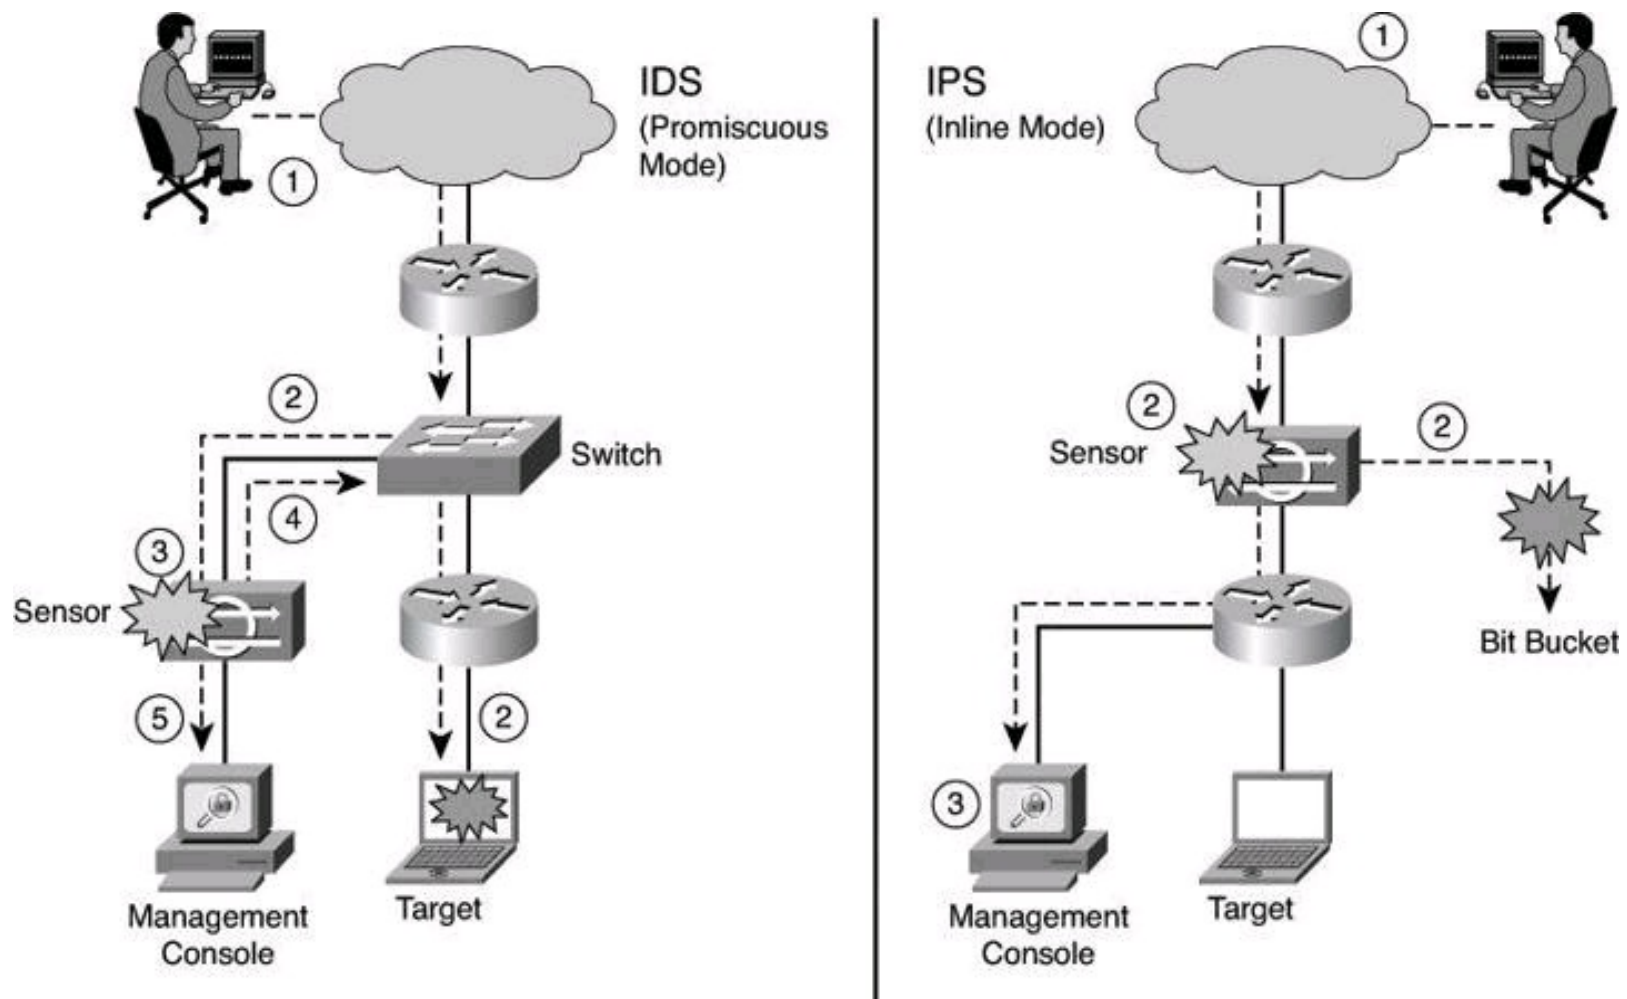


Figure 11-1. IDS and IPS Operational Differences

The following are the steps that occur when an attack is launched in an environment monitored by an IDS:

Step 1. An attack is launched on a network that has a sensor deployed in IDS mode.

Step 2. The switch sends copies of all packets to the IDS sensor (configured in promiscuous mode, which is explained later in this section) to analyze the packets. At the same time, the target machine experiences the malicious attack.

Step 3. The IDS sensor, using a signature, matches the malicious traffic to the signature.

Step 4. The IDS sensor sends to the switch a command to deny access to the malicious traffic.

Step 5. The IDS sends an alarm to a management console for logging and other management purposes.

The following are the steps that occur when an attack is launched in an environment monitored by an IPS:

Step 1. An attack is launched on a network that has a sensor deployed in IPS mode (configured in inline mode, which is explained later in this section).

Step 2. The IPS sensor analyzes the packets as soon as they come into the IPS sensor interface. The IPS sensor, using signatures, matches the malicious traffic to the signature and the attack is stopped immediately. Traffic in violation of policy can be dropped by an

IPS sensor.

Step 3. The IPS sensor can send an alarm to a management console for logging and other management purposes.



Key
Topic

Promiscuous Versus Inline Mode

A sensor can be deployed either in promiscuous mode or inline mode. In promiscuous mode, the sensor receives a copy of the data for analysis, while the original traffic still makes its way to its ultimate destination. By contrast, a sensor working inline analyzes the traffic live and therefore can actively block the packets before they reach their destination.

It is worth mentioning that Cisco appliances such as the Cisco ASA AIP-SSM (discussed later in this chapter), although advertised as IPS devices, can work either in promiscuous mode or in inline mode.



Key
Topic

Management Console

The term *management console*, used in this chapter and seen in [Figure 11-1](#), refers to a separate workstation equipped with software to configure, monitor, and report on events. The section “Monitoring Cisco IOS IPS Alarms and Event Managers” introduces some of the Cisco IPS management solutions.

[Table 11-1](#) lists some of the advantages and limitations of deploying platform sensors in promiscuous mode; in other words, as an IDS.

Table 11-1. Advantages and Limitations of Deploying a Sensor in Promiscuous Mode: IDS

Advantage	Limitation
Deploying the IDS sensor does not have any impact on the network (latency, jitter, and so on).	IDS sensor response actions cannot stop the trigger packet and are not guaranteed to stop a connection. IDS sensor response actions are typically better at stopping an attacker than at stopping a specific attack itself.
The IDS sensor is not inline and, therefore, a sensor failure cannot affect network functionality.	IDS sensor response actions are less helpful in stopping email viruses and automated attackers such as worms, and IDS sensors are more vulnerable to network evasion techniques.
Overrunning the IDS sensor with data does not affect network traffic; however, it does affect the capability of the IDS sensor to analyze the data.	Users deploying IDS sensor response actions must have a well-thought-out security policy combined with a good operational understanding of their IDS deployments. Users must spend time to correctly tune IDS sensors to achieve expected levels of intrusion detection.
	Being out of band (OOB), IDS sensors are more vulnerable to network evasion techniques, which are designed to totally conceal an attack.

[Table 11-2](#) lists some of the advantages and limitations of deploying a platform sensor in inline mode; in other words, as an IPS.

Table 11-2. Advantages and Limitations of Deploying a Sensor in Inline Mode: IPS

Advantage

Limitation

You can configure an IPS sensor to perform a packet drop that can stop the trigger packet, the packets in a connection, or packets from a source IP address.

Because an IPS sensor must be inline, IPS sensor errors or failure can have a negative effect on network traffic.

Being inline, an IPS sensor can use stream normalization techniques, such as fragmentation reassembly, to check the validity of the transmission, thus reducing or eliminating many of the network evasion capabilities that exist.

Overrunning IPS sensor capabilities with too much traffic negatively affects the performance of the network.

Users deploying IPS sensor response actions must have a well-thought-out security policy combined with a good operational understanding of their IPS deployments.

An IPS sensor will affect network timing because of latency, jitter, and so on. An IPS sensor must be appropriately sized and implemented so that time-sensitive applications, such as VoIP, are not negatively affected.

So, IDS or IPS? Why Not Both?

When designing a distributed IPS architecture, both IDS and IPS deployment modes can be used to improve the accuracy, intelligence, response, and performance of the architecture. When combined, IDSs in promiscuous mode and IPSs in inline mode complement each other to accomplish this objective. [Figure 11-2](#) illustrates an example, described next.

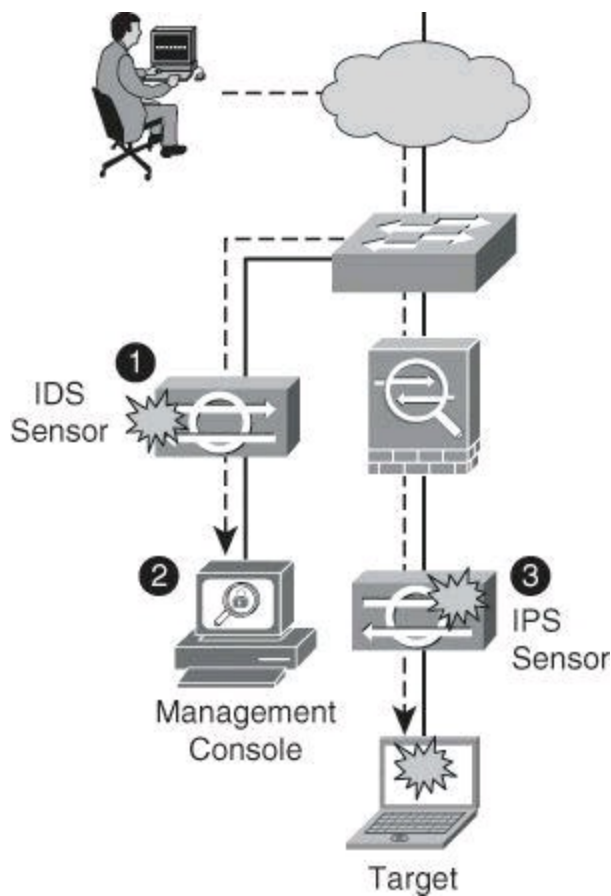


Figure 11-2. Using Both IDS and IPS Devices

The IDS sensor in front of the firewall is deployed in promiscuous mode to monitor traffic in the untrusted network. It is likely that this IDS sensor will have more visibility over a wide variety of events and malicious traffic that have the potential to become security incidents. The internal network is protected initially by the firewall, and the role of the IDS sensor is to gather security intelligence in order to define a baseline of current and relevant threats. This information is then correlated and fed into the detection capabilities of the internal IPS sensor.

The IPS sensor behind the firewall can therefore focus on monitoring the internal network and detect incidents originated inside. It can also implement a more restrictive policy for external traffic, while remaining accurate, using the information gathered and baselined by the IDS sensor.

Alarm Types

As shown in the previous figures, IDS and IPS sensors produce alarms. The capability of IDS and IPS sensors to accurately detect an attack or a policy violation and generate an alarm is critical to the functionality of the sensors. Attacks can generate the following types of alarms:

- **False positive:** A false positive is an alarm triggered by normal traffic or a benign action. Consider this scenario: A signature exists that generates alarms if the enable password of any network devices is entered incorrectly. A network administrator attempts to log in to a Cisco router but enters the wrong password. The IDS cannot distinguish between a rogue user and the network administrator, and it generates an alarm.
- **False negative:** A false negative occurs when a signature is not fired when offending traffic is detected. Offending traffic could range from a corporate user pinging too aggressively a target on the Internet, to attacks against corporate web servers. A false

negative should be considered a software bug if the IDS and IPS have a signature that has been designed to detect the offending traffic but failed to do so. In this situation, the failure to trigger on the offending traffic should be reported to the vendor of the IPS product.

- **True positive:** A true positive occurs when an IDS and IPS signature is correctly fired and an alarm is generated when offending traffic is detected. For example, consider a Unicode attack. Cisco IPS sensors have signatures that detect Unicode attacks against Microsoft Internet Information Services (IIS) web servers. If a Unicode attack is launched against Microsoft IIS web servers, the sensors detect the attack and generate an alarm.
- **True negative:** A true negative occurs when a signature is not fired when non-offending traffic is captured and analyzed. In other words, the sensor does not fire an alarm when it captures and analyzes “normal” network traffic.

[Figure 11-3](#) provides a summary of the alarm types. To understand the terminology, think in terms of the question, “Was the alarm triggered?” A positive means that the alarm was triggered and a negative means that the alarm was not triggered. Thus the expression *false alarm*, which is the same as *false positive* (positive because the alarm was triggered, but false because the intrusion did not happen or the intrusion was not detected by the sensor).

		Actual Offending Traffic?	
		NO	YES
Alarm Triggered?	NO	True Negative	False Negative
	YES	False Positive	True Positive

Figure 11-3. Making Sense of Alarm Types Terminology

Note

Positive or negative refers to whether the alarm was triggered. Positive designates an alarm has been triggered, and negative designates an alarm has not been triggered. *True* or *false* is more complex; this answers the question, “Was the result to trigger (or not trigger) the alarm the right decision?” If the action taken by the sensor was right, then the result is true, and if the sensor took the wrong decision, the result is false. As an example, if malicious traffic travels through your network and the IPS sensor fails to detect it, it is a false negative. The *negative* refers to the alarm not triggering. The *false* means not triggering was the wrong decision for the IPS sensor to take.

Intrusion Prevention Technologies

Multiple detection technologies are typically used to provide an effective intrusion detection architecture. [Table 11-3](#) summarizes the advantages and limitations of the various types of IDS and IPS sensors available. These will be discussed in greater detail in this section.

Table 11-3. Types of IDS and IPS Sensors

Sensor Type	Advantages	Limitations
Signature based	Easy configuration Fewer false positives Good signature design	No detection of unknown signatures Initially a lot of false positives Signatures must be created, updated, and tuned
Policy based	Simple and reliable Customized policies Can detect unknown attacks	Generic output Policy must be created
Anomaly based	Easy configuration Can detect unknown attacks	Difficult to profile typical activity in large networks Traffic profile must be constant
Reputation based	Leverages local, enterprise, and global correlation Improved accuracy and relevancy	More prone to false positives and false negatives Requires timely updates

Signature-Based IDS/IPS

A signature-based approach is usually the starting point of an effective intrusion detection architecture. A *signature* is a set of rules that an IDS and an IPS use to detect typical intrusive activity, such as denial of service (DoS) attacks. You can easily install signatures using IDS and IPS management software such as Cisco IPS Device Manager (IDM) and Cisco IPS Express Manager (IME). Sensors allow you to modify existing signatures and define new ones.

A signature-based IDS or IPS sensor looks for specific, predefined patterns (signatures) in network traffic. It compares the network traffic to a database of known attacks, and triggers an alarm or prevents communication if a match is found. The signature can be based on a single packet or a sequence of packets. New attacks that do not match a signature do not result in detection. For this reason, the signature database needs to be constantly updated.

As sensors scan network packets, they use signatures to detect known attacks and respond with predefined actions. A malicious packet flow has a specific type of activity and signature, and an IDS or IPS sensor examines the data flow using many different signatures. When an IDS or IPS sensor matches a signature with a data flow, the sensor takes action, such as logging the event or sending an alarm to IDS or IPS management software, such as Cisco Configuration Professional (CCP).

Signature-based intrusion detection can produce false positives because certain normal network activity can be misinterpreted as malicious activity. For example, some network applications or operating systems may send out numerous Internet Control Message Protocol (ICMP) messages, which a signature-based detection system might interpret as an attempt by an attacker to map out a network segment. You can minimize false positives by tuning your sensors. You can tune built-in signatures (tuned signatures) by adjusting the many signature parameters.

Note

Protocol analysis–based intrusion detection relies on signature-based intrusion detection, where the signature performs a check to ensure that the data unit header, flags, payload, and so on respect the protocol.

Signature-based pattern matching is an approach that is rigid but simple to employ. In most cases, the pattern is matched against only if the suspect packet is associated with a particular service or, more precisely, destined to and from a particular port. This matching technique helps to lessen the amount of inspection done on every packet. However, it makes it more difficult for systems to deal with protocols that do not reside on well-defined ports, such as Trojan horses and their associated traffic, which can move at will.

At the initial stage of incorporating signature-based IDS or IPS, before the signatures are tuned, there can be many false positives (traffic generating an alert even though it is no threat to the network). After the system is tuned and adjusted to the specific network parameters, there will be fewer false positives than with the policy-based approach.

Policy-Based IDS/IPS

In policy-based systems, the IDS or IPS sensor is configured based on the network security policy. You must create the policies used in a policy-based IDS or IPS. Any traffic detected outside the policy will generate an alarm or will be dropped. Creating a security policy requires detailed knowledge of the network traffic and is a time-consuming task.

Policy-based signatures use an algorithm to determine whether an alarm should be fired. Often, policy-based signature algorithms are statistical evaluations of the traffic flow. For example, in a policy-based signature used to detect a port sweep, the algorithm issues an alarm when the threshold number of unique ports is scanned on a particular machine. Policy-based signature algorithms can be designed to analyze only specific types of packets (for example, SYN packets, where the SYN bit is turned on during the handshaking process at the beginning of the session).

The policy itself might require tuning. For example, you might have to adjust the threshold level of certain types of traffic so that the policy conforms to the utilization patterns on the network that it is monitoring. Policies can be used to look for very complex relationships.

Anomaly-Based IDS/IPS

Anomaly-based (or *profile-based*) signatures typically look for network traffic that deviates from what is seen “normally.” The biggest issue with this methodology is that you first must define what *normal is*. If during the learning phase your network is the victim of an attack and you fail to identify it, the anomaly-based IPS will interpret that malicious traffic as normal, and no alarm will be triggered the next time this same attack takes place. Some systems have hard-coded definitions of normal traffic patterns and, in this case, could be considered heuristic-based systems.

Normal behavior is typically defined based on traffic patterns, traffic and protocol mix, traffic volumes, and other criteria. This is called statistical baselining. A second approach focuses on traffic that deviates from protocol standards. This is called protocol anomaly baselining.

The technique used by anomaly-based IDS/IPS systems is also referred as *network behavior*

Reputation-Based IPS

A more recent approach to intrusion detection is the reputation-based IPS. This technique uses reputation analysis for various traffic descriptors, such as IP addresses, URLs, Domain Name System (DNS) domains, and others. This typically translates into reputation filters, sometimes known as white lists or black lists, that round up a signature-based system by filtering known malicious sources, destinations, or application components.

A reputation-based approach typically requires communication between the sensor and the source of the reputation information. Sources can be local to the device, which builds reputation based on the history of incidents known to the local device. However, the true power of reputation-based systems lies in the correlation of reputation information with enterprise knowledge and even global knowledge, usually provided by early warning systems, honeypots, and managed security intelligence services. Cisco Security Intelligence Operations, introduced later in this chapter, is a reputation-based IPS.

Honeypots

Honeypots can also be considered to be a form of IDS. Honeypot systems use a dummy server to attract attacks. The purpose of the honeypot approach is to distract attacks away from real network devices. By staging different types of vulnerabilities in the honeypot server, you can analyze incoming types of attacks and malicious traffic patterns. You can use this analysis to tune your sensor signatures to detect new types of malicious network traffic.

Honeypot systems are used in production environments, typically by large organizations that are considered interesting targets for hackers, such as financial enterprises, governmental agencies, and so on. Also, antivirus and other security vendors tend to use them for research.

Many security experts preach the use of honeypots as an early warning system to be deployed with your IDS/IPS system, not in lieu of. Honeyd is an example of popular open source honeypot software. Although honeypots are often found as dedicated servers, it is possible to set up virtual honeypots using VMware or Windows Virtual PC. Keep in mind that should the honeypot be successfully hacked and used as a launching platform for an attack on a third party, the honeypot's owner could incur downstream liability.

IPS Attack Responses

Detection is only part of the role of an IPS. You should define and plan the reaction capabilities of sensors by defining the response that is triggered upon detection, no matter what detection techniques (or combination of techniques) you use.

When an IPS sensor detects malicious activity, it can choose from any or all of the following actions:

- **Deny Attacker Inline:** This action terminates the current packet and future packets from this attacker address for a specified period of time. The sensor maintains a list of the

attackers currently being denied by the system. You can remove entries from the list or wait for the timer to expire. The timer is a sliding timer for each entry. Therefore, if attacker A is currently being denied but issues another attack, the timer for attacker A is reset, and attacker A remains on the denied attacker list until the timer expires. If the denied attacker list is at capacity and cannot add a new entry, the packet is still denied.

- **Deny Connection Inline:** This action terminates the current packet and future packets on this TCP flow. The packet for this flow was dropped. This is also referred to as *deny flow*.
- **Deny Packet Inline:** This action terminates the packet. The packets are dropped.
- **Log Attacker Packets:** This action starts IP logging on packets that contain the attacker address and sends an alert. This action causes an alert to be written to the event store, which is local to the Cisco IOS router, even if the produce-alert action is not selected. Produce alert is discussed later in a bullet.
- **Log Pair Packets:** This action starts IP logging on packets that contain the attacker and victim address pair. This action causes an alert to be written to the event store, even if the produce-alert action is not selected.
- **Log Victim Packets:** This action starts IP logging on packets that contain the victim address and sends an alert. This action causes an alert to be written to the event store, even if the produce-alert action is not selected.
- **Produce Alert:** This action writes the event to the event store as an alert.
- **Produce Verbose Alert:** This action includes an encoded dump of the offending packet in the alert. This action causes an alert to be written to the event store, even if the produce-alert action is not selected.
- **Request Block Connection:** This action sends a request to a blocking device to block this connection.
- **Request Block Host:** This action sends a request to a blocking device to block this attacker host.
- **Request SNMP Trap:** This action sends a request to the notification application component of the sensor to perform Simple Network Management Protocol (SNMP) notification. This action causes an alert to be written to the event store, even if the produce-alert action is not selected.
- **Reset TCP Connection:** This action sends TCP resets to hijack and terminate the TCP flow.

Caution

Responses to deny inline actions (packet, connection, or attacker) are not available on sensors in IDS mode.

Note

Packets that are dropped (denied) based on false alarms can result in network disruption if the dropped packets are required for mission-critical applications downstream of the IPS

sensor. Therefore, do not be overly aggressive when assigning the deny-action to signature. Also, “deny” discards the packet without sending a reset. Cisco recommends using “deny and reset” with alarm.

Also note that IP logging and verbose alert traces use a common capture file writing code called libpcap. This is the same format used by the famous packet-capture tool Wireshark (formerly Ethereal); by Snort, a famous freeware IDS; by NMAP, a well-known fingerprinting tool; and by Kismet, a famous wireless sniffing tool.

You can use the reset TCP connection action in conjunction with deny-packet and deny-flow actions. However, deny-packet and deny-connection actions do not automatically cause reset TCP connection actions to occur.

IPS Anti-Evasion Techniques

Attackers will do their best to get around your intrusion prevention architecture. Multiple IPS evasion techniques allow them to deploy exploits using a stealth approach, one that often renders IPS sensors unable to detect and prevent intrusion. These techniques include the following:

- **Traffic fragmentation:** Any evasion attempt where the attacker splits malicious traffic, hoping to avoid detection or filtering by confusing the network IPS sensor reassembly methods, bypassing the network IPS sensor if it does not perform any reassembly at all, or reordering split data if the network IPS sensor does not correctly order it in the reassembly process.
- **Traffic substitution:** The attacker attempts to evade detection by substituting payload data with other data in a different format but with the same meaning. If the IPS sensor does not recognize the true meaning of data, and only looks for data in a particular format, it may miss such malicious payloads. Examples include the use of Unicode strings, exploiting case sensitivity, substitution of spaces with tabs, and others.
- **Protocol-level misinterpretation:** The attacker attempts to evade detection by causing the network IPS sensor to misinterpret the end-to-end meaning of network protocols and see traffic differently from what will actually be seen and processed by the target. Consequently, the sensor will either ignore traffic that should not be ignored or vice versa.
- **Timing attacks:** The attacker attempts to evade detection of correlating signatures by performing their actions more slowly, not exceeding the thresholds inside the time windows that these signatures use to correlate different packets together.
- **Encryption and tunneling:** The attacker attempts to evade detection by encrypting packets. Sensors monitor the network and capture the packets as they traverse the network. Network-based sensors rely on the data that is being transmitted in plaintext. When packets are encrypted, the sensor captures the data but is unable to decrypt it and cannot perform meaningful analysis. This type of evasive technique assumes that the attacker has already established a secure session with the target network or host.
- **Resource exhaustion:** In this less subtle method of evading detection, the attacker relies on extreme resource consumption. It does not matter if such a denial is against the device or the personnel who are managing the device. Specialized tools can be used to create a

tremendous number of alarms that consume the resources of the IPS device and prevent attacks from being logged.

The following anti-evasion features are available on Cisco IPS sensors:

- Complete session reassembly that supports the string and service engines that must examine a reliable byte stream between two network endpoints
- Data normalization (deobfuscation) inside service engines, where all signatures convert network traffic data into a normalized, canonical form before comparing it to the signature-matching rules
- IP Time to Live (TTL) analysis and TCP checksum validation to guard against end-to-end protocol-level traffic interpretation
- Configurable intervals for correlating signatures, or the use of an external correlation that does not require real-time resources
- Inspection of traffic inside Generic Routing Encapsulation (GRE) tunnels to prevent evasion through tunneling
- Smart and dynamic summarization of events to guard against too many alarms for high event rates

[Table 11-4](#) summarizes the evasion methods and the corresponding Cisco IPS anti-evasion features that just described.

Table 11-4. Anti-Evasion Techniques Used by Cisco IPS

Evasion Method	Cisco IPS Anti-Evasion Features
Traffic fragmentation	Complete session reassembly in string and service engines
Traffic substitution and insertion	Data normalization (deobfuscation) in service engines
Protocol-level misinterpretation	IP TTL analysis TCP checksum validation
Timing attacks	Configurable intervals Use of third-party SIEM system for correlation
Encryption and tunneling	GRE tunnel inspection
Resource exhaustion	Smart dynamic event summarization

Note

Security Information and Event Management (SIEM) solutions include SIM (security information management) and SEM (security event management). A SIEM product provides real-time analysis of security alerts generated by network devices such as firewalls, IPS sensors, routers, and so forth. The network devices send their syslog, SNMP, or other format security logs to the SIEM for analysis, correlation, and reporting. SIEM solutions can be in the form of software installed on your server or appliance or event management

services. Cisco's own SIEM solution used to be the Cisco Security Monitoring, Analysis, and Response System (MARS), which is now end-of-life status. Some of the Cisco third-party partners for SIEM solutions, at the time of writing, are Splunk, ArcSoft, LogLogic, and netForensics.

Risk-Based Intrusion Prevention

Intrusion management is all about risk. Threat vectors evolve so rapidly that it becomes impossible to possess all knowledge of all threats and all threat mutations, especially when new mutations occur within hours and even minutes of the seed malware.

A common technique that is used to improve the accuracy and context awareness of the intrusion prevention architecture is illustrated in [Figure 11-4](#). This technique aims to build a risk rating into the detection capabilities, the goal of which is to detect and respond to relevant incidents only. In turn, this should reduce noise and allow IPS administrators to focus on mitigation and not on navigating the vast amounts of information.

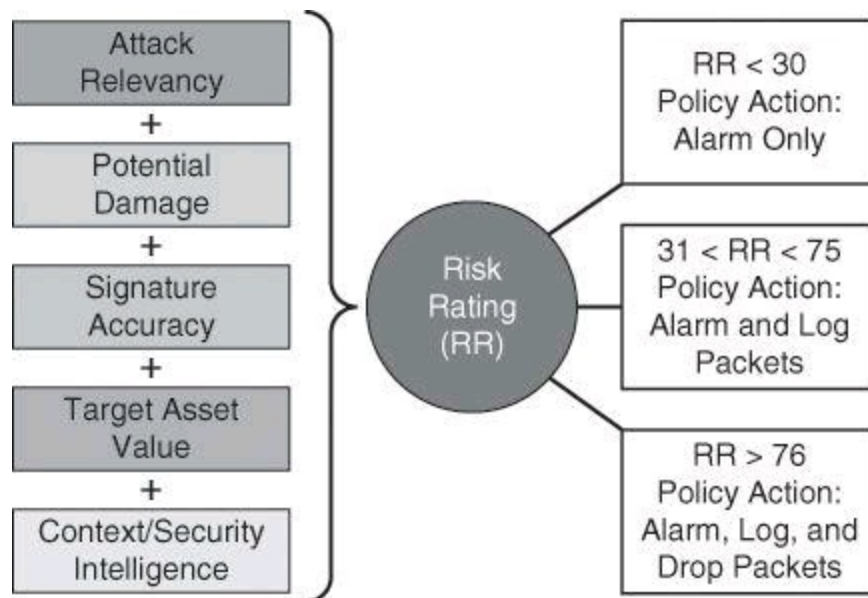


Figure 11-4. Building a Risk Rating into the Detection Capabilities

For example, an IPS sensor should detect patterns of global virus outbreaks. Using risk ratings, however, it would trigger a high-severity alert if the virus exploits the operating systems that are used in the specific network being protected, but perhaps a medium severity alert if the operating system is not in use in the organization. The response would be different for unpatched or more vulnerable servers running the relevant operating system, as opposed to fully patched servers with endpoint protection. Within these subsets, the response would be different for unpatched servers have a high asset value versus unpatched servers that store public or nonvaluable information and have a lower asset value.

Using these considerations, risk ratings typically include several components:

- **Potential damage that could be caused by the activity described by the signature:** Attacks that are potentially more devastating should result in higher risk ratings, and vice versa.
- **Asset value of the target of the attack:** Attacks against higher-valued assets should

result in higher risk ratings, and vice versa.

- **Accuracy of the triggering signature:** Inaccurate signatures lower confidence in correct sensor decisions. Therefore, the triggering of an inaccurate signature should result in lower risk ratings, and vice versa. Additionally, sensors running in inline mode are often more accurate than sensors running in promiscuous mode, due to their traffic normalization functions. Signatures triggering in inline mode should have a correspondingly higher risk rating.
- **Relevancy of the attack to the target:** Attacks that cannot cause damage against their target (for example, a Windows-specific attack targeting a Linux host) should be deprioritized and assigned a lower risk rating value, because they cause no or little damage.
- **Other security countermeasures (controls) in the environment:** Often, these may have some additional clues that can increase the situational awareness of the sensor, and allow it to react to the current context of network events and threats. External sources of information such as honeypots and early warning systems also fit into this category. This security intelligence may also influence risk rating by raising the risk rating for events that are caused by known threats that may be experiencing global or at least significant outbreaks outside of the protected network.

Threat Rating

Cisco IPS provides a threat rating (TR) when it detects an event. The TR is the risk rating (RR) minus values associated with protective responses, called TR adjustments. TR adjustments are actions such as denying an attacker or denying the connection. So, $TR = RR - TR \text{ adjustments}$. For example, if a risk is rated at 80 and the protection measure (such as blocking the attacker) is rated at 45, the TR is 35.

Further discussion on threat ratings is beyond the scope of this book; however, you will find plenty of information on this topic in *CCNP Security IPS 642-627 Official Cert Guide* (Cisco Press, 2011).

IPv6-Aware IPS

IPv6 awareness is another important consideration for IPS architectures. Sensors should be IPv6 aware; they should showcase detection techniques that consider the various threats that are found both in IPv6-only environments and in dual-stack environments where IPv4 and IPv6 coexist. As an example, IPS signatures can detect Teredo tunnels and other IPv4-to-IPv6 transition techniques, often used to exploit networks implementing dual stacks.

Alarms

Alarms fire when specific parameters are met. You must balance the number of incorrect alarms that you can tolerate with the capability of the signature to detect actual intrusions. If you have too few alarms, you might be letting in more suspect packets, but network traffic will flow more quickly. If IPSs use untuned signatures, they produce many false positive alarms, which could result in masking important alarms because they would get lost among the massive number of alarms received. You should consider the following factors when implementing alarms that a signature uses:

- The level assigned to the signature determines the alarm severity level.
- A Cisco IPS signature is assigned one of four severity levels:
 - **Informational:** Activity that triggers the signature is not considered an immediate threat, but the information provided is useful information.
 - **Low:** Abnormal network activity is detected that could be perceived as malicious, but an immediate threat is not likely.
 - **Medium:** Abnormal network activity is detected that could be perceived as malicious, and an immediate threat is likely.
 - **High:** Attacks used to gain access or cause a DoS attack are detected, and an immediate threat is extremely likely.
- You can manually adjust the severity level that an alarm produces.
- To minimize false positives, study your existing network traffic patterns and then tune your signatures to recognize intrusion patterns that are atypical (out of character) for your network traffic patterns. Do not base your signature tuning on traffic patterns that are based only on industry examples. Use industry examples as a starting point, determine what your own network traffic patterns are, and then use the results in your signature alarm tuning efforts.
- As an additional source of information, consider implementing NetFlow on network access devices such as routers and firewalls. NetFlow will collect IP traffic statistics and samples, which you can later export as NetFlow records to a NetFlow collector, such as SolarWinds or one of many others. These NetFlow statistics would assist you in characterizing the type of traffic going through your network.

IPS Alarms: Event Monitoring and Management

Event monitoring and management can be divided into the following two needs:

- Real-time event monitoring and management
- Analysis based on archived information (reporting)

These functions can be handled by a single server, or they can be placed on separate servers to scale the deployment. The number of sensors that should forward alarms to a single IPS management console is a function of the aggregate number of alarms per second generated by those sensors.

There is an important difference between reporting and monitoring. Note that archives are often a significant source of data when producing reports.

- **Reporting:** Analysis based on archived information
- **Event monitoring:** Real-time monitoring

Key
Topic

Archiving is an important aspect for proper forensics and compliance. Ensure that your archives are properly safeguarded from potential tampering and accidental loss. Your organization should have an archiving policy, often dictated by governance of the longevity of the archive (how far back

you must keep your archives). Your policies also should stipulate whether archives can be kept on site or, as usually recommended, off site.

Experience with customer networks has shown that the number of sensors reporting to a single IPS management console should be limited to 25 or fewer. These customers use a mixture of default signature profiles and tuned signatures. The number of alarms generated by each sensor is determined by how sensitively the sensor is tuned; the more sensitive the tuning, the fewer the alarms that are generated and the larger the number of sensors that can report to a single IPS management console.

Note

With the evolution of technology, the limit of 25 sensors reporting to a single IPS management console is constantly being pushed. Check with your vendor for the latest information.

Multiple protocols are available for alarm generation. Common protocols are SNMP, syslog, and Security Device Event Exchange (SDEE), the latter of which is more customized for security monitoring. SDEE is a standard developed to communicate an event generated by security devices.

When implementing multiple IPS management consoles, implement either separate monitoring domains or a hierarchical monitoring structure.

It is essential to tune out false positives to maximize the scalability of the network IPS deployment. Sensors that are expected to generate a large number of alarms, such as those sitting outside the corporate firewall, should log in to a separate IPS management console, because the number of false alarms raised dramatically increases the noise-to-signal ratio and makes it difficult to identify otherwise valid events.

Recall the following:

- False positives happen when the IDS/IPS mistakes legitimate traffic for an attack.
- False negatives happen when the IDS/IPS sensor misses an attack.



Global Correlation

[Figure 11-5](#) illustrates a distributed approach to gathering and analyzing security intelligence and threat information. By combining local, enterprise, and global intelligence, and correlating the information at all levels, you can reduce noise and focus on relevant incidents only.

- Device intelligence is gathered on a per-device basis, and when viewed locally it will be useful for pinpointing the problem to the place in the network where the sensor is located. For instance, a port-scanning reconnaissance attack on the demilitarized zone (DMZ) could be an indication of a full-blown incident, but it could also be a minor event.
- Enterprise intelligence is gathered by all sensors of the organization and typically is centralized in a SIEM system. It provides additional information to categorize events, based on the scale and scope of the event. For instance, a port-scanning attack on the DMZ

network acquires more relevance when it is followed within minutes by a firewall policy violation and a password-guessing incident on a mission-critical server.

- Global intelligence is usually gathered by early warning systems, honeypots, and managed threat services outside the organization. They give incidents a global scope. For instance, if the attack pattern detected by the enterprise SIEM is known to come from the same IP sources at a global scale, this information should trigger a more timely response. The port scanning suddenly becomes a critical incident when local, enterprise, and global information is correlated.

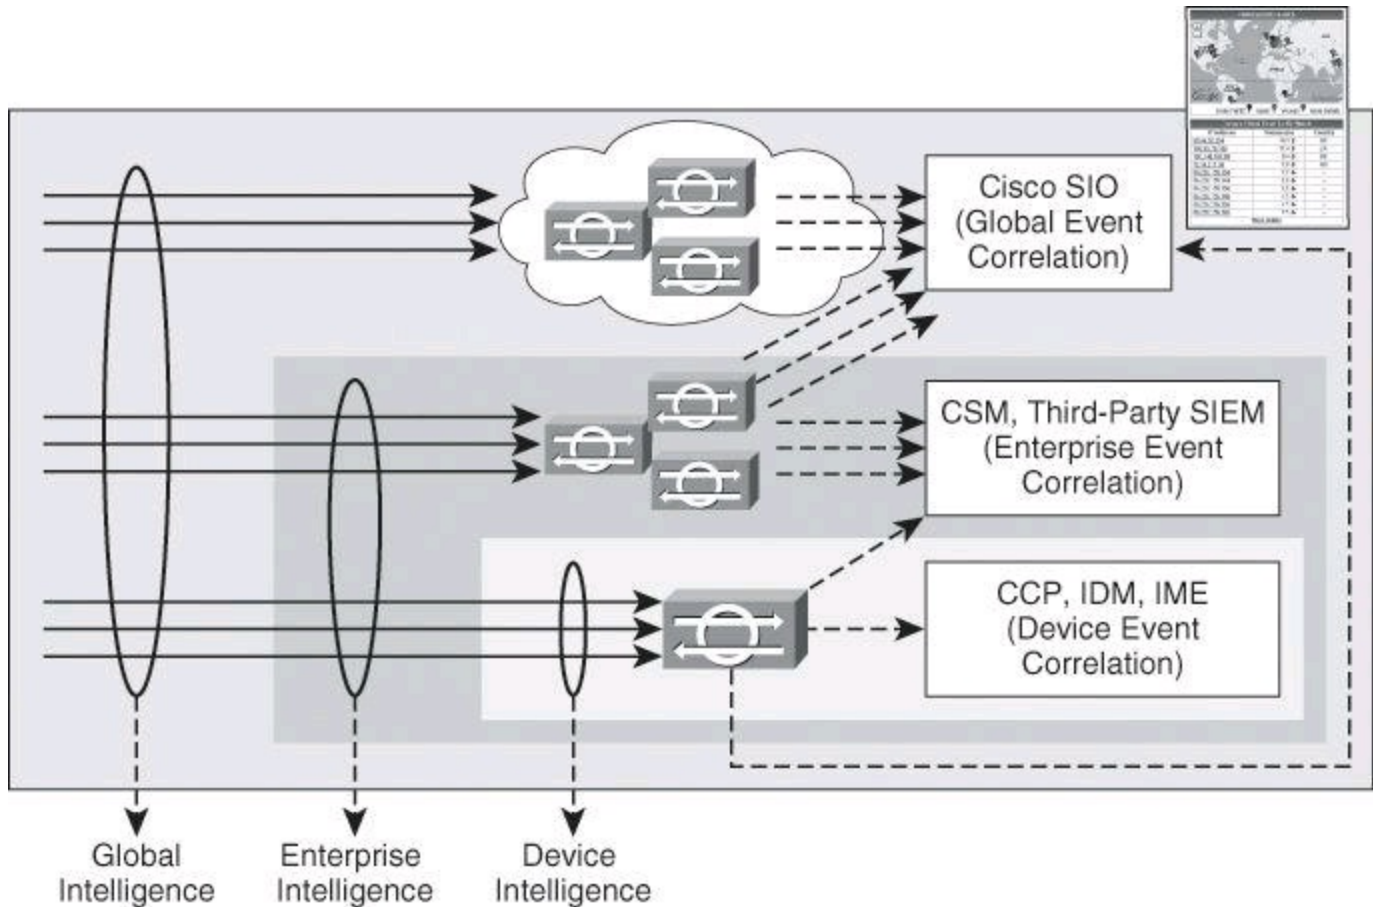


Figure 11-5. Device, Enterprise, and Global Correlation

In [Figure 11-5](#), you will also notice the event management options using Cisco IPS architectures at local, enterprise, and global levels:

- **Device event correlation:** Device-level correlation is typically accomplished with products such as Cisco Configuration Professional, IPS Device Manager, or IPS Manager Express.
- **Enterprise event correlation:** Enterprise correlation can be gathered using Cisco Security Manager or third-party, ecosystem partner SIEM systems for more specialized event management and forensics capabilities, thus providing enterprise wide intelligence. Ecosystem partner products include ArcSight, netForensics, Splunk, and others. Refer to Cisco.com for more information; specifically, look for *Cisco Security Information Event Management Deployment Guide*.
- **Global event correlation:** Global security intelligence correlation is the responsibility of the Cisco Security Intelligence Operations (SIO) cloud-based service that connects global threat information, reputation-based services, and sophisticated analysis to Cisco network

security devices to provide stronger protection with faster response times.

Why Did Cisco Pay \$830M for IronPort?

Why did Cisco pay \$830M to acquire IronPort in 2007? IronPort was selling great email and web security appliances, but these devices would not have justified the high price tag. Cisco bought IronPort for its crown jewel, SenderBase, since renamed Cisco IronPort SensorBase. SensorBase enables your security network devices, such as firewalls, sensors, and so forth, to not only build a risk profile on IP addresses, therefore allowing risk profiles to be dynamically created on HTTP sites and Simple Mail Transfer Protocol (SMTP) email sources, but also keep a list of IP addresses known for generating malicious activities. The information is collected and analyzed globally by Cisco SIO, which is then responsible for disseminating that information to all devices subscribing to SensorBase. You can check out Cisco SenderBase for yourself at <http://www.senderbase.org>. Under Reputation Lookup, plug in your domain name or www.ciscopress.com. The score is neutral and you want some excitement? Type ihaveabadreputation.com. This domain belongs to IronPort and is artificially maintained with a score of Poor to be used as a target by organizations that need to test the efficacy of their Web-Based Reputation Score filtering systems. By the way, at that website you will also find the EICAR test file, used to monitor the efficacy of virus protection.

Would you like to check the score of real domains that have been known to host malware? Visit <http://www.malwaredomainlist.com> for a list of malware domains. Pick an offending IP address from the list, go back to www.senderbase.org, and check the reputation of that offending domain and its history. Fascinating, isn't it?

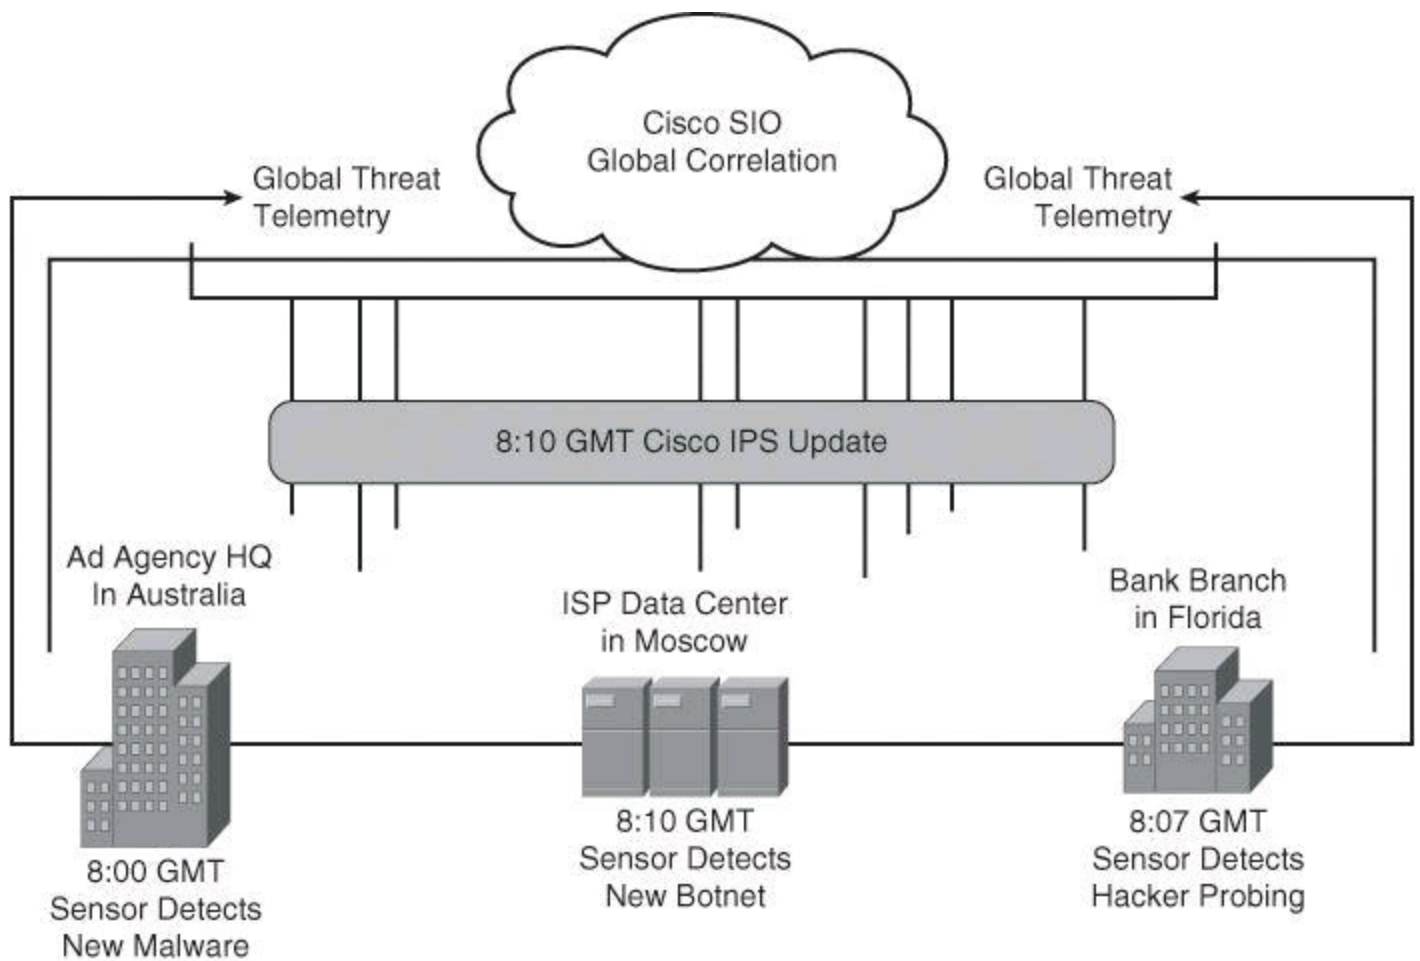
More information can be found on the power of SensorBase and Cisco Global Correlation, considered the best protection against zero-day attacks, at www.cisco.com/go/sio.

IPS Deployment

Different networks require different design and deployment options. Even within the same administrative domain, different network blocks will benefit from different IPS technologies, deployment options, and form factors. You should plan the IPS architecture carefully, using a “places in the network” approach and considering the different levels of risk for the different assets you are trying to protect.

[Figure 11-7](#) illustrates two examples:

- A remote-access SSL VPN architecture, where the potential exists that users will connect to the SSL VPN gateway router using an infected device (PC, laptop, tablet, etc.). The focus of the signature database policy is on application layer inspection, and the objective is to mitigate threats coming from unmanaged and potentially exploited endpoints.
- A private WAN scenario, where a combination of Cisco IOS IPS and integrated hardware modules on ISR G2 routers provides a smaller footprint and cost-effective architecture. The focus on the signature database policy is on DoS, combined with a reputation-based approach on some branches.



In this example, Cisco IPS with Global Correlation detects threats propagating in various countries and creates a new ruleset within 5 minutes to protect against the emerging threat.

Figure 11-6. Global Correlation and Cisco SIO at Work, Preventing Zero-Day Attack

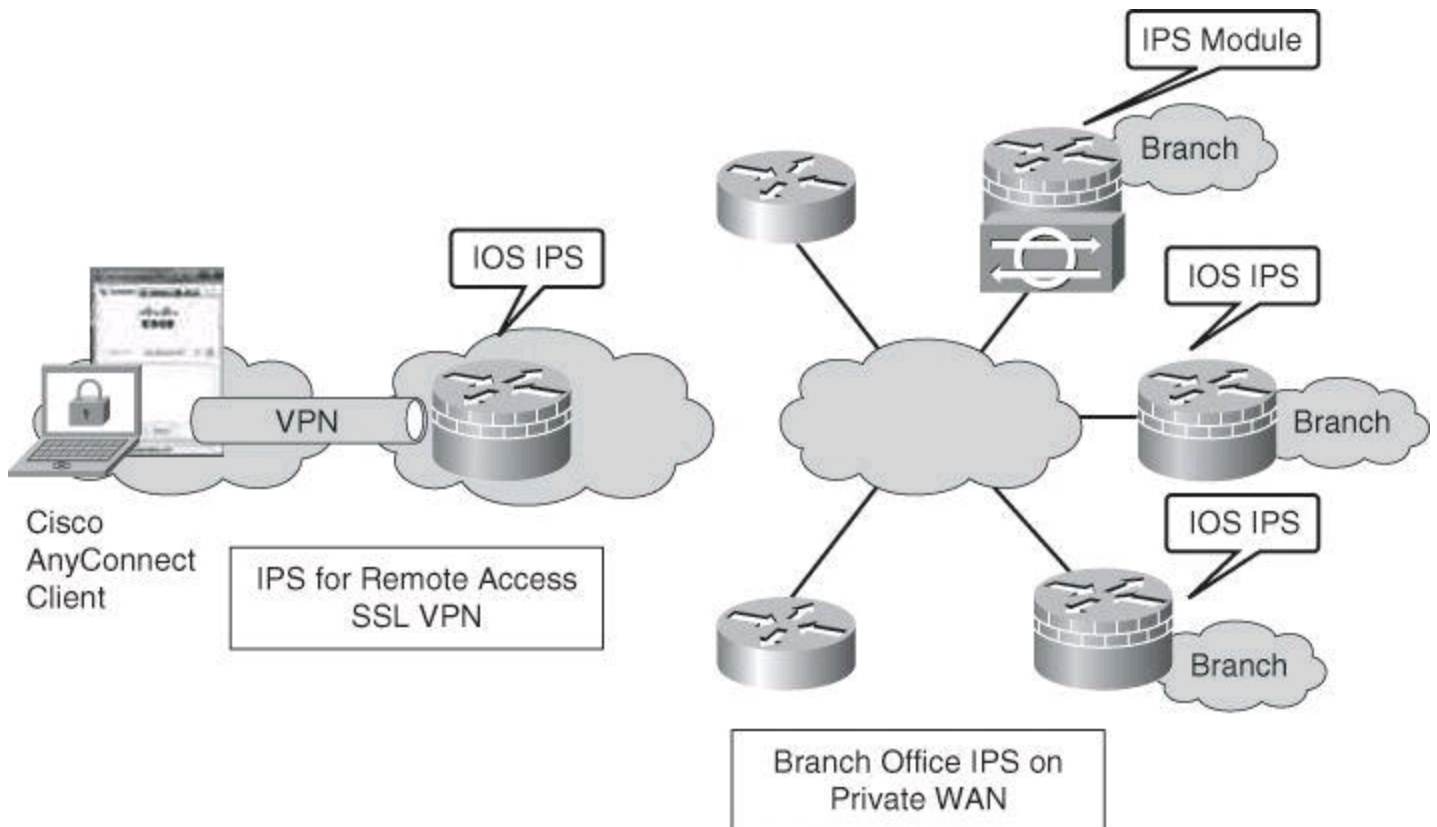


Figure 11-7. Examples of IPS Deployments

Many other places in the network can leverage the available options to deploy an effective IPS architecture.

Cisco IPS Offerings

When deploying an IPS architecture using Cisco technologies, various form factors are available to deploy intrusion prevention sensors using a distributed security intelligence approach, as shown in [Figure 11-8](#).

Cisco IPS Form Factors	Appliances	IPS4240	IPS4255	IPS4260	IPS4270
	HW Modules – Catalyst 6500			IDSM2	IDSM2 Bundle
	HW Modules – ASA	AIP-SSC-5	AIP-SSM-10	AIP-SSM-20	AIP-SSM-40
	ISR-Integrated	IOS IPS	AIM-IPS	NME-IPS	
		Small	Medium	High	
		Throughput			

Figure 11-8. IPS Platforms from Cisco

The following is a brief description of the available options and footprints:

- **Cisco ASA Advanced Inspection and Prevention Security Services Module (ASA AIP-SSM):** The Cisco ASA AIP-SSM uses advanced inspection and prevention technology to provide high-performance security services, such as intrusion prevention services and advanced content security services. The Cisco ASA AIP-SSM products include a Cisco ASA AIP-SSM-10 module with 1 GB of memory, a Cisco ASA AIP-SSM-20 module with 2 GB of memory, and a Cisco ASA AIP-SSM-40 module with 4 GB of memory.
- **Cisco IOS IPS:** Cisco IOS IPS is an inline, deep-packet inspection feature, available in Cisco Integrated Services Routers Generation 2 (ISR G2), which effectively mitigates a wide range of network attacks. A component of the Cisco IOS Integrated Threat Control framework and complemented by the Cisco IOS Flexible Packet Matching feature, Cisco IOS IPS provides your network with the intelligence to accurately identify, classify, and stop or block malicious traffic in real time.
- **Cisco IPS 4200 Series Sensors:** Cisco IPS 4200 Series Sensors offer significant protection to your network by helping to detect, classify, and stop threats, including worms, spyware and adware, network viruses, and application abuse. Using Cisco IPS Sensor Software Version 7.x, the Cisco IPS solution combines inline intrusion prevention services with innovative technologies that improve accuracy. As a result, more threats can be stopped without the risk of dropping legitimate network traffic. Cisco IPS Sensor Software

Version 7.x includes virtualization, enhanced detection capabilities, and improved scalability, resiliency, and performance features.

• **Cisco Catalyst 6500 Series Intrusion Detection System Services Module 2 (IDS-M-2):**

The Catalyst 6500 Series IDS-M-2 is part of the Cisco IPS solution. It works in combination with the other components to efficiently protect your data infrastructure. With the increased complexity of security threats, achieving efficient network intrusion security solutions is critical to maintaining a high level of protection. Vigilant protection ensures business continuity and minimizes the effect of costly intrusions.

• **Cisco IPS Advanced Integration Module (IPS AIM):** Cisco offers various IPS solutions using Cisco ISR G2 platforms. Cisco IPS Sensor Software running on the Cisco IPS AIM provides advanced, enterprise-class IPS functions and meets the ever-increasing security needs of branch offices. The Cisco IPS AIM can scale in performance to match branch office WAN bandwidth requirements today and in the future. At the same time, the integration of IPS onto a Cisco ISR keeps the solution cost low and effective for businesses of all sizes.

Note

Cisco IOS IPS and the Cisco IPS AIM cannot be used together. Cisco IOS IPS must be disabled when the IPS AIM is installed. Cisco IOS IPS is an IPS application that provides inspection capabilities for traffic flowing through the router. Although it is included in the Cisco IOS Advanced Security feature set, it uses the router CPU and shared memory pool to perform the inspection. Cisco IOS IPS also runs a subset of IPS signatures. The Cisco IPS AIM runs with a dedicated CPU and memory, offloading all processing of IPS signatures from the router CPU. It can load a full signature set and provide enhanced IPS features not available on Cisco IOS IPS.

New Cisco IPS Technologies and Sensors

Cisco introduced two additional IPS technologies in 2012:

- Cisco IPS 4300 Series Sensors, capable of 1-Gbps throughput.
- Cisco ASA 5500-X Series firewalls with IPS technology, including zero-day attack protection with anomaly detection, all built into the OS. These new firewalls are based on the SecureX technology and run Cisco ASA 8.6(1) code and Cisco IPS 7.1 code.

For more information on these new technologies and on all current Cisco IPS technologies, visit <http://www.cisco.com/go/ips> and <http://www.cisco.com/go/asa>.

IPS Best Practices

The following are the recommended practices for designing and deploying IPS architecture:

- Use a combination of detection technologies.
- Take advantage of multiple form factors to deploy a distributed and cost-effective IPS architecture.

- Use a “places in the network” approach, which, for Cisco, refers to the building blocks of a corporate network, such as a data center, a campus, and a branch office.
- Enable anti-evasion techniques.
- Take advantage of local, enterprise, and global correlation.
- Use a risk-based approach to improve accuracy and simplify management.
- When deploying a large number of sensors, automatically update signature packages instead of manually upgrading every sensor.
- Place the signature packages on a dedicated FTP server within the management network.
- Tune the IPS architecture constantly.

Note

A Great Debate: Fail-Close or Fail-Open?

This is another topic of discussion you could add to the list of IPS best practices. Fail-close or fail-open is a philosophical debate your organization needs to engage in: “If the IPS sensor stops working, do we let the traffic go through or do we stop the traffic?” The two opposing philosophies are represented in [Figure 11-9](#), where the network administrator needs to decide whether the traffic will be allowed to flow into the DMZ if the Cisco ASA AIP SSM fails.

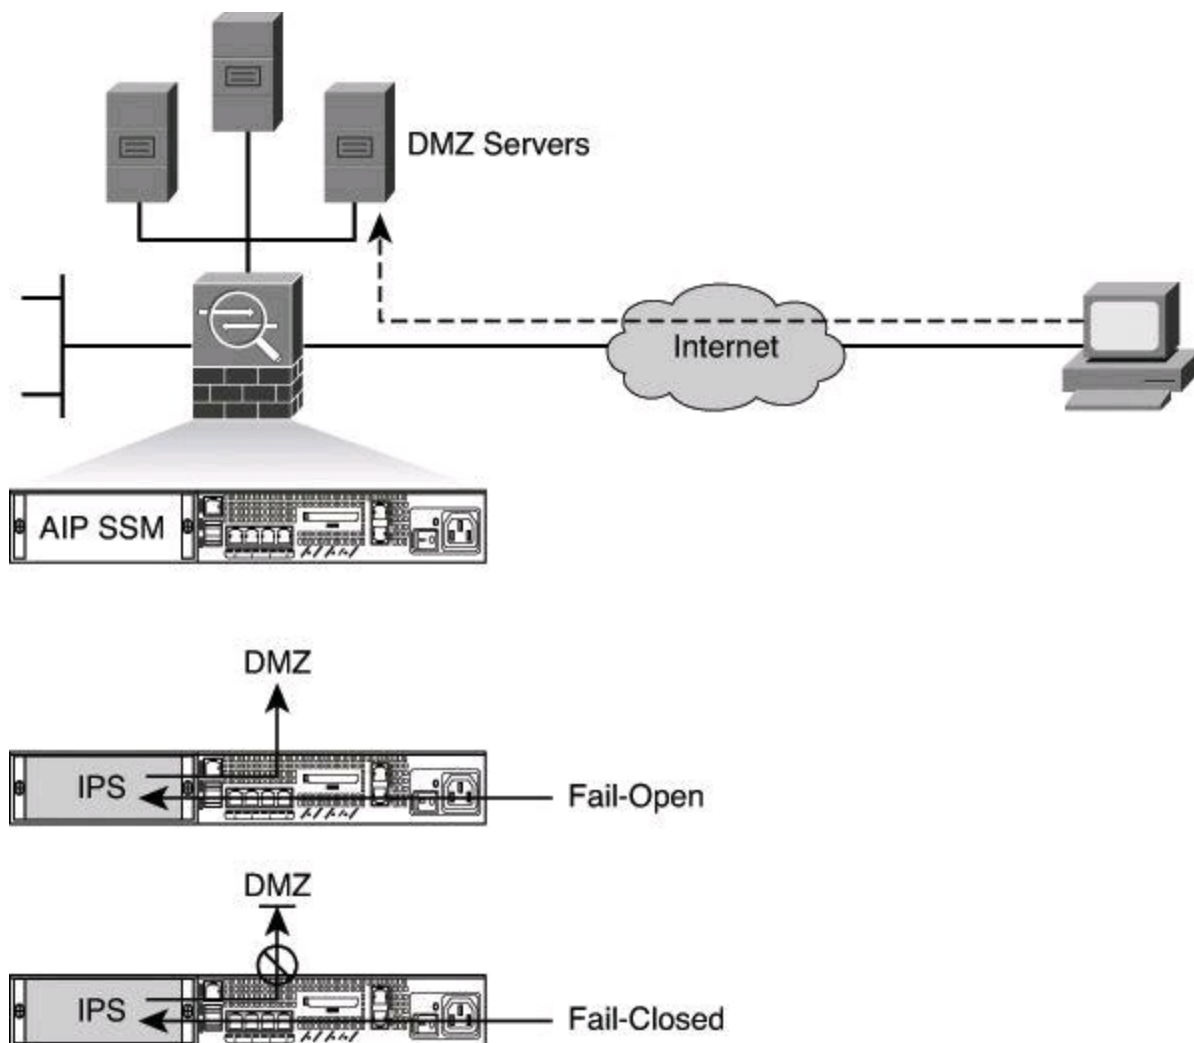


Figure 11-9. Fail-Open or Fail-Close Approach

It seems that the balance is tilting in favor of the fail-open approach with security experts, especially when dealing with for-profit organizations, but each organization has to define and enforce its own policy.

The default with Cisco IOS IPS is fail-open. You can verify this with the command **show ip ips configuration**, which should display “IPS fail closed is disabled.” The output of this command is shown later in [Example 11-2](#).

Note

Readers interested in learning more about generic topics regarding IDS/IPS should consider visiting <http://www.searchsecurity.com>, more precisely, “Security Topics” and “Tutorials.”

Recommended practices are based on a series of key factors in current IPS architectures, required to manage evolving threats and business requirements:

- Intelligent, distributed detection
 - Vulnerability- and exploit-specific signatures
 - Protocol anomaly detection
 - Knowledge base anomaly detection
 - Reputation filters
- Accurate, precise response to relevant attacks
 - Risk management–based policy
 - Global correlation adding reputation
 - On-box correlation
 - “Trustworthiness” linkages with the endpoint
- Flexible deployment options
 - Passive and/or inline with flexible response (IDS/IPS)
 - Sensor virtualization
 - Physical and logical (VLAN) interface support
 - Software and hardware bypass

Cisco IPS Architecture

Cisco IPS architectures provide an effective approach to threat control and containment by implementing IPS technologies that address the key factors mentioned in the previous section. [Figure 11-10](#) illustrates the various IPS functions that are applied to traffic as it is processed by a Cisco IPS sensor, starting with the virtual sensors. The traffic can then be analyzed by reputation filters, anti-evasion techniques, modular signature engines with automatic signature and engine updates, local and global device correlation, and risk ratings. Depending on the findings, alarm management and/or logging and forensics capabilities will be activated.

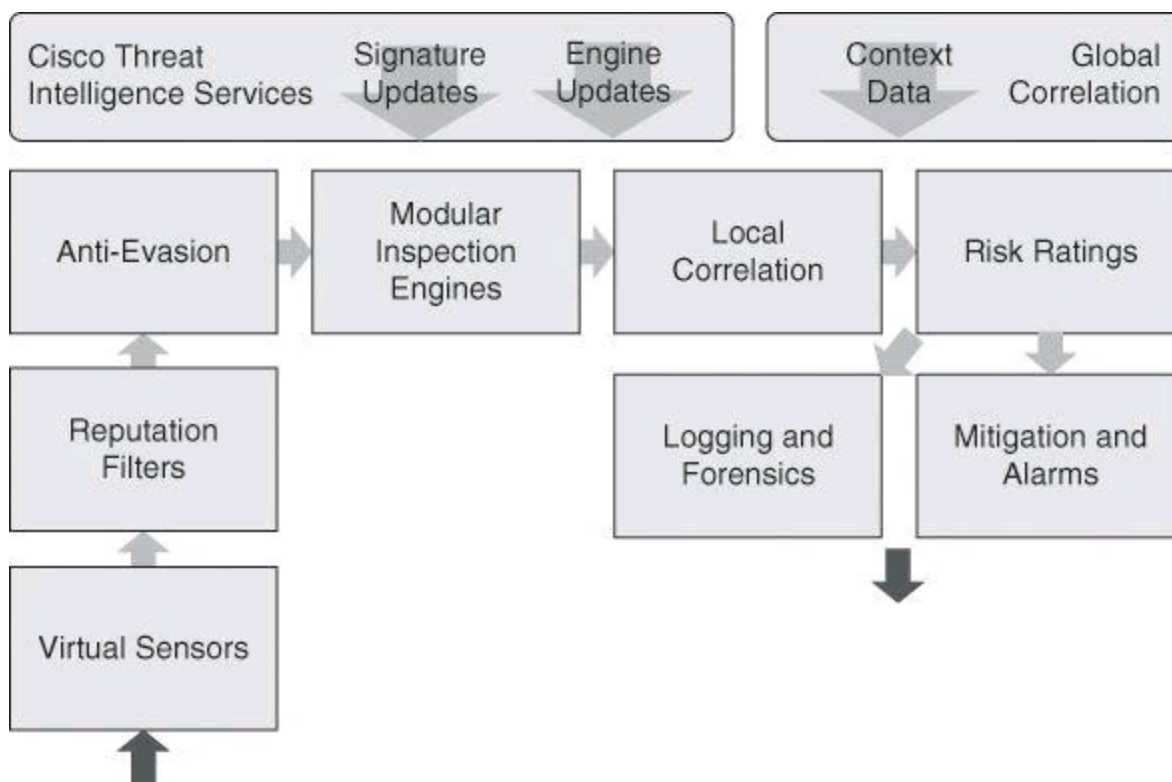


Figure 11-10. Cisco IPS Architecture

Note

Not all IPS form factors support all IPS capabilities that are shown in [Figure 11-10](#). Please refer to the proper documentation for each IPS sensor product for specific platform details. As an example, virtual sensors are used with Cisco IPS 4x00 sensor and can be thought of as a collection of data that is defined by a set of configuration policies. A discussion on virtual sensors is beyond the scope of this book.

Cisco IOS IPS

Configuring Cisco IOS Intrusion Prevention System (IPS) is a core competency for a network security administrator. This section describes how to configure Cisco IOS IPS on routers using Cisco Configuration Professional, and includes a description of the building blocks of Cisco IOS IPS, its deployment options, and guidelines for signature tuning.

Cisco IOS IPS Features

Cisco has implemented IPS functions into its Cisco IOS Software. Cisco IOS IPS uses technology from Cisco Intrusion Detection System (IDS) and IPS Sensor product lines, including Cisco IPS 4200 Series Sensors and Cisco Catalyst 6500 Series Intrusion Detection System Services Module 2 (IDSM-2). Cisco IOS IPS combines existing Cisco IDS and IPS product features with the following three intrusion detection techniques:

- Profile-based intrusion detection:** Profile-based intrusion detection generates an alarm when activity on the network goes outside a defined profile. With anomaly detection, profiles are created for each user or user group on your system. These profiles are then used as a baseline to define normal user and network activity. A profile could be created to monitor web traffic.

- **Signature-based intrusion detection:** Signature-based intrusion detection is less prone to triggering a false alarm when detecting unauthorized activity. A signature is a set of rules pertaining to typical intrusion activity. Signature-based intrusion detection uses signatures that are based on values in IP, TCP, UDP, and ICMP headers. Network engineers research known attacks and vulnerabilities and then develop signatures to detect these attacks and vulnerabilities on the network. These attack signatures encompass specific traffic or activity that is based on known intrusive activity.

Cisco IOS IPS implements signatures that can look at every packet going through the network and generate alarms when necessary. Cisco IOS IPS generates alarms when a specific pattern of traffic is matched or a signature is triggered. You can configure Cisco IOS IPS to exclude signatures and modify signature parameters to work optimally in your network environment.

A pattern-matching approach searches for a fixed sequence of bytes in a single packet. Pattern matching is a rigid approach but is simple to employ. In most cases, the pattern is matched against a packet only if the suspect packet is associated with a particular service or, more precisely, destined to or from a particular port. For example, a signature might be based on a simple pattern-matching approach such as the following: If <the packet is IPv4 and TCP> and <the destination port is 2222> and <the payload contains the string “foo”> then <fire an alarm>.

- **Protocol analysis–based intrusion detection:** Protocol analysis–based intrusion detection is similar to signature-based intrusion detection but performs a more in-depth analysis of the protocols specified in the packets. A deeper analysis examines the payloads within TCP and UDP packets, which contain other protocols. For example, a protocol such as DNS is contained within TCP or UDP, which itself is contained within IP.

The first step of protocol analysis is to decode the packet IP header information and determine whether the payload contains TCP, UDP, or another protocol. For example, if the payload is TCP, some of the TCP header information within the IP payload is processed before the TCP payload is accessed

Protocol analysis requires the IPS sensor to know how various protocols work so that it can more closely analyze the traffic of those protocols to look for suspicious or abnormal activity. For each protocol, the analysis is based not only on protocol standards, particularly the RFCs, but also on how things are implemented in the real world. Many implementations violate protocol standards. It is important that signatures reflect common and accepted practice rather than the RFC-specified ideal; otherwise, false results can be reported.

The following are key points of Cisco IOS IPS:

A circular icon with a dotted border containing the text "Key Topic".

Key
Topic

- Software-based inline intrusion prevention sensor
- Supports same software and signature format as other Cisco IPS products starting with 12.4(11)T

- Supports signature-based scanning, but uses a blend of detection technologies:
 - Profile based
 - Signature based
 - Protocol analysis based

The following attributes describe the primary benefits of the Cisco IOS IPS solution:

- Cisco IOS IPS uses the underlying routing infrastructure to provide an additional layer of security with investment protection.
- Because Cisco IOS IPS is inline and is supported on a broad range of routing platforms, attacks can be effectively mitigated to deny malicious traffic from both inside and outside the network.
- When used in combination with Cisco IDS, Cisco IOS Firewall, virtual private network (VPN), and Network Admission Control (NAC) solutions, Cisco IOS IPS provides superior threat protection at all entry points to the network.
- Cisco IOS IPS is supported by easy and effective management tools, such as Cisco CCP, IDM, IME, and Cisco Security Manager.
- Whether threats are targeted at endpoints, servers, or the network infrastructure, Cisco IPS offers pervasive intrusion prevention solutions that are designed to integrate smoothly into the network infrastructure and to proactively protect vital resources.
- Cisco IOS IPS supports attack signatures from the same signature database that is available for Cisco IPS appliances.

Scenario: Protecting the Branch Office Against Inside Attack

The Cisco IOS IPS solution enables a distributed approach to threat mitigation by allowing intrusion prevention to be enabled on a software footprint on existing routers. [Figure 11-11](#) depicts a typical branch protection scenario, where VPN routers connect to corporate assets across a public network. Cisco IOS IPS can be used to implement a focused intrusion prevention strategy, mitigating malware outbreaks that could originate in less protected remote locations, or partner locations that fall outside of the administration domain of corporate security teams. The smaller footprint takes full advantage of the hardware and architectural resources of ISR G2 routers and allows for inbound IPS functionality to protect the branch itself from targeted external attacks.

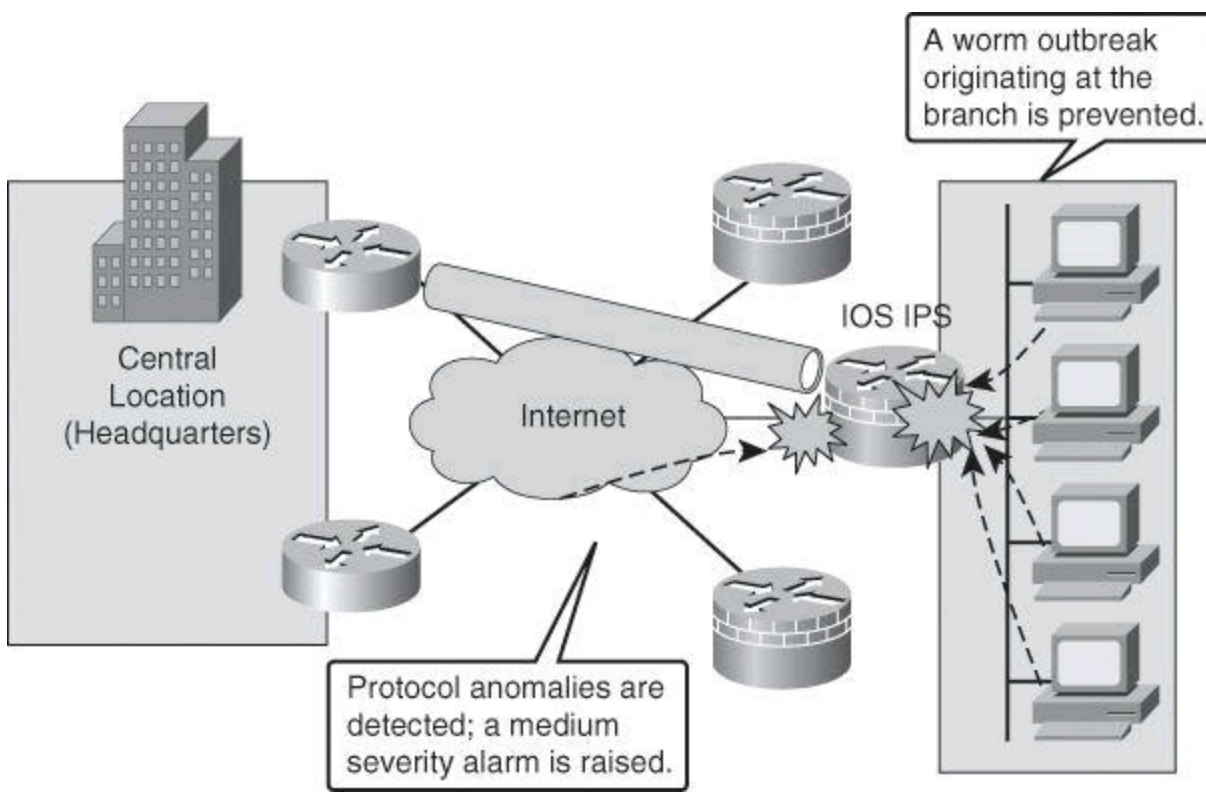


Figure 11-11. Cisco IOS IPS Protecting Against an Inside Attack

Signatures

Many of the features that make the Cisco IPS architecture effective are inherited by Cisco IOS IPS. [Table 11-5](#) describes the features of Cisco IOS IPS–based signatures.

Table 11-5. Cisco IOS IPS Signature Features

Cisco IOS IPS Signature Feature	Description
Regular expression string pattern matching	Enables the creation of string patterns using regular expressions.
Response actions	Enables the sensor to take an action when the signature is triggered.
Alarm summarization	Enables the sensor to aggregate alarms. It does this to limit the number of times an alarm is sent when the signature is triggered.
Threshold configuration	Enables a signature to be tuned to perform optimally in a network.
Anti-evasive techniques	Enables a signature to defeat evasive techniques used by an attacker.
Risk ratings	Enables risk-based tuning using Event Risk Rating (ERR), which is based on attack severity, signature fidelity, and asset value.

Signature Files

As explained earlier in this chapter, a signature is a set of rules that an IDS and an IPS use to detect

typical intrusive activity, such as DoS attacks.

IPS signatures are dynamically updated and posted to Cisco.com on a regular basis so that customers can access signatures that help protect their network from the latest known network attacks. The updates take the form of signature files, also known as signature packages or simply signature updates. Manual updates are also possible.

Cisco IPS devices use signature engines to load the signature files and scan signatures. Signatures that are contained within the signature packages are managed by various signature engines. The packages typically contain signature definitions for multiple engines.

Signature engines typically correspond to the protocol in which the signature occurs and look for malicious activity in that protocol. Each signature engine provides a common set of signature parameters that can be used to tune the sensitivity, scope, and actions of that particular signature engine, instead of making changes to individual signatures.

Signature Engines = Types of Interrogation

One way to think of signature engines is as types of interrogation. Each engine is an individual process that specializes in a particular type of interrogation. For example, one engine might specialize in IP headers, while another engine specializes in TCP headers, another in UDP headers, and yet another in ICMP traffic. Other engines look at application layer traffic such as SMTP or DNS. Data is processed in parallel by all the appropriate engines. Each engine specifies a set of available parameters to match malware traffic patterns.

A signature package has definitions for each signature it contains. After signatures are loaded and compiled onto a router running Cisco IOS IPS, IPS can begin detecting the new signatures immediately. Routers access signature definition information through a directory in flash that contains three configuration files—the default configuration, the delta configuration, and the Signature Event Action Processor (SEAP) configuration. SEAP is the control unit responsible for coordinating the data flow of a signature event.

Within a signatures package, signatures are categorized using signature engines and top-level categories. The first criterion to categorize signatures uses the basic and advanced signature categories. The basic signature category is appropriate for routers with less than 128 MB of available flash memory. The advanced signature category is appropriate for routers with more than 128 MB of available flash memory. Signatures are organized into signature microengines, explained further later in this section.

IOS IPS Default Category

Starting with Cisco IOS Release 15.0(1)M, in addition to the basic and advanced signature categories, there is a new category called IOS IPS default. This default category of signatures (including some lightweight signatures) is updated frequently by the Cisco signature team.

Signature files also come in different formats for the different methods to load them into the

router's flash. The file naming convention indicates which method to use. The following file naming conventions are available:

- Cisco Configuration Professional signature updates use the sigvX-SDM-S####.zip name format, where *X* is the IPS architecture version, and #### is the update ID, a number that increases with every new signature package.
- You can copy files from the CLI, in which case the naming conventions change: IOS-S####-CLI.pkg applies to CLI updates, where #### is the update ID.

Unified Signatures

The signatures became unified across all Cisco platforms with version 5.x of Cisco IOS IPS. Also, you will notice in the bullets regarding the signature updates the SDM reference in the naming convention. This is a legacy from the days of Cisco Security Device Manager, the predecessor to CCP.

Signature Management

Multiple enhanced features are available in Cisco IOS IPS to simplify the signature management process. The features include the following:

- **Encrypted signature support:** Digitally signed signatures allow for trusted authentication and nonrepudiation of the source.
- **Lightweight signatures:** Lightweight signatures allow loading of larger signatures sets, without consuming significant additional memory or reducing the memory that is consumed by an existing signature set.
- **Direct download from Cisco.com capability:** This feature allows an administrator to use the CLI to specify, download, and upgrade to new signatures posted for Cisco IOS routers directly from Cisco.com. An administrator can also configure the router through the CLI to receive future periodic signature downloads automatically to eliminate the manual maintenance efforts and costs of changing or tuning IPS signatures whenever a new update is posted.
- **Tuning per top-level categories:** Top-level signature categories classify signatures for easy grouping and tuning. Group-wide parameters, such as signature event actions, can be applied to a group instead of per signature.
- **Signature tuning inheritance:** This feature allows a network administrator to preserve the customization done on a signature when a signature package is updated. A local file includes locally tuned signature parameters, taking precedence over globally administered signature updates that may reset those parameters. Network administrators can either preserve the current configuration of signature actions or override the configuration with the default settings in future signature updates.

Examining Signature Microengines

A signature microengine (SME) is a component of an IDS and IPS sensor that supports a group of signatures that are in a common category. Each SME is customized for the protocol and fields that it is designed to inspect and defines a set of legal parameters that have allowable ranges or sets of

values. The SMEs look for malicious activity in a specific protocol. Signatures can be defined for any of the supported SMEs by using the parameters offered by the supporting SME. Packets are scanned by the SMEs that understand the protocols contained in the packet.

Cisco SMEs implement parallel scanning. All the signatures in a given SME are scanned in parallel fashion, rather than serially. Each SME extracts values from the packet and passes portions of the packet to the regular expression engine. The regular expression engine can search for multiple patterns at the same time (in parallel). Parallel scanning increases efficiency and results in higher throughput.

When IDS (promiscuous mode) or IPS (inline mode) is enabled, an SME is loaded (or built) on to the router. When an SME is built, the router may need to compile the regular expression found in a signature. Compiling a regular expression requires more memory than the final storage of the regular expression. Be sure to determine the final memory requirements of the finished signature before loading and merging signatures.

Note

A regular expression is a systematic way to specify a search for a pattern in a series of bytes. As an example, a regular expression to be used to prevent data containing .exe or .com or .bat from crossing the firewall could look like this:

```
.*\.[Ee][Xx][Ee]|\.*\.[Cc][Oo][Mm]|\.*\.[Bb][Aa][Tt]
```

Note

For the list of currently supported SMEs, refer to Cisco.com. Specifically, see “Signature Microengines Overview and Lists of Supported Engines” at http://www.cisco.com/en/US/docs/ios-xml/ios/sec_data_ios_ips/configuration/15-0m/sec-cfg-ips.html#GUID-E402D400-0222-4206-8B2D-3B2368C7EABB

[Table 11-6](#) summarizes the types of signature engines available in Cisco IOS IPS.

Table 11-6. Summary of Types of Supported Signature Engines

Signature Engine	Description
Atomic	Signatures that examine single packets, for ICMP, TCP, and UDP. The atomic engine does not store persistent data across packets. Each packet is evaluated independently of preceding or following packets.
Service	Signatures that examine the many services that are attacked, such as the service HTTP engine.
String	Signatures that use regular expression-based patterns to detect intrusions. Used with TCP, UDP, and ICMP. As an example, it could be looking for destination port 23 in the TCP header.
Multi-string	Supports flexible pattern matching and supports Trend Labs signatures. Multi-string signatures inspect Layer 4 transport protocol (ICMP, TCP, and UDP) payloads using multiple string matches for one signature.
Other	Internal engine to handle miscellaneous signatures.

[Table 11-7](#) provides more details on signature engines.

Table 11-7. Details on Signature Microengines

Signature Microengine	Description
ATOMIC.IP	Provides simple Layer 3 IP alarms
ATOMIC.ICMP	Provides simple ICMP alarms based on these parameters: type, code, sequence, and ID
ATOMIC.IPOPTIONS	Provides simple alarms based on the decoding of Layer 3 options
ATOMIC.UDP	Provides simple UDP packet alarms based on these parameters: port, direction, and data length
ATOMIC.TCP	Provides simple TCP packet alarms based on these parameters: port, destination, and flags
SERVICE.DNS	Analyzes the DNS service
SERVICE.RPC	Analyzes the remote procedure call (RPC) service
SERVICE.SMTP	Inspects SMTP
SERVICE.HTTP	Provides HTTP protocol decode-based string engine; includes anti-evasive URL deobfuscation
SERVICE.FTP	Provides FTP service special decode alarms
STRING.TCP	Offers TCP regular expression–based pattern inspection engine services
STRING.UDP	Offers UDP regular expression–based pattern inspection engine services
STRING.ICMP	Provides ICMP regular expression–based pattern inspection engine services
MULTI-STRING	Supports flexible pattern matching and supports Trend Labs signatures
Other	Provides internal engine to handle miscellaneous signatures

Signature Tuning

Router memory and resource constraints prevent a router from staging all Cisco IOS IPS signatures in such a way that they start utilizing router resources. When signature packages are copied into flash, the signatures are not yet staged in RAM. When Cisco IOS IPS is initialized, the signatures are loaded into the signature database. *Loading* refers to the process where Cisco IOS IPS parses the signature files and populates the signature database. However, it is the process of compiling, as shown in [Figure 11-12](#), that stages the signatures and makes them utilize router resources. *Compiling* refers to the process where the parameter values from signatures are compiled into a regular expression table.

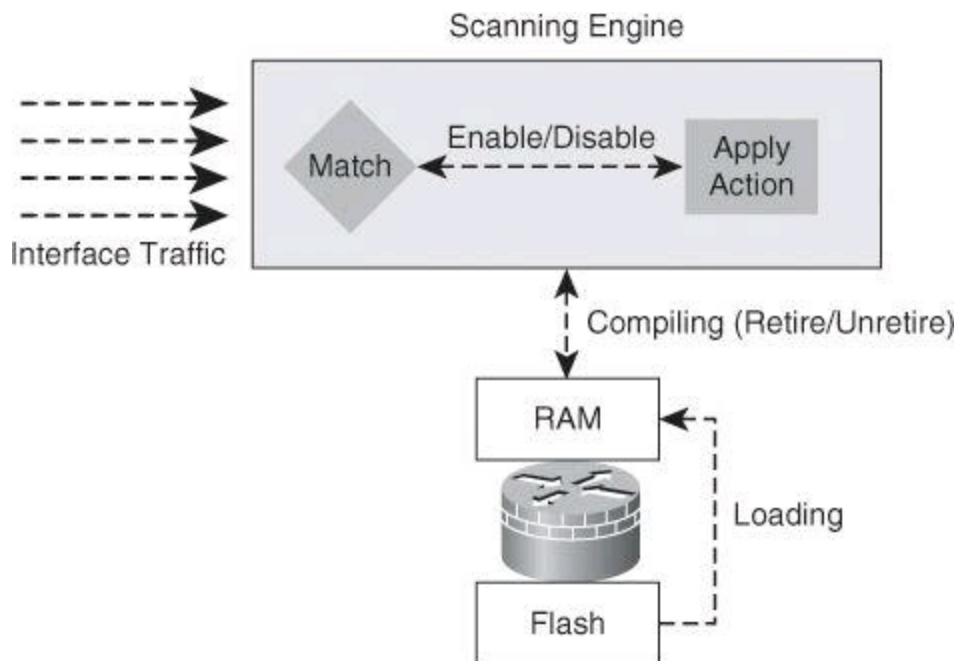


Figure 11-12. Signatures Interactions with Cisco IOS

When signatures are compiled, they are not scanned on a serial basis. A parallel signature scanning engine is used to scan for multiple patterns within a signature engine at any given time, and match intrusion activity for traffic flows that are processed by the router's interfaces. If traffic matches those patterns, an action is performed.

You can alter this process to preserve router resources and tune your intrusion prevention system. Cisco IOS IPS allows administrators to retire signatures, which prevents them from being compiled and using router resources. Similarly, unretiring signatures would compile them again. Retiring preserves the router's resources.

The system also allows administrators to enable and disable signatures. An enabled signature performs the actions that are associated with the signature, while a disabled signature does not.

[Table 11-8](#) provides a description of the different states of signatures.

Table 11-8. Signature States

Signature State	Description
Retired	Cisco IOS IPS will not compile that signature into memory for scanning.
Unretired	Cisco IOS IPS will compile the signature into memory and use the signature to scan traffic.
Enabled	When triggered by a matching packet (or packet flow), the signature takes the action associated with it.
Disabled	When triggered by a matching packet (or packet flow), the signature does not take the action associated with it.

By default, most signatures are retired to preserve router resources. The first level of tuning in Cisco IOS IPS environments is aimed at striking a balance between performance and IPS effectiveness. Relevant signatures should be unretired and enabled to increase the effectiveness of the system. When doing this, memory and computing resources should be closely monitored. There is no magic number or rule to design the appropriate configuration for your environment. Careful planning is required. You should start with the predefined basic and advanced signature definitions that are based on the hardware resources on the router, and build from there.

It is worth noting that disabling signatures does not preserve computing resources if the signature is not retired. An unretired disabled signature, similar to an unretired enabled signature, is still processed against inspected traffic by the scanning engine. The only difference is that the associated action will not be performed in the case of the disabled signature.

It is also worth noting that only unretired *and* successfully compiled signatures will take the action when they are enabled. When a signature is retired, even though it may be enabled, it will not be compiled and will not take the action that is associated with it. Also, when a signature is disabled, even though it may be unretired and successfully compiled, it will not take the action that is associated with it.

The combinations of signature compilations and states are outlined in [Table 11-9](#).

Table 11-9. Combinations of Signature Compilations and States

	Enabled	Disabled
Retired	No memory consumption No action	No memory consumption No action
Unretired	Consumes memory IPS action performed	Consumes memory No action

The predefined Cisco IOS IPS basic and advanced signature categories contain an optimum combination of signatures for all standard memory configurations, providing a good starting point.

The following list summarizes the guidelines for planning an efficient and effective Cisco IOS IPS signature definition:

- The number of signatures that can be compiled depends on the free memory available on the router.
- For routers with 128 MB of flash, start with the basic signature category.
- For routers with 256 MB+ of flash, start with the advanced signature category.
- Retire risk-irrelevant signatures according to your needs.
- Monitor free memory when retiring or unretiring signatures.
- In restrictive policies, define a fail-closed action if signatures fail to compile. This setting instructs the router to drop all packets until the signature engine is built and ready to scan traffic. If this command is issued, one of the following scenarios occurs:
 - If IPS fails to load the signature package, all packets are dropped—unless the

user specifies an access control list (ACL) for packets to send to IPS.

- If IPS successfully loads the signature package, but fails to build a signature engine, all packets that are destined for that engine are dropped.
- If this command is not issued, all packets are passed without scanning if the signature engine fails to build.
- Disabled signatures are still scanned and processed, and will consume resources.
- Never unretire the “All” signature category.

Other signature tuning opportunities are outside the scope of this book. However, some guidelines and recommended practices are listed here. It is recommended that you tune signature parameters to customize the sensitivity, scope, response, and overall accuracy of your intrusion prevention system.

- Apply Cisco IOS IPS policies on interfaces for ingress traffic.
- Tune signature parameters to define the sensitivity, scope, response, and accuracy of signatures.
 - Matching criteria (patterns, protocol parameters, and so on)
 - Signature action
 - Thresholds and counters
- Add or remove actions based on risk rating.
 - Example: For an alarm-only signature, add a deny packet inline action if the risk rating is above a certain threshold.
- Create custom signatures for vulnerabilities and threats unique to your environment.
- Configure the router not to override tuned parameters in the next signature package update.

Monitoring IPS Alarms and Event Management

An important part of network security is the alarms generated by IPS. These alarms need to be collected and analyzed. In this section, we discuss the alarms and where they are sent.

Cisco IOS IPS Alarms Monitoring

[Figure 11-13](#) shows how to use the Security Device Event Exchange (SDEE) protocol and a syslog-based approach to send Cisco IPS alerts. The sensor generates an alarm when an enabled signature is triggered. Alarms are stored on the sensor. A host can pull the alarms from the sensor using SDEE. Pulling alarms from a sensor allows multiple hosts to subscribe to the event “feed” to allow a host or hosts to subscribe on an as-needed basis.




Figure 11-13. Support for SDEE and Syslog

The support for SDEE and syslog in the Cisco IOS IPS solution is as follows:

- Cisco IOS Software supports the SDEE protocol. When Cisco SDEE notification is enabled (by using the **ip ips notify sdee** command), by default 200 events can be stored in the event buffer, whose size can be increased to hold a maximum of 1000 events. When you disable Cisco SDEE notification, all stored events are lost. A new buffer is allocated when the notifications are reenabled.
- SDEE uses a pull mechanism. That is, requests come from the network management application, and the IDS and IPS router responds.
- SDEE becomes the standard format for all vendors to communicate events to a network management application.
- You must also enable HTTP or HTTPS on the router, using the **ip http server** command, when you enable SDEE. The use of HTTPS ensures that data is secured as it traverses the network.
- The Cisco IOS IPS router still sends IPS alerts via syslog.

When you use CCP, you can keep track of alarms that are common in SDEE system messages, including IPS signature alarms. The following is an example of an SDEE system alarm message:

```
%IPS-4-SIGNATURE:Sig:1107 Subsig:0 Sev:2 RFC1918 address
[192.168.121.1:137 ->192.168.121.255:137]
```

The preceding alarm was triggered by the fact that a packet with a private address, as listed in RFC 1918, traversed the IPS sensor.

Note

For a complete list of the Cisco IOS IPS system messages, refer to the “Interpreting Cisco IOS IPS System Messages” section in the *Cisco IOS Security Configuration Guide, Release 12.4*.

Event Management

Using SDEE and syslog, event monitoring can be accomplished at various levels, depending on the availability of local, enterprise, or global correlation and event management systems. In addition to the router's buffers, the following management applications can be used to store a single router's alarms, as also previously shown in [Figure 11-5](#):

- Local event management and correlation
 - Cisco Configuration Professional
 - IPS Device Manager
 - IPS Manager Express
- Enterprise event management and correlation
 - Cisco Security Manager
 - Third-party ecosystem partner SIEM systems
- Global event management and correlation
 - Cisco Security Intelligence Operations (SIO)

IPS Manager Express

IPS Manager Express (IME) is a free application from Cisco that can manage alarms from up to 10 SDEE sources—either Cisco IPS Sensors or Cisco IOS IPS routers. Many of the features in IME work only with the IPS Sensors, and not with the Cisco IOS IPS routers.

Configuring Cisco IOS IPS Using Cisco Configuration Professional

Cisco IOS IPS allows you to manage intrusion prevention on routers that use Cisco IOS Software Release 12.3(8)T4 or later. Cisco Configuration Professional lets you control the application of Cisco IOS IPS on interfaces, import and edit signature files from Cisco.com, and configure the action that Cisco IOS IPS should take if a threat is detected.

Following are the configuration steps to deploy Cisco IOS IPS using CCP:

- Step 1.** Download the latest Cisco IOS IPS signature package to a local PC using Cisco Configuration Professional Auto Update.
- Step 2.** Launch the IPS Policies Wizard to configure Cisco IOS IPS.
- Step 3.** Verify that Cisco IOS IPS configuration and signatures are properly loaded.
- Step 4.** Perform signature tuning.
- Step 5.** Verify alarms.

The configuration scenario is shown [Figure 11-14](#). Configure this branch router to enable Cisco IOS IPS on both inside and outside interfaces, enabling the advanced signature category. Then, you can perform basic tuning, unretiring signatures and changing the default action.

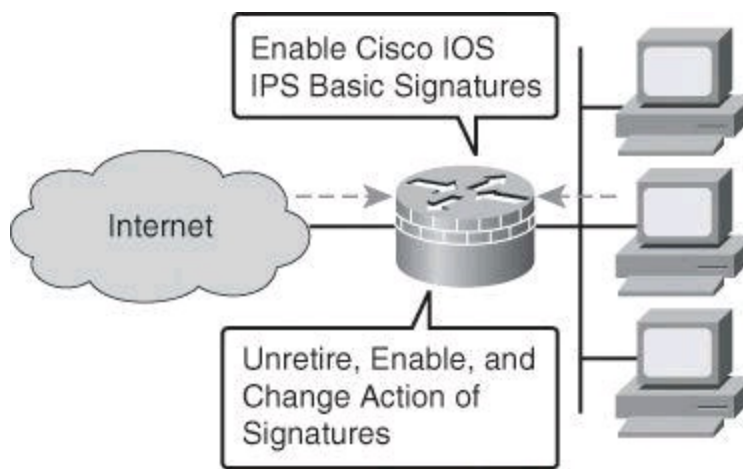


Figure 11-14. Cisco IOS IPS Configuration Scenario

Step 1: Download Cisco IOS IPS Signature Package

In the first step, you navigate to the Auto Update screen. From the Cisco Configuration Professional Home page, click **Configure**, choose **Security > Intrusion Prevention** in the left navigation pane, click the **Edit IPS** tab on the right, and then click **Download** in the left column. If SDEE notification is not enabled on the router, an Information dialog box appears, as shown in [Figure 11-15](#). Click **OK** to enable SDEE notification.

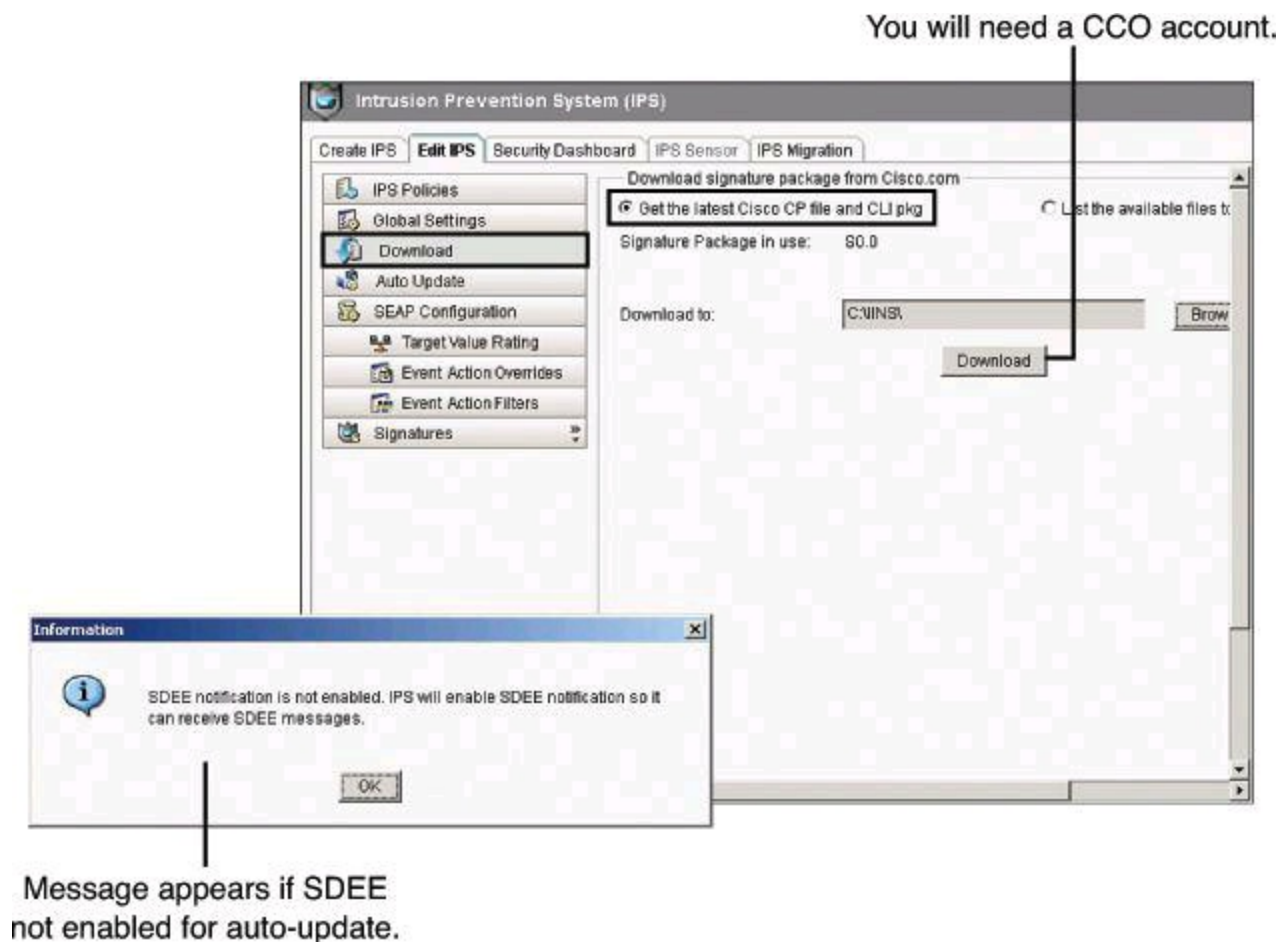


Figure 11-15. Downloading Signature Package Using CCP

Download the latest Cisco IOS IPS signature package to a local TFTP or FTP server. On the Download panel, click the **Get the Latest Cisco CP File and CLI pkg** radio button, as shown in [Figure 11-15](#). Next, click the **Browse** button and select a directory on your local PC in which to save

the downloaded files. You can choose the TFTP or FTP server root directory, which will be used later when deploying the signature package to the router. Next click **Download**, as shown in [Figure 11-15](#).

When prompted to provide Cisco.com login credentials, enter your Cisco.com registered username and password.

CCP connects to Cisco.com, as shown in [Figure 11-16](#), and starts to download both the CCP signature file (for example, sigv5-SDM-S353.zip) and the CLI signature pkg file (for example, IOS-S353-CLI.pkg) to the directory you selected in the previous step. After both files are downloaded, CCP prompts you to push the downloaded signature package to the router. Choose **No** if, as with our scenario, Cisco IOS IPS has not yet been initialized on the router.



Figure 11-16. Updating the Cisco IOS IPS Signature Package

Step 2: Launch IPS Policies Wizard

After CCP downloads the latest Cisco IOS CLI signature package, go to the Create IPS tab to create the initial Cisco IOS IPS configuration, as shown in [Figure 11-17](#) (**Configure > Security > Intrusion Prevention > Create IPS**). If prompted to apply changes to the router, click **Apply Changes**. Next, click **Launch IPS Rule Wizard** (which in fact launches the IPS Policies Wizard). An Information dialog box pops up to inform you that Cisco Configuration Professional needs to establish an SDEE subscription to the router to retrieve alerts. Click **OK**.

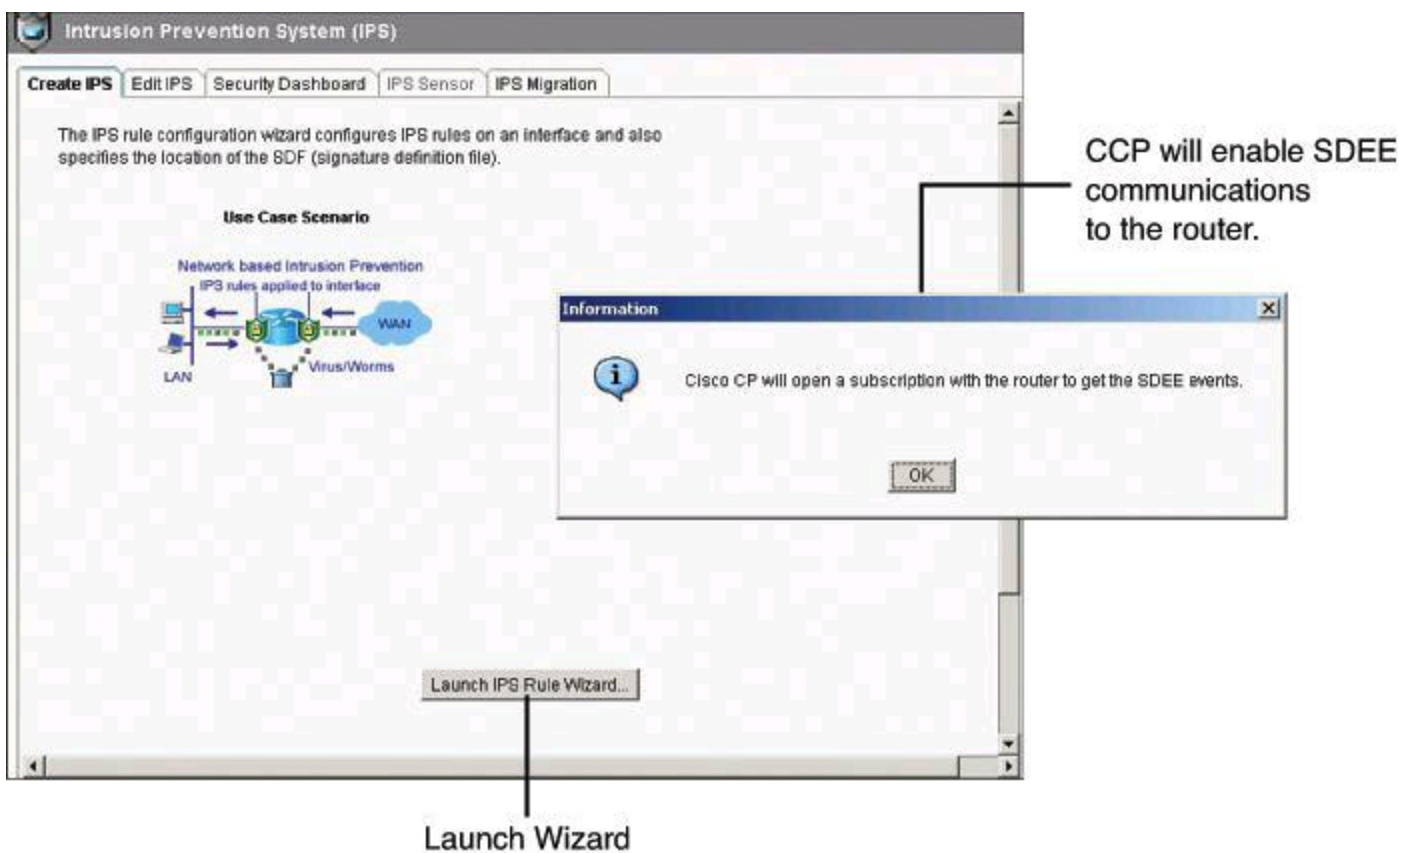


Figure 11-17. Creating an IPS Policy by Launching the IPS Policies Wizard in CCP

On the Select Interfaces screen of the IPS Policies Wizard, choose the interfaces on which you want to apply the Cisco IOS IPS rule by specifying whether the rule is to be applied to inbound traffic or outbound traffic, as shown in [Figure 11-18](#). If you check both the inbound and the outbound boxes, the rule applies to traffic flowing in both directions. Click **Next** to continue.

Select interfaces to enable intrusion prevention.

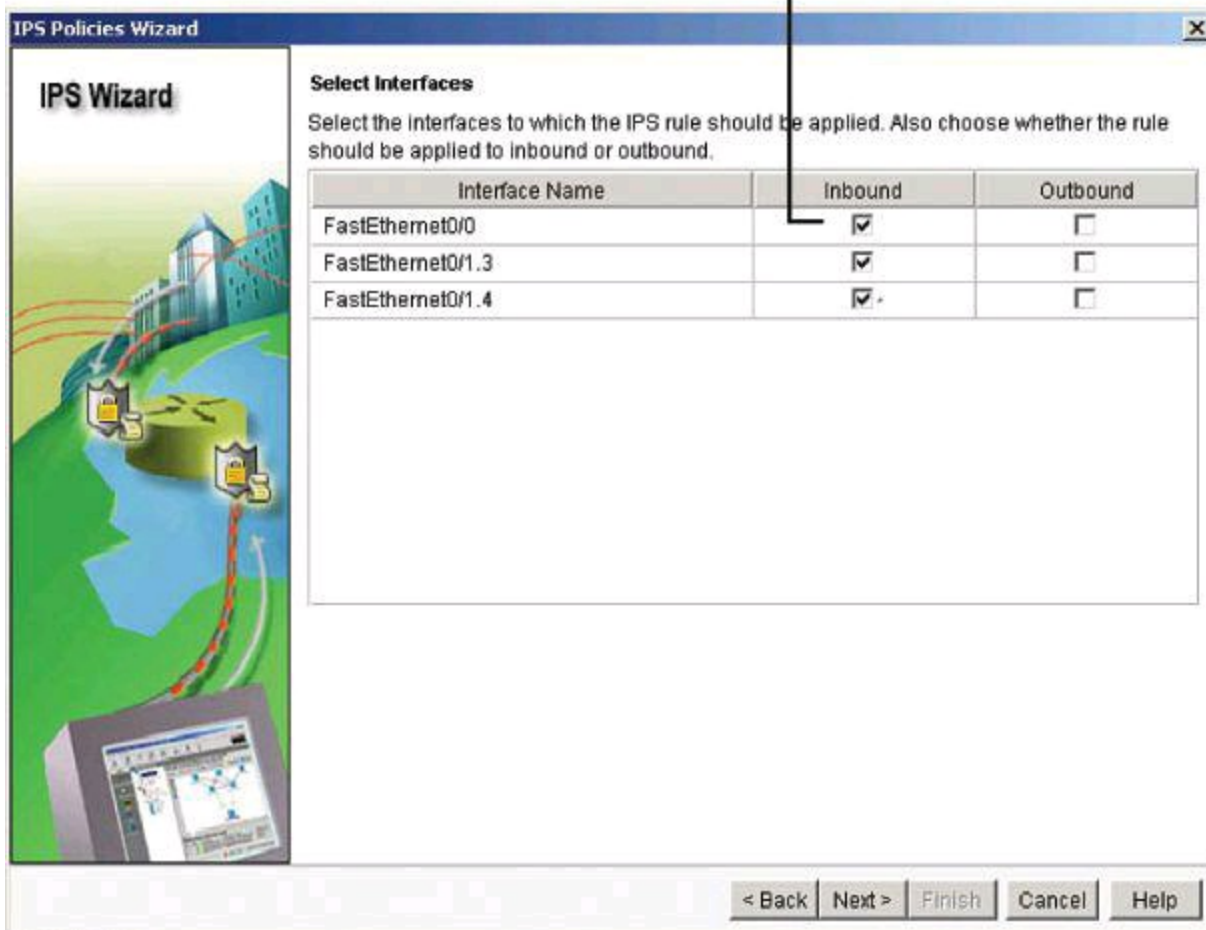


Figure 11-18. IPS Policies Wizard: Selecting the Interfaces

FastEthernet 0/0 Inbound and FastEthernet 0/1 Outbound

What would happen if you configured the IPS rules to be applied to FastEthernet 0/0 inbound and FastEthernet 0/1 outbound? You could get two alarms, one generated by the ingress rule and one by the egress rule.

On the Signature File and Public Key screen of the IPS Policies Wizard, you make the signature file available to the router. This is demonstrated in [Figure 11-19](#). In the Signature File section, click the first radio button, **Specify the Signature File You Want to Use with IOS IPS**. Then click the ... button to open the Specify Signature File dialog box, in which you specify the location of the signature package file.

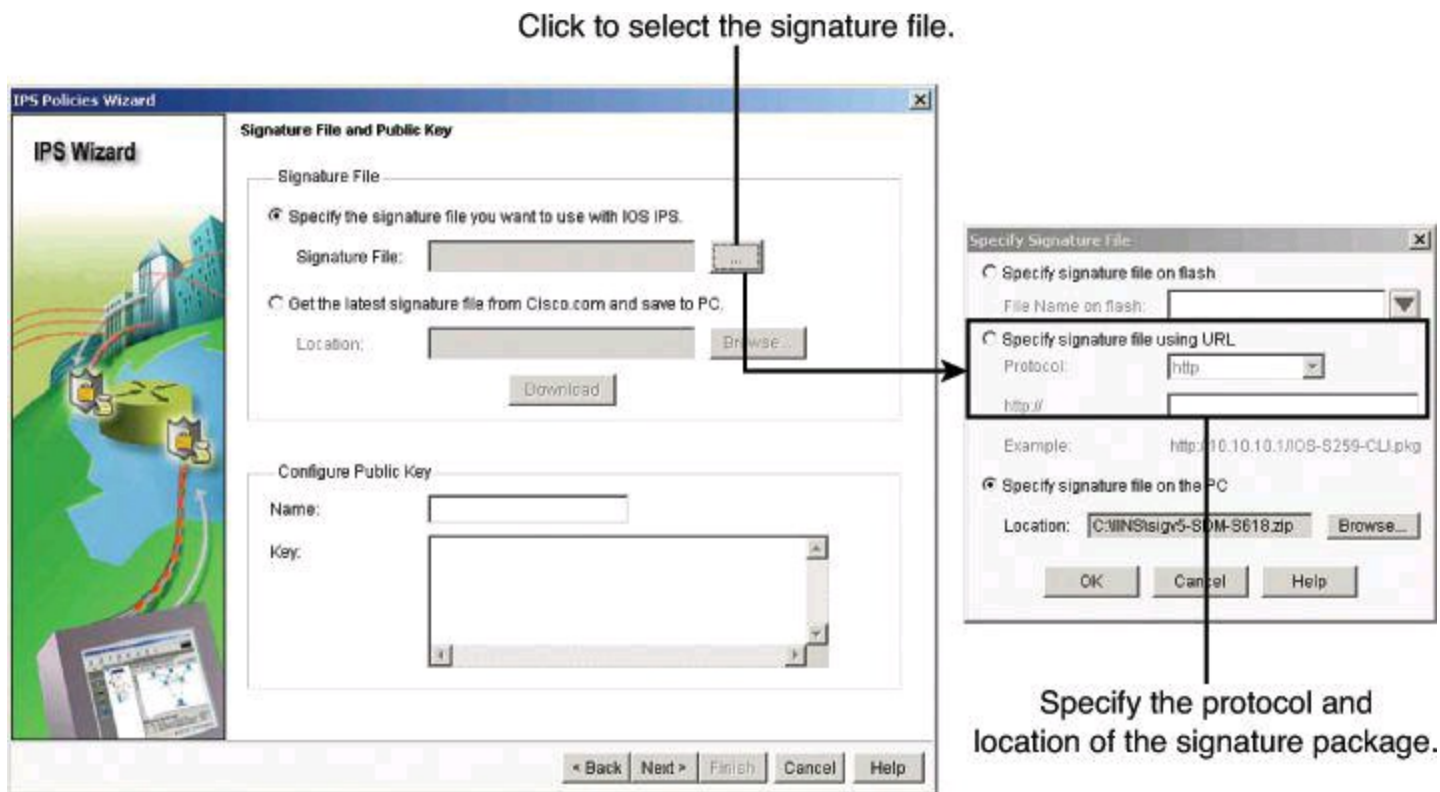
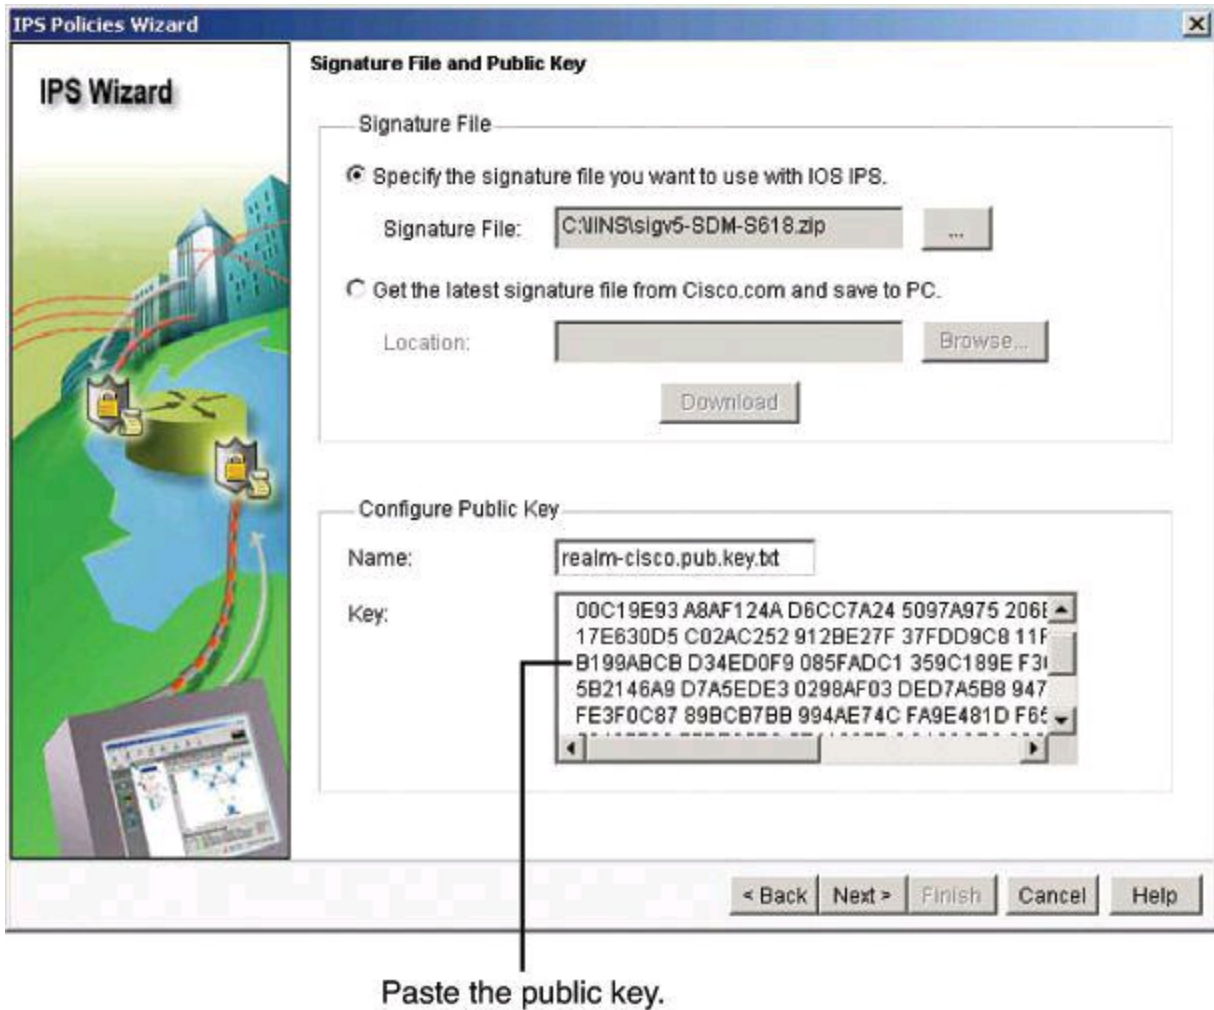


Figure 11-19. IPS Policies Wizard: Selecting the Signature File

You will need to make the previously downloaded signature file available via TFTP, FTP, and other options, such as having it resident on the local PC. In the Specify Signature File dialog box, choose the protocol and location URL. In this example, we are uploading the signature file from the local PC. When you are finished, click **OK** to return to the Signature File and Public Key screen.

Each change to the signature configuration is saved in the delta file. This file must be digitally signed with a public key. You can obtain a key from Cisco.com and copy and paste the information in the Name and Key fields. Public keys will be discussed in [Chapter 12](#), “[Fundamentals of Cryptography and VPN Technologies](#).”

[Figure 11-20](#) shows **realm-cisco.pub.key.txt** entered in the Name field of the Configure Public Key section, with the appropriate public key string copied and pasted in the Key field. This public key can be downloaded from Cisco.com at <http://download-sj.cisco.com/cisco/ciscosecure/ids/sigup/5.0/ios/real-cisco.pub.key.txt>.



Paste the public key.

Figure 11-20. IPS Policies Wizard: Downloading and Installing Cisco’s Public Key

Click **Next** to continue.

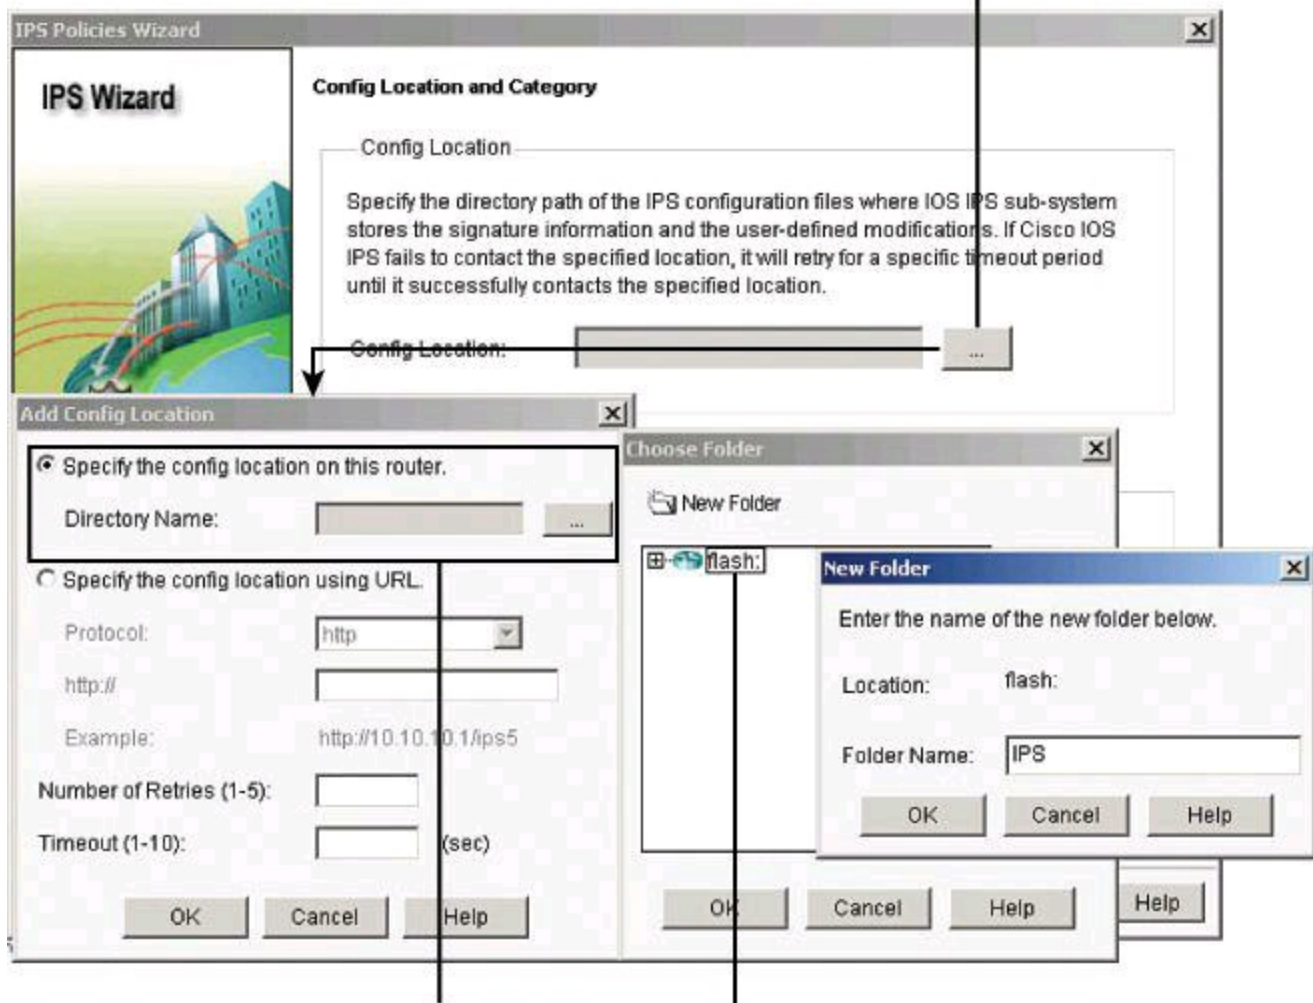
In the next wizard screen, Config Location and Category, you specify a location in the router flash for storing the signature information that Cisco IOS IPS will use. This information consists of the signature file and the delta file that is created when changes are made to the signature information. Follow these steps, shown in [Figure 11-21](#):

Step 1. In the Config Location and Category screen, select a location where the signature definition and configuration files will be stored by clicking the ... button to the right of the Config Location field.

Step 2. In the Add Config Location dialog box, choose the first option, **Specify the Config Location on This Router**, and then click the ... button to the right of the Directory Name field.

Step 3. In the Choose Folder dialog box, choose the router flash folder where the IPS files will be stored. If you want IPS files to be in their own folder, click on the **New Folder** button and give an intuitive name to the new folder. In our example, we called it **IPS**.

1. Click to select the IPS file's location on the router.

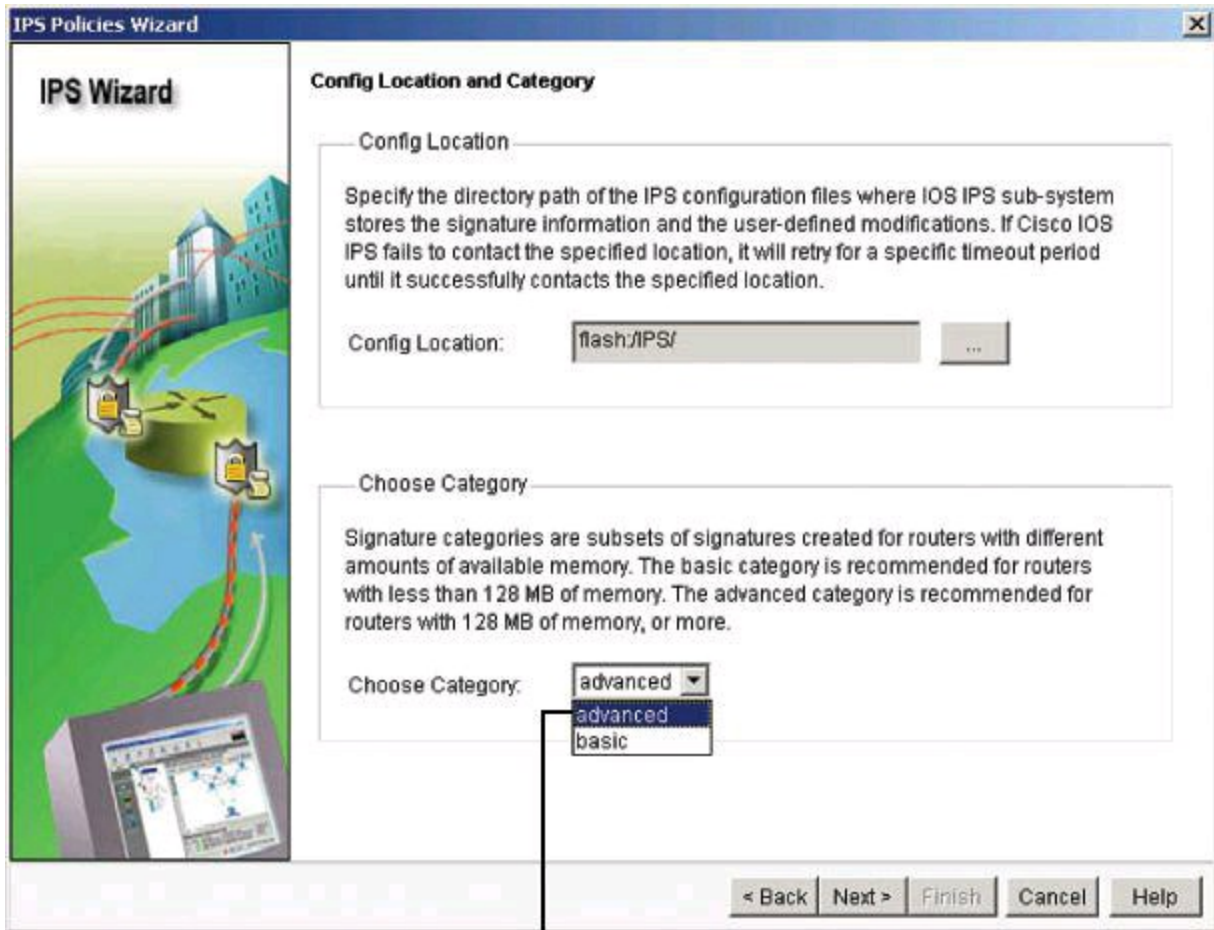


2. Click to select the location in Flash.

3. Click to select the specific folder.

Figure 11-21. IPS Policies Wizard: Storing Signature Information

After you click OK in the various dialog boxes to return to the Config Location and Category screen of the IPS Policies Wizard, choose a signature category from the Choose Category drop-down list, shown in [Figure 11-22](#), according to the amount of flash memory that is installed on the router. Because router memory and resource constraints may prevent the use of all the available signatures, there are two categories of signatures—basic and advanced. The basic category is appropriate for routers with less than 128 MB of available flash memory. The advanced category is appropriate for routers with more than 128 MB of available flash memory. The advanced category requires additional licenses to be obtained.



Select the advanced signature category.

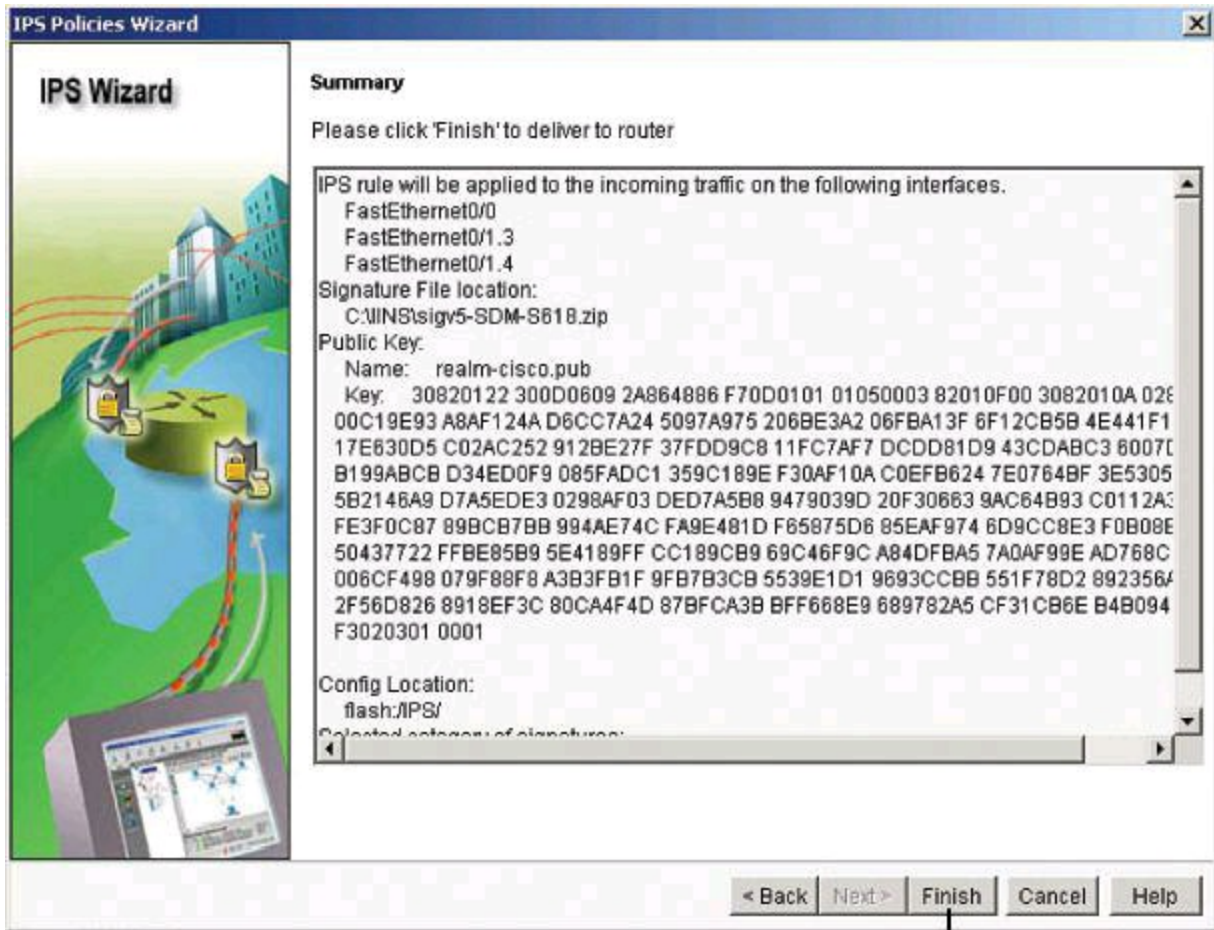
Figure 11-22. IPS Policies Wizard: Configuring Location and Signature Category

Note

If you have not activated the Cisco IOS IPS License, then installation of advanced signature packages will fail. Other signature packages can be installed without the license.

In this example, select **advanced**, since we have enough memory, and click **Next**.

The Summary screen, shown in [Figure 11-23](#), provides a brief description of the tasks for the Cisco IOS IPS initial configuration. Click **Finish** on the Summary screen to deliver the configurations and signature package to the router.



Review settings and click Finish.

Figure 11-23. IPS Policies Wizard: Summary Configuration

Step 3: Verify Configuration and Signature Files

When the signatures are loaded, Cisco Configuration Professional then displays the Edit IPS tab (**Configure > Security > Intrusion Prevention > Edit IPS**) with the current configuration. Verify the configuration by checking which interface and in what direction Cisco IOS IPS is enabled.

On the Edit IPS tab, you can also enable or disable Cisco IOS IPS on an interface and view information about how Cisco IOS IPS is applied, by clicking IPS Policies as shown in [Figure 11-24](#). If you enable Cisco IOS IPS on an interface, you can specify which traffic to examine for intrusion.

Intrusion Prevention System (IPS)

Create IPS **Edit IPS** Security Dashboard IPS Sensor IPS Migration

Interfaces: All Interfaces Enable Edit Disable Disable All

Interface Name	IP	Inbound	Outbound	VFR status	Description
FastEthernet0/0	200.200.1.2	Enabled	Disabled	on	Outside
FastEthernet0/1.3	172.16.1.1	Enabled	Disabled	on	DMZ
FastEthernet0/1.4	10.10.0.1	Enabled	Disabled	on	Inside

Review the IPS status per interface.

IPS Filter Details: Inbound Filter Outbound Filter

⚠️ IPS rule is enabled, but there is no filter configured for this rule. IPS will allow all Inbound traffic.

IPS Policies
Global Settings
Download
Auto Update
SEAP Configuration
Target Value Rating
Event Action Overrides
Event Action Filters
Signatures

Figure 11-24. Reviewing IPS Configuration and Interface Status

On the Edit IPS tab, you can click **Signatures** to verify the signature numbers and start planning your resource preservation strategy based on the router model, as shown in [Figure 11-25](#).

Review the number of signatures.

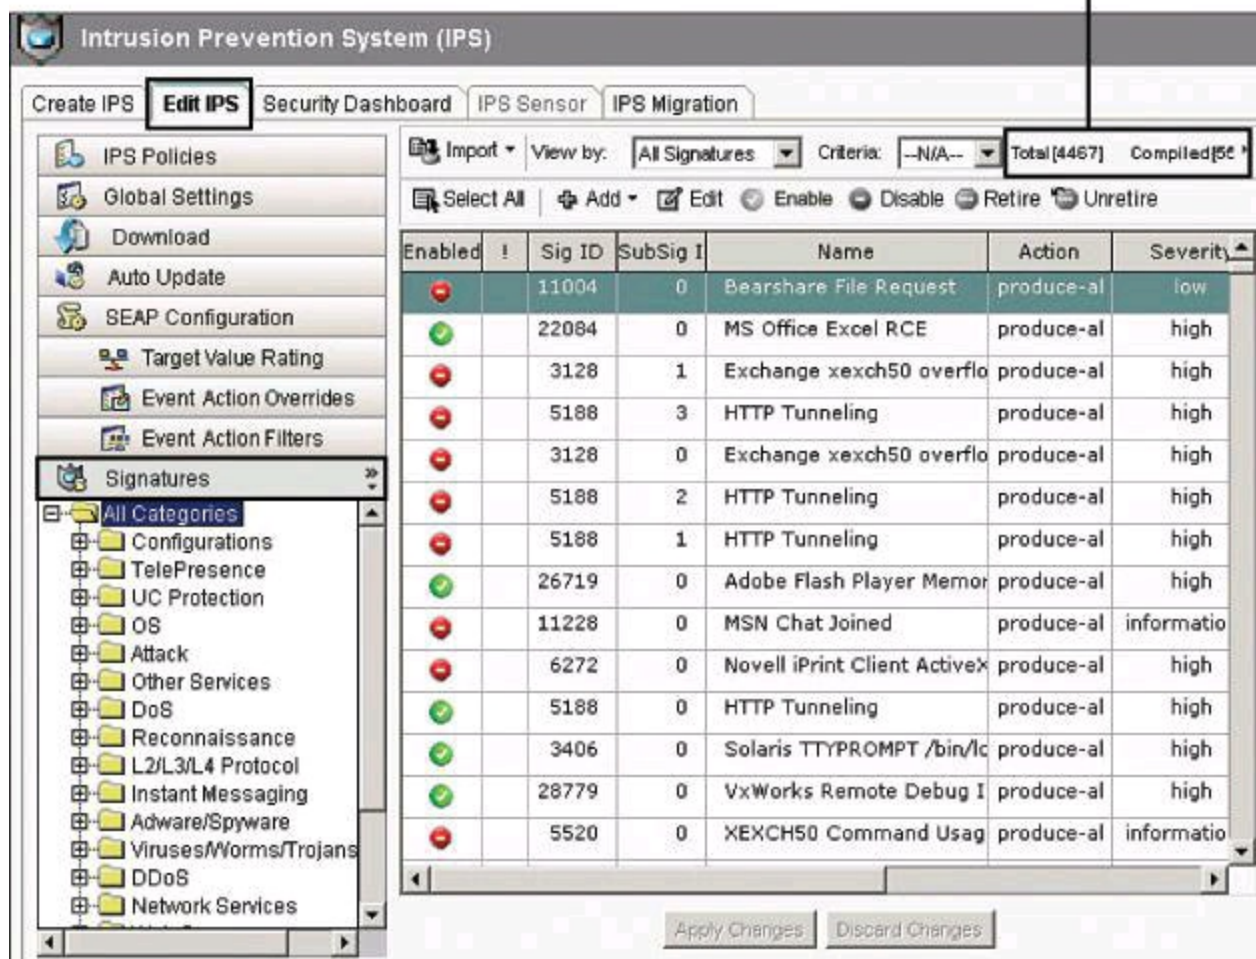


Figure 11-25. Reviewing IPS Signatures

Compiled Signatures

In [Figure 11-25](#), notice the Sub-Sig ID column. When there are multiple ways for Cisco IPS to recognize the same kind of attack, all those signatures will have the same Signature ID number (Sig ID), but each will have a different Sub-Signature ID number (Sub-Sig ID).

Also notice the Compiled count in the upper-right corner of [Figure 11-25](#). This represents the number of unretired signatures. “Compiled” and “unretired” are synonyms in this context.

Step 4: Perform Signature Tuning

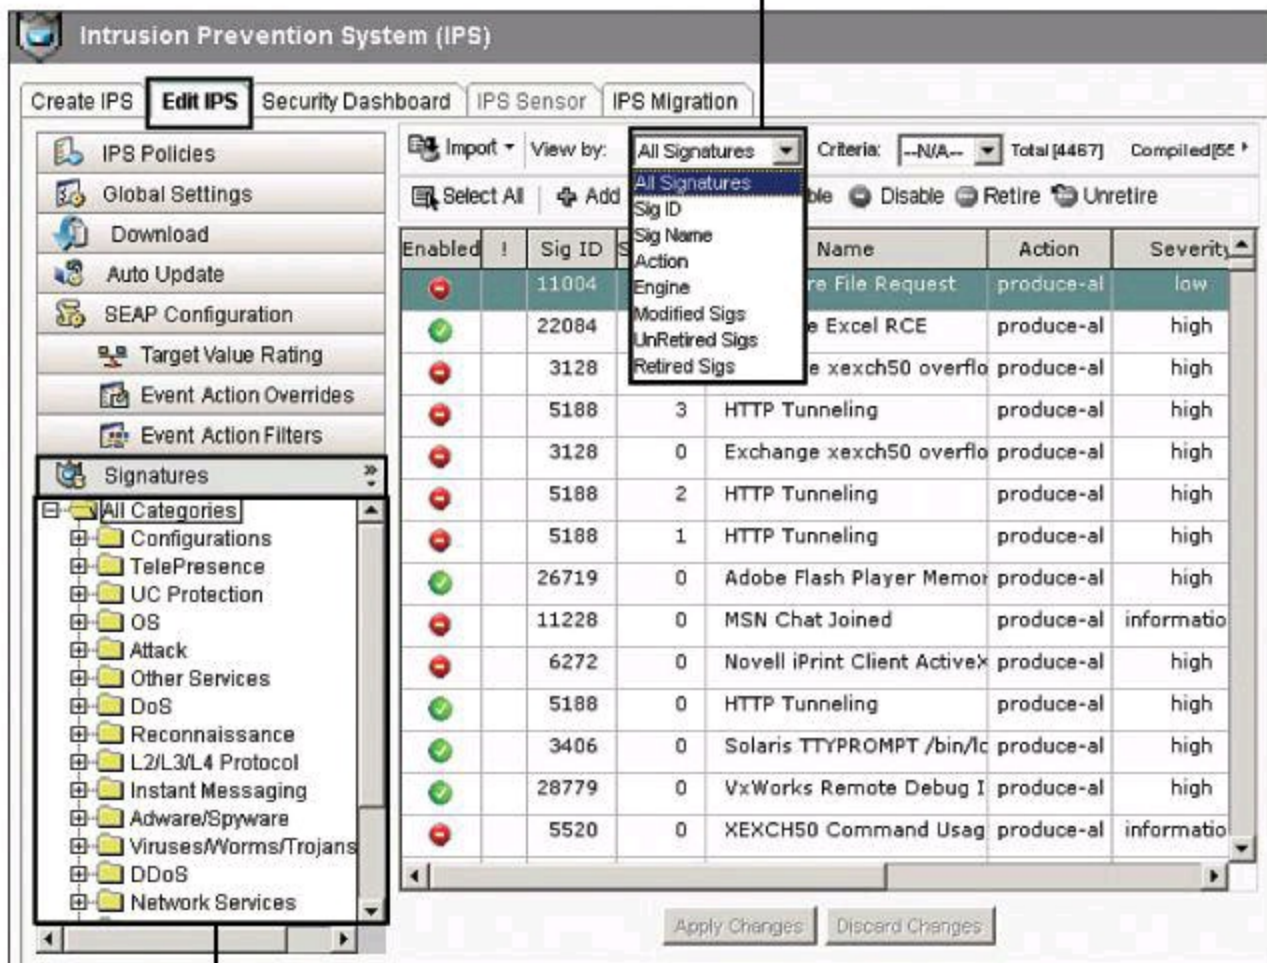
The most basic level of Cisco IOS IPS tuning manages resource preservation. Navigate to **Configure > Security > Intrusion Prevention > Edit IPS > Signatures** to tune the signatures, as shown in [Figure 11-26](#). To find specific signatures or groups of signatures, you can use these filtering mechanisms:

- The signature tree enables you to filter the signature list on the right according to the type of signature that you want to view. First choose the branch for the general type of signature that you want to display. The signature list displays the configured signatures for the type that you chose. If a plus (+) sign appears to the left of the branch, there are subcategories that you can use to refine the filter. Click the + sign to expand the branch and then choose

the signature subcategory that you want to display. If the signature list is empty, there are no configured signatures available for that type. For example:

- To display all attack signatures, click the Attack branch folder.
- To see the subcategories that you can use to filter the display of attack signatures, click the + sign next to the Attack folder.
- To see denial of service signatures, click the DoS folder.
- The View By and Criteria drop-down lists enable you to filter the display according to the types of signatures that you want to view. First choose the criteria in the View By drop-down list, and then choose the value for that criteria in the Criteria drop-down list. For example, if you choose Engine in View By, Criteria changes to Engine, and you can choose among the available engines, such as ATOMIC.ICMP and SERVICE.DNS. If you choose Sig ID or Sig Name in the View By drop-down list, you must enter a value in the Criteria field.

Find signatures using these criteria.



Use signature tree to filter by category.

Figure 11-26. Looking at the Signatures

To retire or unretire and enable or disable signatures, select the signatures and then click Enable, Disable, Retire, or Unretire, based on your needs and the resources available on the router.

Notice how the status changes in the Enabled or the Retired column. A yellow icon appears for the signature in the column next to Enabled, as shown in [Figure 11-27](#). The yellow icon means that

changes have been made to the signature but have not been applied. Click **Apply Changes** to make the changes take effect.

The yellow icon indicates changes have been made but not applied.

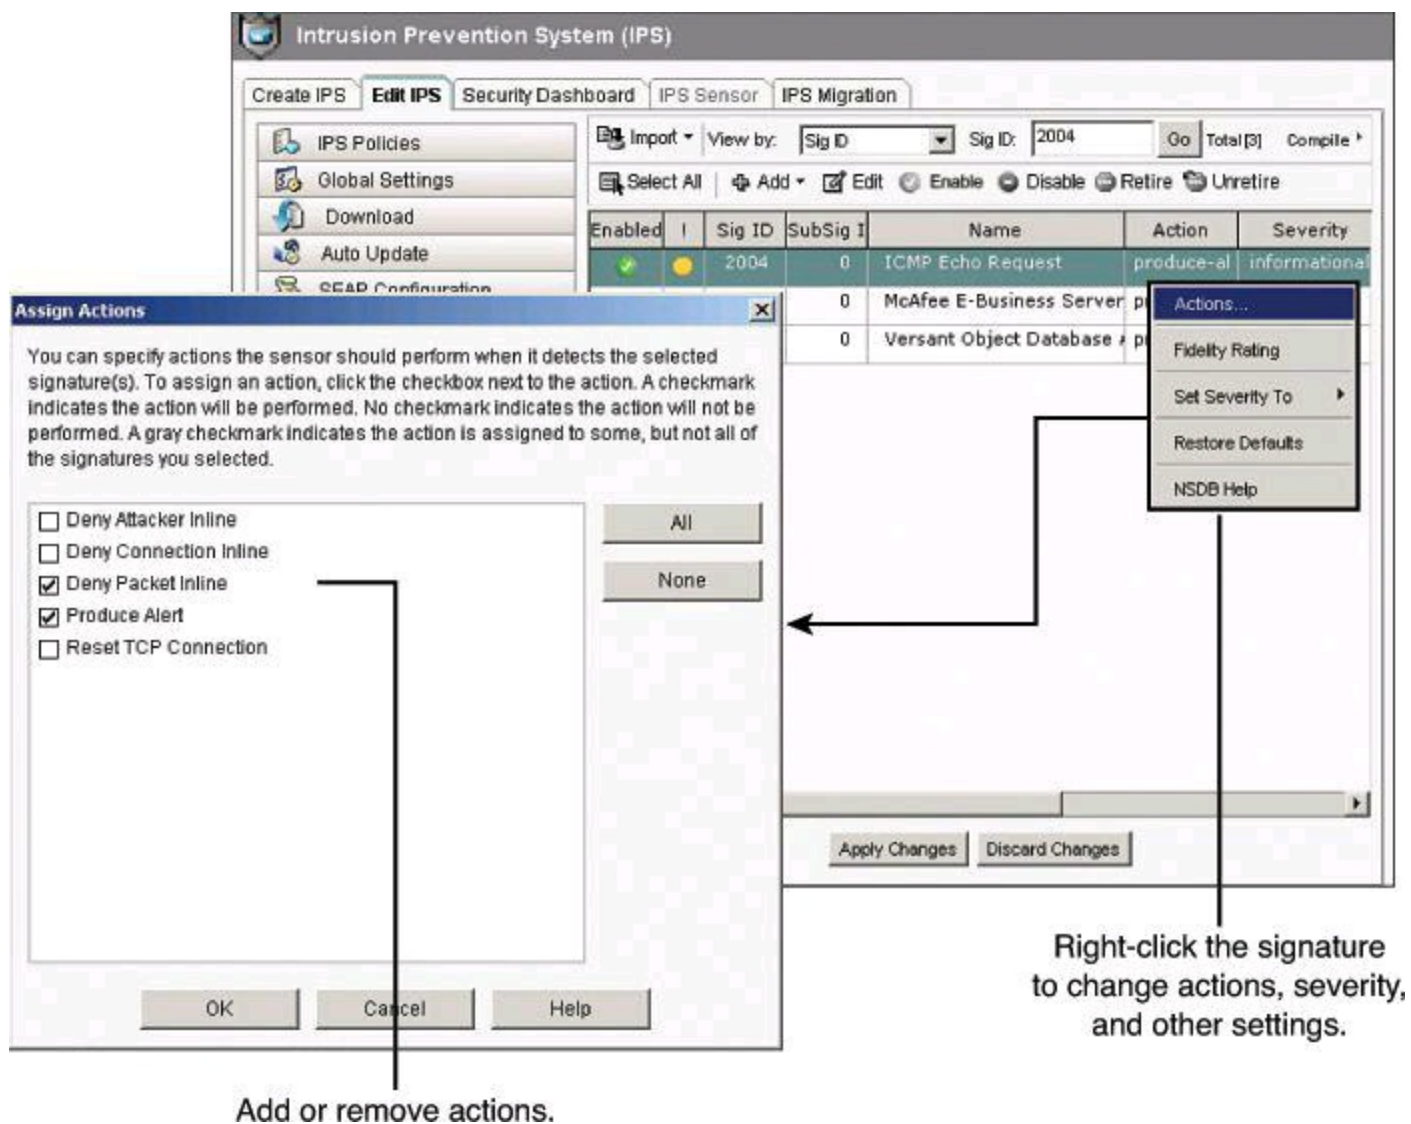
Edit, Enable, Disable, Retire, Unretire

Enabled	Sig ID	SubSig ID	Name	Action	Severity
<input checked="" type="checkbox"/>	2004	0	ICMP Echo Request	produce-al	informational
<input type="checkbox"/>	20041	0	McAfee E-Business Server	produce-al	high
<input type="checkbox"/>	20040	0	Versant Object Database	produce-al	high

Apply changes.

Figure 11-27. Enable, Disable, Retire, or Unretire Signatures

To change the action that is associated with a signature, select the signature, right-click and choose **Actions**, and then check or uncheck the check boxes next to the actions to choose the actions to be associated with this signature, as shown in [Figure 11-28](#).



Add or remove actions.

Right-click the signature to change actions, severity, and other settings.

Figure 11-28. Changing Action of Signatures

Following is a summary of the options available in the context menu when you right-click a signature:

- **Actions:** Click to choose the actions to be taken when the signature is matched.
- **Set Severity To:** Click to set the severity level of a signature to high, medium, low, or informational.
- **Restore Defaults:** Click to restore the default values of the signature.
- **Remove Filter:** Click to remove a filter applied to the signature.
- **NSDB Help:** Click to display help on the Network Security Database (NSDB) (requires a Cisco.com account).

Signature Actions

With Cisco IOS IPS, protection can be done in two ways. You can enable explicit protection on specific signatures, as shown in [Figure 11-28](#), or it can be enabled implicitly using Event Action Overrides to protect whenever the risk rating is at or above a specified threshold. Determining which signatures should have blocking actions can be a daunting task, so the simpler alternative is to use Event Action Overrides, which uses Event Risk Rating (ERR) to make the decision on your behalf.

Cisco IOS IPS Configuration = XML

When you apply changes to the Cisco IOS IPS configuration, you will not see a command preview in CCP. This is due to the fact that Cisco IOS IPS configuration is stored in XML documents in flash, and not in the running config. The files are in flash:/IPS, as you specified from the Config Location and Category screen of the IPS Policies Wizard.

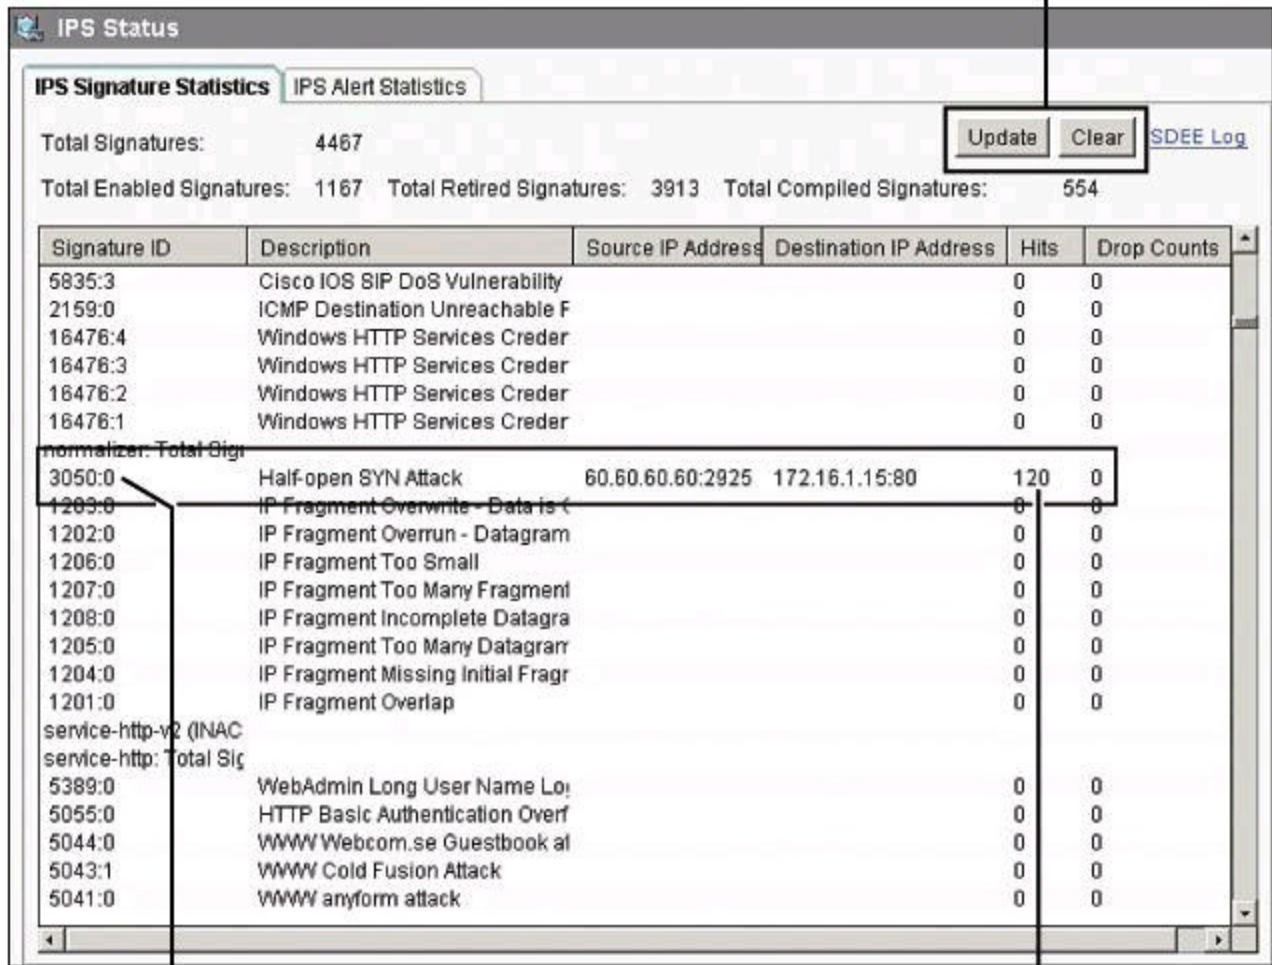
Step 5: Verify Alarms

Cisco Configuration Professional can be used as an event viewer for Cisco IOS IPS alarms. You can view the log by navigating to **Monitor > Security > IPS Status** and clicking the **IPS Signature Statistics** tab. (The IPS Status pane also includes an IPS Alert Statistics tab, which is discussed a bit later in this section.)

[Figure 11-29](#) illustrates the contents of the IPS Signature Statistics tab. Statistics are displayed for each enabled signature in the Cisco IOS IPS configuration. The top of the tab displays signature totals to provide a snapshot of the signature configuration. The following totals are provided:

- Total Signatures
- Total Enabled Signatures
- Total Retired Signatures
- Total Compiled Signatures

Update refreshes the view;
Clear resets counters.



Alarms are displayed per signature.

Hit Count per Source/Destination
Combination

Figure 11-29. Monitoring IPS Signature Statistics from CCP

The statistics table displays alarm counters per signature. The table includes the signature ID, signature description, number of hits or times the signature has been matched, and a drop count that represents the number of times the packet has been dropped due to this event. If a packet that matches a signature arrives, the source and destination IP addresses are listed as well.

The table is not updated automatically. You can click **Update** to check for and include the latest signature statistics. You can also click **Clear** to set all signature statistic counters to 0.

More alarm detail can be observed by displaying the SDEE log. Click the **SDEE Log** link, shown in [Figure 11-30](#), of the IPS Signature Statistics tab to display the log. CCP automatically navigates to **Monitor > Router > Logging** and displays the SDEE Message Log tab, shown in [Figure 11-30](#). Some of the alarms relate to the status of the Cisco IOS IPS service, while some others relate to Cisco IOS IPS service error conditions. Click the SDEE Messages drop-down menu to filter and display Alerts only.

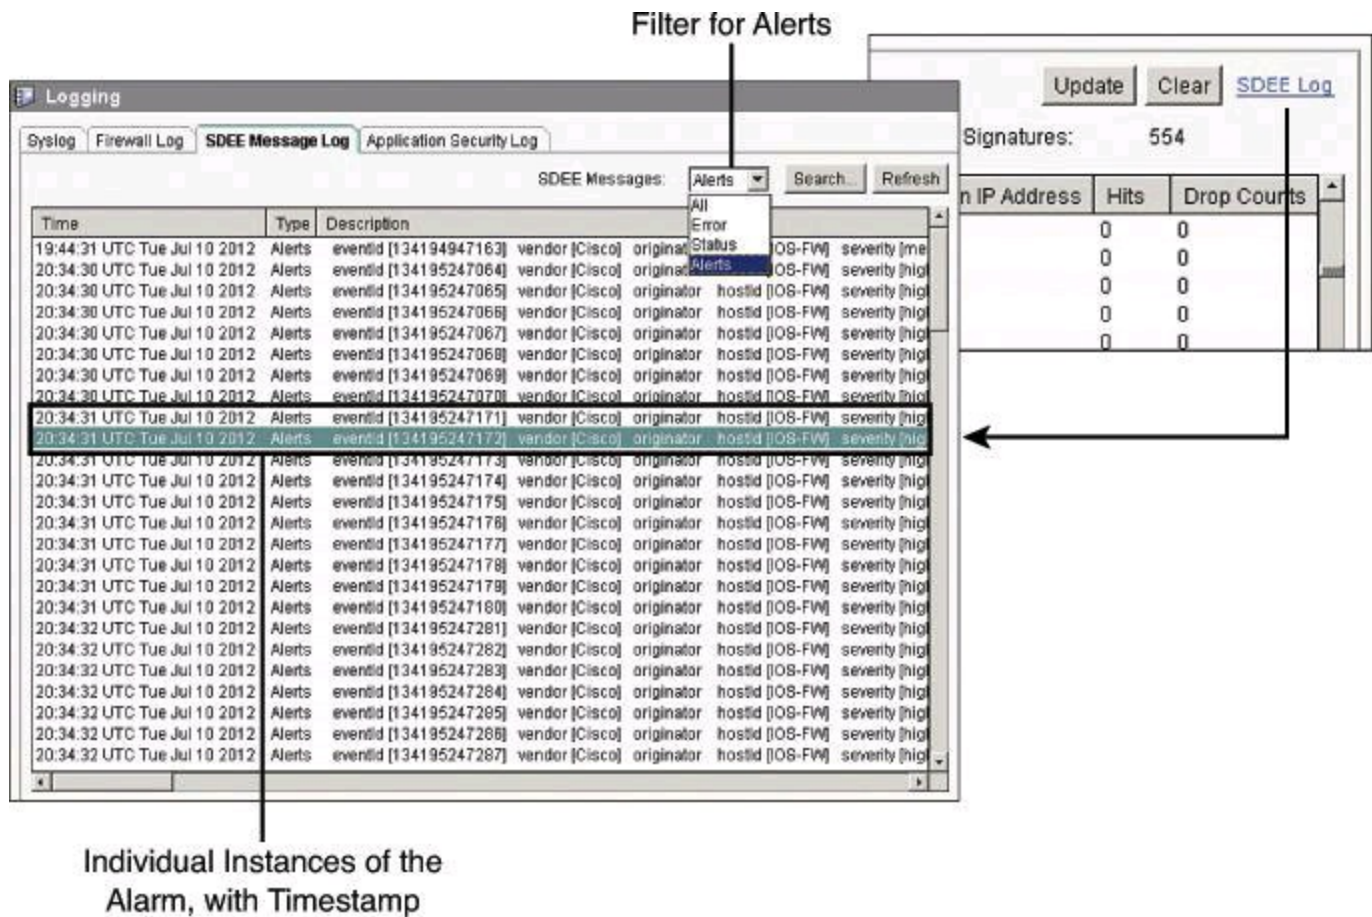


Figure 11-30. Monitoring IPS Alarms from CCP

The list includes all alarms in the router buffer, including the following components:

- **Time:** Timestamp for the individual alert
- **Type:** Alerts, Error, or Status entry
- **Description:** Includes event ID, signature IDs, risk rating, and other pieces of alarm-related information

The list is searchable. Clicking Search opens a search window, where you can choose a search type from the Search menu, enter the appropriate text in the Search field, and then click Find to display matching log entries.

The search types include the following:

- Source IP Address
- Destination IP Address
- Text

Searches are not case sensitive.

The IPS Alert Statistics tab, on the other hand, displays a view from the alerts and risk rating perspective. The panel displays alert statistics in a color-coded format for easy recognition. The top part of the screen shown in [Figure 11-31](#) displays a legend that explains the use of colors in the display.

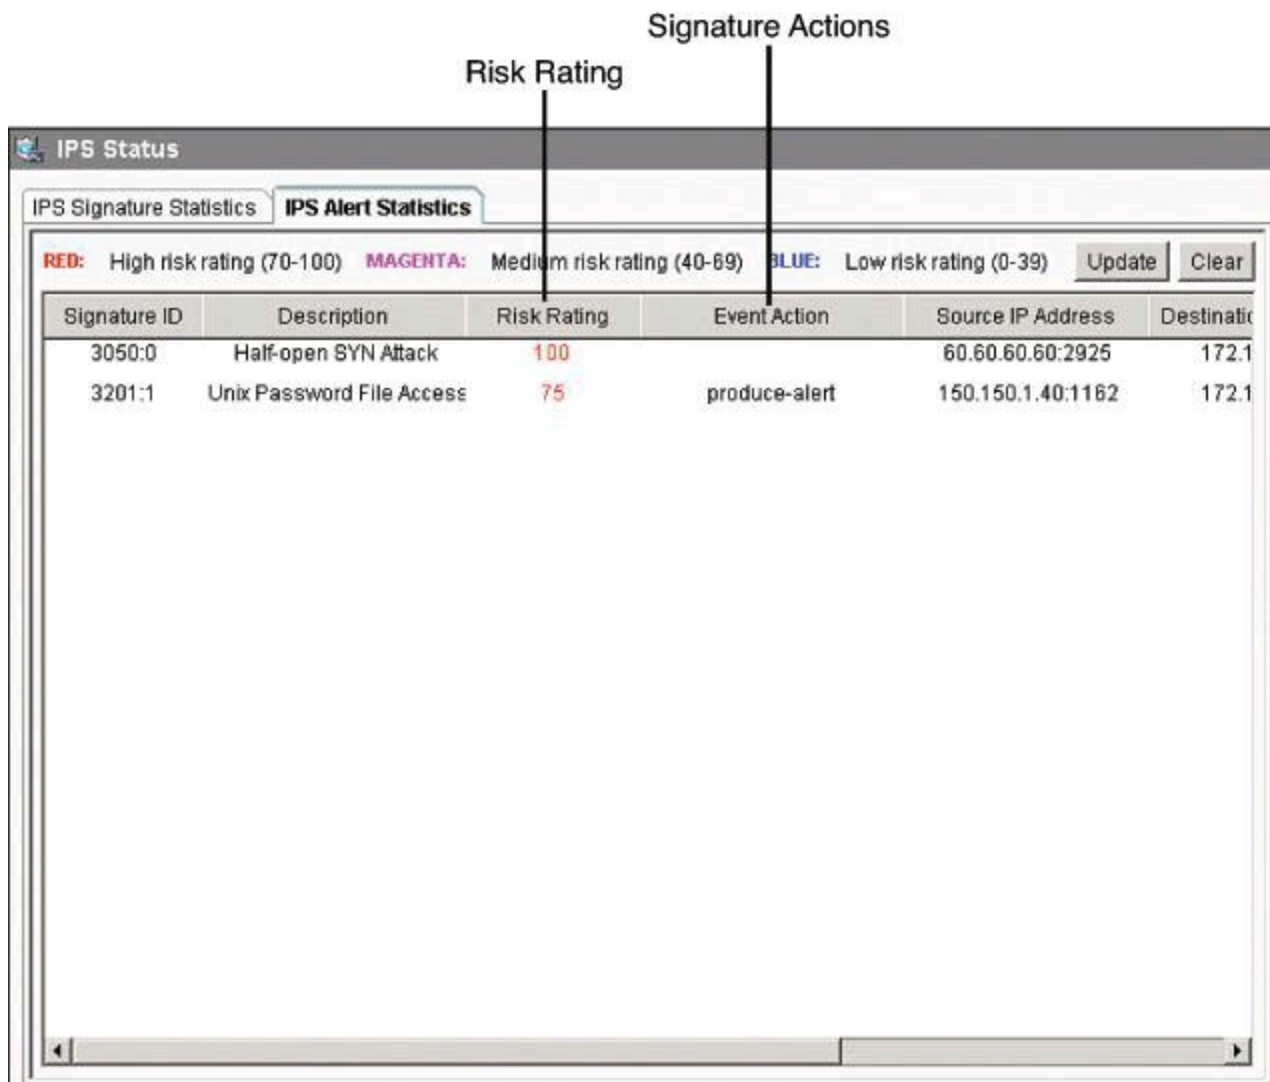


Figure 11-31. IPS Signature Statistics

The alert color coding is presented in [Table 11-10](#).

Table 11-10. Alert Color Coding

Color	Description
Red	The event that generated the alert has a high risk rating in the range of 70 to 100.
Magenta	The event that generated the alert has a medium risk rating in the range of 40 to 69.
Blue	The event that generated the alert has a low risk rating in the range of 0 to 39.

By clicking a column heading, you sort the display based on the values of that parameter. For example, click the Signature ID heading to sort the display in ascending or descending numerical order of signature IDs. Each column is described in the following list:

- **Signature ID:** Numerical signature identifier.
- **Description:** Description of the signature.
- **Risk Rating:** Value between 0 and 100 that represents a numerical quantification of the risk associated with a particular event on the network.
- **Event Action:** Action that Cisco IOS IPS takes when an event matching the signature

occurs.

- **Source IP Address:** The IP address from which the packet originated.
- **Destination IP Address:** The IP address to which the packet was addressed. If the packet is malicious, the destination IP address can be considered the target.
- **Hits:** Number of matching packets.
- **Drop Count:** The number of matching packets dropped.
- **Engine:** The signature engine associated with the signature.

Configuring Cisco IOS IPS Using the CLI

To use the command-line interface (CLI) to specify an IPS rule, use the **ip ips name name** command in global configuration mode as follows:

```
router(config)# ip ips name sdm_ips_rule
```

In the preceding command, the IPS rule had been created using CCP. Commands created by CCP appear in the CLI with the notation “sdm,” for Security Device Manager, which was the predecessor to CCP.

To specify the location of the IPS configuration, use the **ip ips config location location** global configuration command, as demonstrated here:

```
router(config)# ip ips config location flash:/ips/retries 1
```

To specify the method of event notification, use the **ip ips notify** global configuration command. The following is an example of event notification sent using SDEE:

```
router(config)# ip ips notify SDEE
```

Note

Examples in this section that deal with Cisco IOS IPS CLI configuration assume that the signature files are already on the router.

To configure the router to support the default basic signature set, use the **ip ips signature-category** global configuration command as follows:

```
Router(config)# ip ips signature-category
Router(config-ips-category)# category all
Router(config-ips-category-action)# retired true
Router(config-ips-category-action)# exit
Router(config-ips-category)# category ios_ips basic
Router(config-ips-category-action)# retired false
```

To apply an IPS rule to an interface, use the **ip ips ips_rule_name** command in interface configuration mode as demonstrated here:

```
router(config)# interface FastEthernet0/0
router(config-if)# ip ips sdm_ips_rule in
```

You should also consider turning on packets virtual reassembly on the sensing interface. Virtual

Fragment Reassembly (VFR) enables the Cisco IOS Firewall to examine out-of-sequence fragments and reorder the packets into the correct order. It examines the number of fragments from a same single IP address. When VFR is enabled on the Cisco IOS Firewall, it creates the appropriate dynamic ACLs, thereby protecting the network from various fragmentation attacks. To enable VFR on an interface, use the **ip virtual-reassembly** command in interface configuration mode, as demonstrated here:

```
Router(config)# interface FastEthernet0/0
Router(config-if)# ip virtual-reassembly
```

[Example 11-1](#) provides a combined view of the commands shown in the preceding paragraphs.

Example 11-1. Cisco IOS IPS CLI Configuration

[Click here to view code image](#)

```
Router(config)# ip ips name sdm_ips_rule
Router(config)# ip ips config location flash:/ips/ retries 1
Router(config)# ip ips notify SDEE
!
Router(config)# ip ips signature-category
Router(config-ips-category)# category all
Router(config-ips-category-action)# retired true
Router(config-ips-category-action)# exit
Router(config-ips-category)# category ios_ips basic
Router(config-ips-category-action)# retired false
!
Router(config)# interface Ethernet0/0
Router(config-if)# ip ips sdm_ips_rule in
Router(config-if)# ip virtual-reassembly
```

Use the **show ip ips configuration** command to display additional configuration data that is not displayed with the **show running-config** command. [Example 11-2](#) shows some sample output from the **show ip ips configuration** command.

Example 11-2. *show ip ips configuration* Command Output

[Click here to view code image](#)

```
Router# show ip ips configuration
IPS Signature File Configuration Status
Configured Config Locations: flash:/ips/
Last signature default load time: 04:39:33 UTC Oct 19 2011
Last signature delta load time: 04:41:43 UTC Oct 19 2011
Last event action (SEAP) load time: -none-

General SEAP Config:
Global Deny Timeout: 3600 seconds
Global Overrides Status: Enabled
Global Filters Status: Enabled
```

```
IPS Auto Update is not currently configured
```

```
IPS Syslog and SDEE Notification Status  
Event notification through syslog is enabled  
Event notification through SDEE is enabled
```

```
IPS Signature Status  
Total Active Signatures: 295  
Total Inactive Signatures: 4008
```

```
IPS Packet Scanning and Interface Status  
IPS Rule Configuration  
IPS name sdm_ips_rule  
IPS fail closed is disabled  
IPS deny-action ips-interface is false  
Obsolete tuning is disabled  
Regex compile threshold (MB) 14  
Interface Configuration  
Interface FastEthernet0/0  
Inbound IPS rule is sdm_ips_rule  
Outgoing IPS rule is not set  
Interface FastEthernet0/1  
Inbound IPS rule is sdm_ips_rule  
Outgoing IPS rule is not set
```

```
IPS Category CLI Configuration:  
Category all:  
Retire: True  
Category ios_ips basic:  
Retire: False
```

You could use the **show ip ips all** command (not shown here) to display additional configuration data that is not displayed with the **show ip ips configuration** command.

Use the **show ip ips interface** command to display interface configuration data. [Example 11-3](#) displays output from the **show ip ips interface** command, revealing that the inbound IPS audit rule **sdm_ips_rule** is applied to FastEthernet 0/0 and FastEthernet 0/1. There is no rule applied for outgoing traffic on either interface.

Example 11-3. *show ip ips interfaces* Command Output

[Click here to view code image](#)

```
Router# show ip ips interfaces  
Interface Configuration  
Interface FastEthernet0/0  
Inbound IPS rule is sdm_ips_rule  
Outgoing IPS rule is not set  
Interface FastEthernet0/1  
Inbound IPS rule is sdm_ips_rule  
Outgoing IPS rule is not set
```

In [Example 11-4](#), the output from the **show ip ips signature count** command displays summary signature information. The output shows the signature update ID, as well as the number of total, enabled, retired, compiled, and obsolete signatures. Remember that “compiled” means unretired.

Example 11-4. *show ip ips signature count* Command Output

[Click here to view code image](#)

```
Router# show ip ips signature count
Cisco SDF release version S594.0
Trend SDF release version V0.0
...
<output omitted>
...
Total Signatures: 4303
Total Enabled Signatures: 1218
Total Retired Signatures: 4008
Total Compiled Signatures: 295
Total Obsoleted Signatures: 13
Total Disallowed Signatures: 3
:
.
```

The router will also generate system log messages, providing an additional tool to monitor the health of the IPS feature. The following message is displayed when an IPS signature engine has been built and is ready to scan packets:

```
%IPS-6-ENGINE_READY:SERVICE.HTTP - 183136 ms - packets for this engine
will be
  scanned
```

The following message is displayed when packets are being dropped due to a failed IPS module, when the **ip ips fail closed** command is configured. The **ip ips fail closed** command implements a restrictive fail-closed policy when the Cisco IOS IPS feature fails.

```
%IPS-5-PACKET_DROP:SERVICE.DNS - packets dropped while engine is
building
```

The following message, discussed previously in this chapter, is displayed when an IPS signature has been triggered. The signature and subsignature IDs are shown. In this particular example, the alarm is related to a potential spoofing attack, as explained earlier.

```
%IPS-4-SIGNATURE:Sig:1107 Subsig:0 Sev:2 RFC1918 address
[192.168.121.1:137
->192.168.121.255:137]
```

Use the **show ip ips statistics** command to show a global view of signatures that have triggered, along with valuable information such as the number of packets checked per signature, the number of alarms triggered per signature, and the number of packet drops per signature. The output of [Example 11-5](#) also shows the number of interfaces enabled for Cisco IOS IPS.

Example 11-5. *show ip ips statistics* Command Output

[Click here to view code image](#)

```
Router# show ip ips statistics
Signature statistics [process switch:fast switch]
signature 3041:0: packets checked [0:4] alarmed [0:4] dropped [0:0]
signature 3040:0: packets checked [0:1] alarmed [0:1] dropped [0:0]
signature 6062:1: packets checked [0:1] alarmed [0:1] dropped [0:0]
signature 6054:0: packets checked [0:3] alarmed [0:3] dropped [0:0]
Interfaces configured for ips 1
Session creations since subsystem startup or last reset 10101
Current session counts (estab/half-open/terminating) [15:764:0]
Maxever session counts (estab/half-open/terminating) [22:1182:0]
Last session created 00:00:08
Last statistic reset never
TCP reassembly statistics
received 2 packets out-of-order; dropped 0
peak memory usage 1 KB; current usage: 0 KB
```

Use the **show ip sdee alerts** command to see a detailed view of the alarm log, as shown in [Example 11-6](#).

Example 11-6. *show ip sdee alerts* Command Output

[Click here to view code image](#)

```
router# show ip sdee alerts
Alert storage: 200 alerts using 96000 bytes of memory
SDEE Alerts
SigID      Sig Name          SrcIP:SrcPort      DstIP:DstPort      VRF
           or Summary Info
1: 3040:0  TCP NULL Packet   172.17.44.101:36044 10.5.5.5:25        NONE
2: 3041:0  TCP SYN/FIN Packet 172.17.44.101:52623 10.5.5.5:25        NONE
3: 6054:0  DNS Version
Request 172.17.44.101:4745 172.17.22.103:53  NONE
4: 6062:1  DNS Authors
Request 172.17.44.101:4756 172.17.22.103:53  NONE
```

Summary

This chapter described how intrusion detection system (IDS) and intrusion prevention system (IPS) technology embedded in Cisco host- and network-based IDS and IPS solutions fights malware in real time. More precisely, you have learned how

- Cisco IOS IPS prevents intrusion by comparing traffic against the signatures of known attacks. A signature is a set of rules that an IDS and an IPS use to detect typical intrusive activity and how these files can be updated manually or automatically. A public key is required to use digitally signed signature files.
- Signatures can be retired or unretired, enabled or disabled.
- Cisco IOS IPS alarms are communicated using SDEE and syslog.
- The easiest way to configure Cisco IOS IPS on the router or security device is to click the

Launch IPS Rule Wizard button in Cisco Configuration Professional to launch the IPS Policies Wizard.

- Cisco Configuration Professional, which includes comprehensive options for configuring IPS and for signature tuning, can also be used to manage Cisco IOS IPS events.
- CCP tools can be used to monitor IPS operations.

References

For additional information, refer to these resources.

Cisco.com Resources

“Cisco IOS IPS Q&A,”
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6634/prod_qas09
Cisco IOS Security Configuration Guide, Release 12.4, “Configuring Cisco IOS Intrusion Prevention System (IPS),”
[http://www.cisco.com/en/US/docs/ios/security/configuration/guide/12_4/sec_12_4_book.h](http://www.cisco.com/en/US/docs/ios/security/configuration/guide/12_4/sec_12_4_book.html)
Cisco Security Information Event Management Deployment Guide,
[http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns982/sbaSIEM_dep](http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns982/sbaSIEM_dep.html)
“Getting Started with IOS IPS: A Step-by-Step Guide,”
http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6634/prod_white
“Intrusion Prevention System,” <http://www.cisco.com/go/ips>

General IDS/IPS Resource

SearchSecurity.com, <http://searchsecurity.techtarget.com/>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

1. Which two modes of IPS operation are currently available with Cisco IDS and IPS solutions?
 - a. Out-of-band
 - b. Promiscuous
 - c. Multicasting
 - d. Inline
2. IDS solutions are typically related to inline deployment of sensors. True or false?
 - a. True
 - b. False
3. Which general patterns of misuse do IDS and IPS technologies look for? (Choose all that apply.)
 - a. Atomic pattern
 - b. Molecular pattern

- c. Intrusive nonces
 - d. Composite (compound) pattern
 - e. Composition pattern
4. You need to customize signatures on Cisco IPS devices to detect ICMP ping exploits packet by packet. Which signature microengine should you use?
- a. Service
 - b. Multi-string
 - c. Atomic
 - d. Sub-atomic
5. Which of the following are a type of IDS or IPS sensor? (Choose all that apply.)
- a. Signature based
 - b. Policy based
 - c. Transgression based
 - d. Anomaly based
6. Match the technology with its definition?
- a. Signature-based IPS
 - b. Policy-based IPS
 - c. Reputation-based IPS
 - d. Anomaly-based IPS
1. Normal behavior typically defined based on traffic patterns, traffic and protocol mix, traffic volumes, and other criteria
 2. Typically implemented in the form of white lists or black lists
 3. Can produce false positives because certain normal network activity can be misinterpreted as malicious activity
 4. Similar to implementing a restrictive firewall policy
7. What is a signature engine?
- a. A set of rules that an IDS and an IPS use to detect typical intrusive activity
 - b. A full-feature intrusion prevention tool located in the core network fabric device
 - c. An internal security service module that provides dedicated CPU and memory to offload intrusion prevention processing.
 - d. A component of an IDS and IPS sensor that supports a group of signatures in a common category
8. Which IPS card could integrate into an ISR router?
- a. Cisco IDSM-2
 - b. Cisco ASA AIP SSM
 - c. Cisco IPS AIM
 - d. Cisco IPS 4200 Series Sensor

- 9.** What is an IPS signature?
- a.** A message digest encrypted with the sender's private key
 - b.** A set of rules used to detect typical intrusive activity
 - c.** A binary pattern specific to a virus
 - d.** An appliance that provides anti-x services
- 10.** Compiling a regular expression found in a signature requires more memory than the final storage of the regular expression. True or false?
- a.** True
 - b.** False
- 11.** Which two choices are examples of techniques used to mitigate an evasion attack?
- a.** Data normalization
 - b.** Botnet detection
 - c.** TCP intercept
 - d.** Session reassembly
 - e.** False positive mitigation
- 12.** Cisco IOS IPS supports risk-based signature tuning using Cisco Event Risk Rating. True or false?
- a.** True
 - b.** False
- 13.** The signature package named sigv5-SDM-S592.zip is an example of what type of signature package in Cisco IOS IPS?
- a.** Cisco Configuration Professional signature package
 - b.** CLI signature package
 - c.** Basic signature package
 - d.** Advanced signature package
 - e.** Default signature package
- 14.** You are tuning IPS signatures using Cisco Configuration Professional and see a yellow icon next to the signature ID. What is the reason for this icon?
- a.** The signature is a candidate for false positives.
 - b.** The signature is categorized with the wrong microengine.
 - c.** You need to click Apply.
 - d.** You need to enable the signature.
 - e.** The signature is disabled and you need to unretire it.
- 15.** Which statement is true about this Cisco IOS IPS syslog message?

```
%IPS-5-PACKET_DROP:SERVICE.DNS - packets dropped while engine is building
```

- a.** A fail-closed strategy is configured.

- b.** The service.dns microengine is disabled.
- c.** Packets are blocked by the action of a signature matching the traffic pattern.
- d.** An inline deployment option is being used.

Part IV: Secure Connectivity

Chapter 12. Fundamentals of Cryptography and VPN Technologies

This chapter introduces the concepts of cryptography and VPN technologies. It covers the following topics:

- Need for VPN and VPN deployment models
- Encryption, hashing, and digital signatures and how they provide confidentiality, integrity, and nonrepudiation
- Methods, algorithms, and purposes of symmetric encryption
- Use and purpose of hashes and digital signatures in providing integrity and nonrepudiation
- Use and purpose of asymmetric encryption and Public Key Infrastructure (PKI)

An IP Security (IPsec) virtual private network (VPN) is an integral part of the security architecture of most organizations. It is used to connect branch offices, remote employees, and business partners to the resources of the organization. Providing confidentiality, integrity, and endpoint authentication, VPNs are ubiquitous and provide data loss prevention mechanisms for data in transit at multiple levels. From Secure Sockets Layer (SSL) VPNs to IPsec VPNs, site-to-site VPNs or remote access options, this security control is now embedded in networks and applications and should be made available in a transparent and manageable fashion. This chapter introduces the cryptographic elements used by VPNs, including symmetric and asymmetric algorithms, and describes the components, deployment options, and operational framework of VPN technologies.

VPN Overview

Historically, a VPN was an IP tunnel. Therefore, a generic routing encapsulation (GRE) tunnel is technically a VPN, even though GRE does not encrypt. Today, the use of a VPN implies the use of encryption. With a VPN, the information from a private network is transported over a public network, such as the Internet, to form a virtual network instead of using a dedicated Layer 2 connection, as shown in [Figure 12-1](#). To remain private, the traffic is encrypted to keep the data confidential. VPNs are *virtual* because they are not deployed over a physical private network end to end. At some point, they will traverse a public network of some sort.

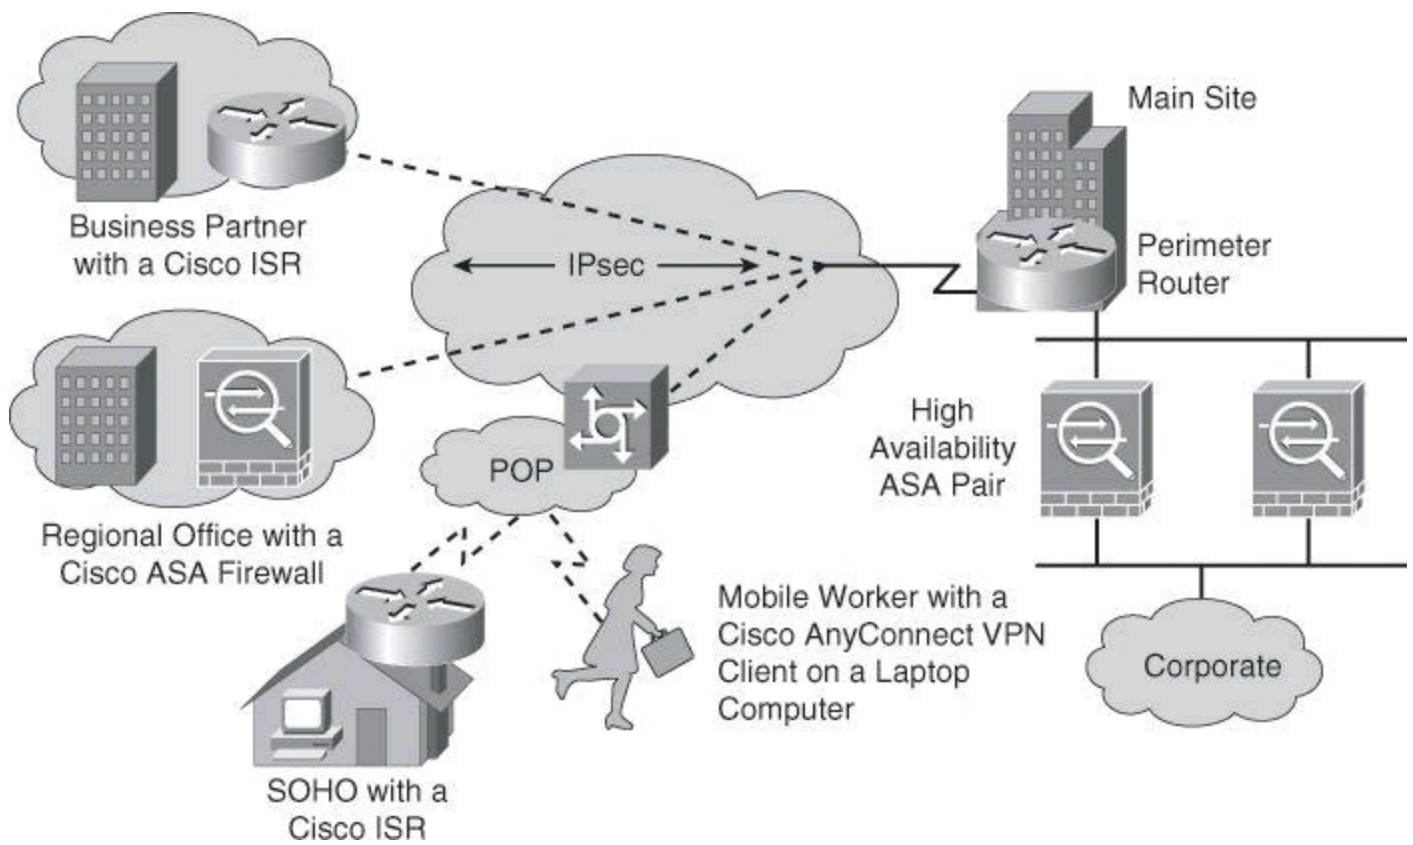


Figure 12-1. Where VPNs Are Found

VPNs are *private* because they use different methods of traffic segregation to keep the transferred traffic private. For instance, legacy Frame Relay networks could be considered VPNs in that they use a public network (the Frame Relay provider network), and they segregate traffic that shares the public network by encapsulating traffic within Frame Relay frames and tagging traffic that belongs to different customers with different virtual circuit identifiers. More recently, Multiprotocol Label Switching (MPLS) networks can also implement VPN-type connectivity because the customers share the provider's network (hence considered public) and their traffic is segregated using MPLS labels. It is worth mentioning that MPLS VPN does not provide confidentiality because the payload is not encrypted.

The term *private* can optionally refer to confidentiality. Even if traffic is segregated along the path of the VPN, malicious and nonmalicious attackers could technically intercept it at some point on the path and see its contents. To mitigate this risk, and to ensure traffic remains private, traffic is often encrypted to keep the data confidential. Encryption also provides other benefits, such as data integrity and authentication of the entities involved in the VPN. Encryption mechanisms are key in IPsec and SSL VPNs, the main subjects in this chapter, and will be explained in more detail throughout the rest of this book.

VPNs have many benefits:

- **Cost savings:** VPNs enable organizations to use cost-effective third-party Internet transport to connect remote offices and remote users to the main corporate site, thus eliminating expensive dedicated WAN links and modem banks. Furthermore, with the advent of cost-effective high-bandwidth technologies, such as digital subscriber line (DSL), organizations can use VPNs to reduce their connectivity costs while simultaneously increasing remote connection bandwidth.

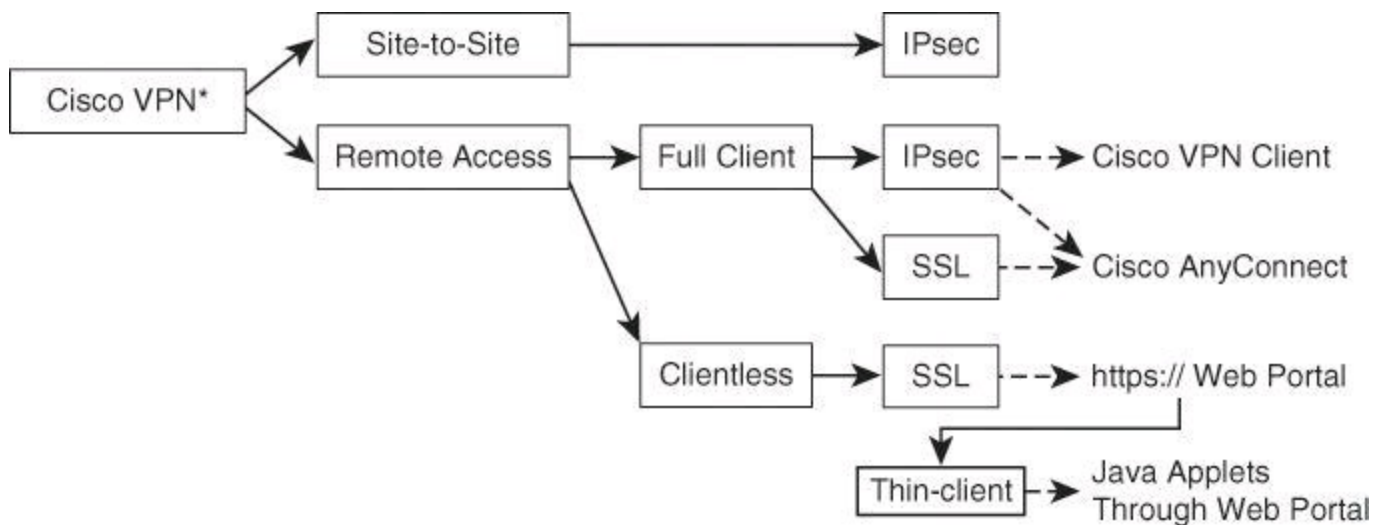
- **Scalability:** VPNs enable corporations to use the Internet infrastructure within ISPs and devices, which makes it easy to add new users. Therefore, corporations are able to add large amounts of capacity without adding significant infrastructure.
- **Compatibility with broadband technology:** VPNs allow mobile workers, telecommuters, and people who want to extend their workday to take advantage of high-speed, broadband connectivity, such as DSL and cable, to gain access to their corporate networks, providing workers significant flexibility and efficiency. Furthermore, high-speed broadband connections provide a cost-effective solution for connecting remote offices.
- **Security:** Optionally, VPNs provide the highest level of security by using advanced encryption and authentication protocols that protect data from unauthorized access.

VPN Types

There are different types of commercially deployed VPNs. VPN are classified according to the following criteria:

- **Based on deployment mode:** Site-to-site VPN and remote-access VPN
- **Based on Open Systems Interconnection (OSI) layer:** Layer 2 VPN (legacy protocols such as Frame Relay or ATM, and Layer 2 MPLS VPN), Layer 3 VPN (IPsec and MPLS Layer 3 VPN), and Layer 7 VPN (SSL VPN)
- **Based on underlying technology:** IPsec VPN, SSL VPN, MPLS VPN, other Layer 2 technologies such as Frame Relay or ATM, and hybrid VPNs combining multiple technologies

The two basic VPN deployment models typically use either IPsec or SSL technologies to keep the communications secured. The Cisco VPN solutions, illustrated in [Figure 12-2](#), and their associated technologies will be discussed throughout the rest of this book. Other VPN technologies, such as MPLS, are beyond the scope of this book.



*This figure shows the most commonly used Cisco VPN technologies and those covered in this book. Other technologies exist but they are beyond the mandate of this book.

Figure 12-2. Cisco VPN Solutions

Site-to-Site VPNs

A site-to-site VPN, shown in [Figure 12-3](#), is an extension of a classic WAN network. Site-to-site VPNs connect entire networks to each other; for example, they can connect a branch office network to a company headquarters network. In the past, a leased line or Frame Relay connection was required to connect sites, but because most corporations now have Internet access, these connections can be replaced with site-to-site VPNs.




Figure 12-3. Site-to-Site VPNs

In a site-to-site VPN, hosts do not have Cisco VPN Client software; they send and receive normal TCP/IP traffic through a VPN “gateway,” which could be a router, firewall, Cisco VPN Concentrator, or Cisco ASA 5500 Series Adaptive Security Appliance. The VPN gateway is responsible for encapsulating and encrypting outbound traffic for all the traffic from a particular site and sending it through a VPN tunnel over the Internet to a peer VPN gateway at the target site. Upon receipt, the peer VPN gateway strips the headers, decrypts the content, and relays the packet toward the target host inside its private network.

Remote-Access VPNs

In the past, corporations supported remote users by using dial-in networks and ISDN. With the advent of VPNs, a mobile user simply needs access to the Internet to communicate with the central office. In the case of telecommuters, their Internet connectivity is typically a broadband connection such as DSL or cable. Remote-access VPNs, shown in [Figure 12-4](#), can support the needs of telecommuters, mobile users, and extranet consumer-to-business traffic. Remote-access VPNs connect individual hosts who must access their company network securely over the Internet.

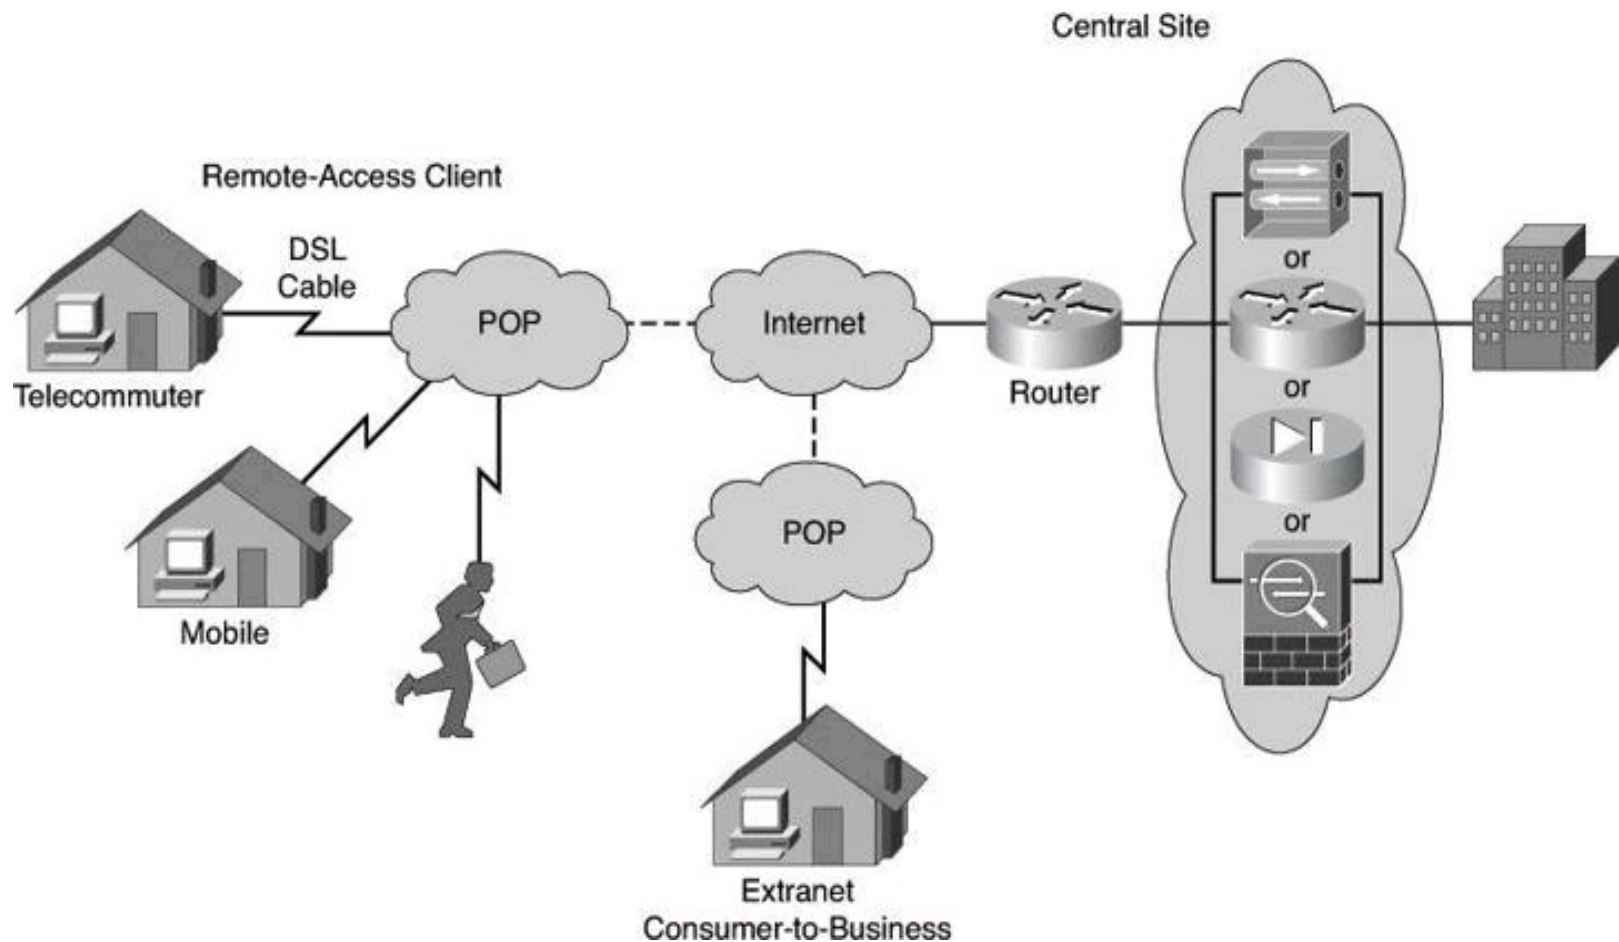


Figure 12-4. Remote-Access VPN

Hosts in a remote-access VPN commonly have a VPN client installed. As shown in [Figure 12-2](#), this is called “Full Client.” Cisco produces two clients: Cisco AnyConnect and Cisco VPN Client. These clients use a virtual network adapter (virtual network interface card) in the host computer to send traffic back to the head office. The VPN client software encapsulates and encrypts the traffic received on the virtual NIC before sending it over the Internet to the VPN gateway at the edge of the target network. Upon receipt, the VPN gateway behaves as it does for site-to-site VPNs.

Examining Cryptographic Services

Cryptographic services are the foundation for many security implementations. The key services provided by cryptography are as follows:

- **Confidentiality:** The assurance that no one can read a particular piece of data except the receivers explicitly intended.
- **Integrity or data authentication:** The assurance that data has not been altered in transit, intentionally or unintentionally.
- **Peer authentication:** The assurance that the other entity is who he, she, or it claims to be.
- **Nonrepudiation:** A proof of the integrity and origin of data. The sender can’t repudiate that he or she is the person who sent the data.
- **Key management:** The generation, exchange, storage, safeguarding, use, vetting, and replacement of keys.

VPN services use a combination of cryptographic technologies and algorithms to accomplish their

goals. In the router-to-router VPN, also known as site-to-site VPN, IP packets use symmetric encryption algorithms to encrypt the payload, with keys negotiated by key management protocols. They also use asymmetric encryption algorithms to create digital signatures and authenticate the VPN peers, and use hashing functions to provide checksum-type integrity checks.

This combination of algorithms and cryptographic methods, when used as a unit to negotiate the security settings to accomplish confidentiality, integrity, and authentication, is known as a cipher suite.

Data Authentication Versus User Authentication

Data authentication, also known as message authentication, is not the same as user authentication. With user authentication, as an example, an administrator is requested to provide his username and password to gain administrative access to a network device. Simply put, data authentication is a means for senders and receivers to validate the origin of each segment of a message that is exchanged. Data authentication also ensures data integrity, a way to prove that the payload wasn't tampered with during transmission. Data authentication is covered in greater detail in this chapter when we discuss cryptographic hashes.

Cryptology Overview

Cryptology is the science of making and breaking secret codes. Cryptology is broken into two separate disciplines: *Cryptography* is the development and use of codes, and *cryptanalysis* is the breaking of those codes. A symbiotic relationship exists between the two disciplines because each makes the other one better. National security organizations employ members of both disciplines and put them to work against each other.

In the past, there have been times when each discipline has been ahead of the other. For example, during the Hundred Years' War between France and England, the cryptanalysts were ahead of the cryptographers. France believed that its cipher was unbreakable; however, the Englishmen were able to break it. Some historians believe that World War II largely turned on the fact that the winning side on both fronts was much more successful than the losing side at cracking the encryption of its adversary.

It is an ironic fact of cryptography that it is impossible to prove that an algorithm is secure. You can prove only that it is not vulnerable to known cryptanalytic attacks. If there are methods that have been developed but are unknown to the cryptographer, an algorithm might be able to be cracked. You can prove only invulnerability to known attacks, except for a brute-force attack.

All algorithms are vulnerable to brute force. If every possible key is tried, one of the keys has to work. Therefore, no algorithm is unbreakable. The best you can hope for are algorithms that are vulnerable only to brute-force attacks.

Note

Two separate techniques can be used to try to achieve secure communication.

The first one is discussed in this chapter: cryptography, which is the science of encrypting a

message.

The second technique is called steganography, which pertains to the method used to hide a message. The message is hidden either within another message or by other means. For example, the German Embassy in Washington, D.C., used steganography by hiding a message within a message. The first letter of each word of a telegram would make up a secret message. This method is referred to as using a null cipher. Another example of steganography was by Histiaeus, a Greek tyrant, who had the head of his most trusted slave shaved, tattooed a message on the scalp, waited for the hair to grow back, and then dispatched the slave to deliver the secret message. Obviously, the element of urgency wasn't there! Another example of steganography is encoding a message in a JPG or BMP image. The "noise" that is induced in the picture by encoding the message is not noticeable to the naked eye.

By comparison, when using cryptography, others are aware that a secret message is being transmitted, but hopefully they can't decipher the message. With steganography, third parties don't know that a message is being transmitted.

If the topic of cryptography and steganography interests you, you must read this fascinating book: *The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography*, by Simon Singh (First Anchor Books Edition, 2000).

The History of Cryptography

The history of cryptography starts in diplomatic circles thousands of years ago. Messengers from a king's court would take encrypted messages to other courts. Occasionally, other courts not involved in the communication would attempt to steal any message sent to a kingdom they considered an adversary. Encryption was first used to prevent this information theft.

Not long after, military commanders started using encryption to secure messages. The messengers had far greater challenges than those of the diplomatic messengers because killing the messenger to get the message was common. With the stakes so high, military commanders resorted to encryption to secure their military communications.

The cipher attributed to Julius Caesar was a simple substitution cipher that he used on the battlefield to quickly encrypt a message that could easily be decrypted by his commanders. Thomas Jefferson, the third president of the United States, was also an inventor, and one of his inventions was an encryption system that he was believed to have used when serving as Secretary of State from 1790 to 1793.

Arthur Scherbius invented a machine in 1918 that served as a template for the machines that all the major participants in World War II used. He called the machine Enigma and sold it to Germany, estimating that if 1000 cryptanalysts tested four keys per minute, all day, every day, it would take 1.8 billion years to try them all.

During World War II, both the Germans and Allies had machines modeled after the Scherbius machine. These were the most sophisticated encryption devices ever developed, and in response, the British invented arguably the world's first computer, the Colossus, to break the encryption that was used by the Germans' Enigma machine.

Ciphers

A cipher is an algorithm for performing encryption and decryption. It is a series of well-defined steps that you can follow as a procedure. Substitution ciphers simply substitute one letter for another. In their simplest form, substitution ciphers retain the letter frequency of the original message. [Table 12-1](#) provides a summary of popular cipher categories.

Table 12-1. Cipher Categories

Cipher	Description
Substitution	Substitute one character for another, such as a = e, b = f, c = g, and so on.
Polyalphabetic	Based on substitution, using multiple substitution alphabets.
Transposition	Also known as permutation. Rather than replacing characters, characters are permuted or rearranged.
One-time pad	Also known as Vernam ciphers. Use keys to apply logical operations to plaintext.

Substitution Cipher

The cipher attributed to Julius Caesar is a simple substitution cipher. Every day has a different key, and that key is used to adjust the alphabet accordingly. For example, if today's key is five, an A is moved five spaces, resulting in an encoded message using F instead; a B is a G, a C is an H, and so forth. The next day the key might be eight, and the process begins again, so A is now I, B is J, and so on.

For example, if a message has 25 occurrences of the letter S, and S is replaced by the letter Q, there will be 25 occurrences of the letter Q. If the message is long enough, it will be vulnerable to frequency analysis because it retains the frequency patterns found in the language. Because of this weakness, polyalphabetic ciphers were invented.

Polyalphabetic Cipher

The Vigenère cipher, shown in [Figure 12-5](#), is a polyalphabetic cipher that encrypts text by using a series of different Caesar ciphers based on the letters of a keyword. It is a simple form of polyalphabetic substitution and therefore invulnerable to frequency analysis.

P L A I N T E X T L E T T E R

KEYWORD LETTER

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
A	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
B	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A
C	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B
D	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C
E	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D
F	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E
G	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F
H	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G
I	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H
J	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I
K	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J
L	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K
M	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L
N	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M
O	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N
P	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
Q	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
R	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
S	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
T	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
U	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
V	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
W	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V
X	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
Y	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X
Z	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y

Figure 12-5. The Vigenère Cipher

The method was originally described in 1553 in a book by Giovan Batista Belaso. Mistakenly, Blaise de Vigenère, a French cryptographer, has been attributed with its invention, hence its name.

To illustrate how it works, suppose that a key of SECRETKEY is used to encode ATTACK AT DAWN. The A (the first letter of ATTACK) is encoded by looking at the row starting with S (the first letter of the key SECRETKEY) for the letter in the A column. In this case, the A is replaced with S. Then you look for the row that begins with E for the letter T. This results in X as your second character. If you continue this encoding method, the message ATTACK AT DAWN is encrypted as SXVRGDKXBSAP.

Note

When using the Vigenère cipher and the message is longer than the key, just repeat the key.

Transposition Ciphers

In transposition ciphers, letters are simply rearranged rather than replaced. An example of this type of cipher is taking the message THE PACKAGE IS DELIVERED and transposing it to read DEREVILEDSIEGAKCAPEHT. In this example, the key is to reverse the letters.

Another example of a transposition cipher is the rail fence cipher, shown in [Figure 12-6](#). In this transposition, the words are spelled out as if they were a rail fence. The example in [Figure 12-6](#) uses a key of three to illustrate how this could be done.

Cipher Text:

T...O...I...O...L...T...E.H.C.V.R.S.L.W.F.E.A.O.C...E...E...B...N...E...N..

Rail Fence Cipher with a Key of Three:

```
T...O...I...O...L...T...E
.H.C.V.R.S.L.W.F.E.A.O.C.
..E...E...B...N...E...N..
```

In order to read the message, simply look diagonally, following the rail fence.

Clear Text = THE COVER IS BLOWN FLEE AT ONCE

Figure 12-6. Transposition Cipher: Rail Fence Cipher

The message THE COVER IS BLOWN FLEE AT ONCE would be encoded as TOIOLTEHCVRSLWFEOCEEENEN using this method of transposition. Once again, no letters were changed; they were just rearranged.

One-Time Pad Cipher

The one-time pad was invented and patented by Gilbert Vernam in 1917 while working at AT&T. The primary design of the one-time pad was meant to overcome the weaknesses of using the Vigenère cipher. Vernam’s idea was a stream cipher that would apply the exclusive OR (XOR) operation to plaintext with a key, but still using the Vigenère cipher. Joseph Maubourgne, a captain in the U.S. Army Signal Corps, contributed the idea of using random data as a key. This combined idea is so significant that the U.S. National Security Agency (NSA) has called this patent “perhaps the most important in the history of cryptography.”

[Figure 12-7](#) shows three one-time pads used in conjunction with the Vigenère cipher.

E	M	Z	F	O
Y	P	G	A	Q
R	N	C	V	X
T	L	B	I	H
J	S	U	D	K

N	S	D	F	Z
V	T	A	M	P
G	L	Y	E	X
B	R	W	U	I
Q	H	O	K	C

T	E	K	C	W
N	X	P	R	Z
I	B	Y	U	L
V	O	S	H	Q
G	M	F	J	A

Figure 12-7. One-Time Pads

There are several difficulties inherent in using one-time pads in the real world. The first of these is the challenge of creating random data. Computers, because they have a mathematical foundation, are incapable of creating random data. In addition, if the key is used more than once, it is trivial to break. Key distribution is also challenging (for example, how do you distribute the pads?).

Note

Rivest Cipher 4 (RC4) is an example of a Vernam cipher that is widely used on the Internet. It is not a one-time-pad because the key used is not random. If you have used Wired Equivalent Privacy (WEP) with a wireless network, you have used RC4.

In [Figure 12-8](#), the plaintext “at dawn attack the high plains” is encrypted with Sheet 2 of the one-time pads, using the Vigenère cipher.

Plain text: A T D A W N A T T A C K T H E H I G H P L A I N S
 key (sheet 2): N S D F Z V T A M P G L Y E X B R W U I Q H O K C
 Cipher text: N L G F V I T T F P I V R L B I Z C B X B H W X U

Using Vigenere below showing the three first letters substitutions:

		P L A I N T E X T L E T T E R																									
		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
A		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
B		B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A
C		C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B
D		D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C
E		E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D
K		F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E
E		G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F
Y		H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G
W		I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H
O		J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I
R		K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J
D		L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K
M		M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L
L		N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M
E		O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N
T		P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
T		Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
E		R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
R		S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
T		T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
U		U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
V		V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
W		W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V
X		X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
Y		Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X
Z		Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y

Figure 12-8. Encryption Using One-Time Pad

In [Figure 12-9](#), the ciphertext is decrypted using Sheet 2 of the one-time pads.

Cipher text: N L G F V I T T F P I V R L B I Z C B X B H W X U
 key (sheet 2): N S D F Z V T A M P G L Y E X B R W U I Q H O K C
 Plain text: A T D A W N A T T A C K T H E H I G H P L A I N S
 below, showing the deciphering of the first four letters

		P L A I N T E X T L E T T E R																									
		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
K E Y W O R D L E T T E R	A	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
	B	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A
	C	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B
	D	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C
	E	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D
	F	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E
	G	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F
	H	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G
	I	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H
	J	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I
	K	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J
	L	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K
	M	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L
	N	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M
	O	O	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N
	P	P	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	Q	Q	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
	R	R	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
	S	S	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
	T	T	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S
	U	U	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
	V	V	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
	W	W	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V
	X	X	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W
	Y	Y	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X
	Z	Z	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y

Figure 12-9. Decryption Using One-Time Pad

Classical ciphers are either transposition ciphers or substitution ciphers.



Recall that a transposition cipher exchanges the position of letters, whereas a substitution cipher replaces one letter with another letter based on a secret key.

[Figure 12-10](#) illustrates an example of a simple computer version of a substitution cipher, where the word WOLFE is translated in ASCII. The key MAJOR is also translated in ASCII. The key is applied to the cleartext using in this example the logical operation XOR. The result is ciphertext. The result of applying the key, using again the XOR operation, is the deciphering of the ciphertext to its original cleartext.

Clear Text = WOLFE

Key = MAJOR

Substitution Algorithm: XOR

Exclusive or = Logical Operation of **Exclusive Disjunction**

(When the elements are the same, the result is 0, and when the elements are different, the result is 1.)

WOLFE in ASCII:	01010111	01001111	01001100	01000110	01000101
MAJOR in ASCII:	01001101	01000001	01001010	01001111	01010010
Cipher Text:	00011010	00001110	00000110	00001001	00010111

Re-applying the Key On the Cipher Text to Find the Clear Text:

Cipher Text:	00011010	00001110	00000110	00001001	00010111
MAJOR in ASCII:	01001101	01000001	01001010	01001111	01010010
Clear Text ASCII:	01010111	01001111	01001100	01000110	01000101

The clear text ASCII found in the last operation corresponds to WOLFE in ASCII.

Figure 12-10. Computer Version of a Substitution Cipher

XOR is a type of logical disjunction on two operands that results in a value of true if only one of the operands has a value of true. In other words, for the result to be true (result of 1), the two operands must be different (one operand must be 0 and the other operand must be 1). In mathematics, operands are the input values used on the operation. With the formula $4 * 5$, 4 and 5 are operands, and $*$ (multiplication) is the operation.

[Figure 12-10](#) is an example of both a substitution cipher and symmetric encryption. Symmetric encryption is discussed later in this chapter.

Block and Stream Ciphers

Algorithms can operate in two modes:

- **Block mode:** The algorithm can work on only fixed chunks of data.
- **Stream mode:** The algorithm can process data bit by bit.

Block Ciphers

Block ciphers transform a fixed-length block of plaintext into a block of ciphertext. Applying the reverse transformation to the ciphertext block and using the same secret key results in decryption. Currently, the fixed length (also known as the block size) for many block ciphers is typically 128 bits. DES has a block size of 64 bits.

Note

Block size refers to how much data is encrypted at any one time, whereas *key length* refers to the size of the encryption key. For example, DES encrypts blocks in 64-bit chunks, including an 8-bit parity check, thus yielding a 56-bit effective key strength.

Block ciphers usually result in output data that is larger than the input data because the ciphertext must be a multiple of the block size. To accomplish this, block algorithms take data one chunk at a time (for example, 8 bytes) and use padding to add artificial data (blanks) if there is less input data than one full block.

The following are common block ciphers:

- DES and 3DES, running in either Electronic Code Book (ECB) mode or Cipher Block Chaining (CBC) mode
- Advanced Encryption Standard (AES)
- International Data Encryption Algorithm (IDEA)
- Secure and Fast Encryption Routine (SAFER)
- Skipjack
- Blowfish
- Rivest-Shamir-Alderman (RSA)

[Figure 12-11](#) illustrates the differences between ECB mode and CBC mode.

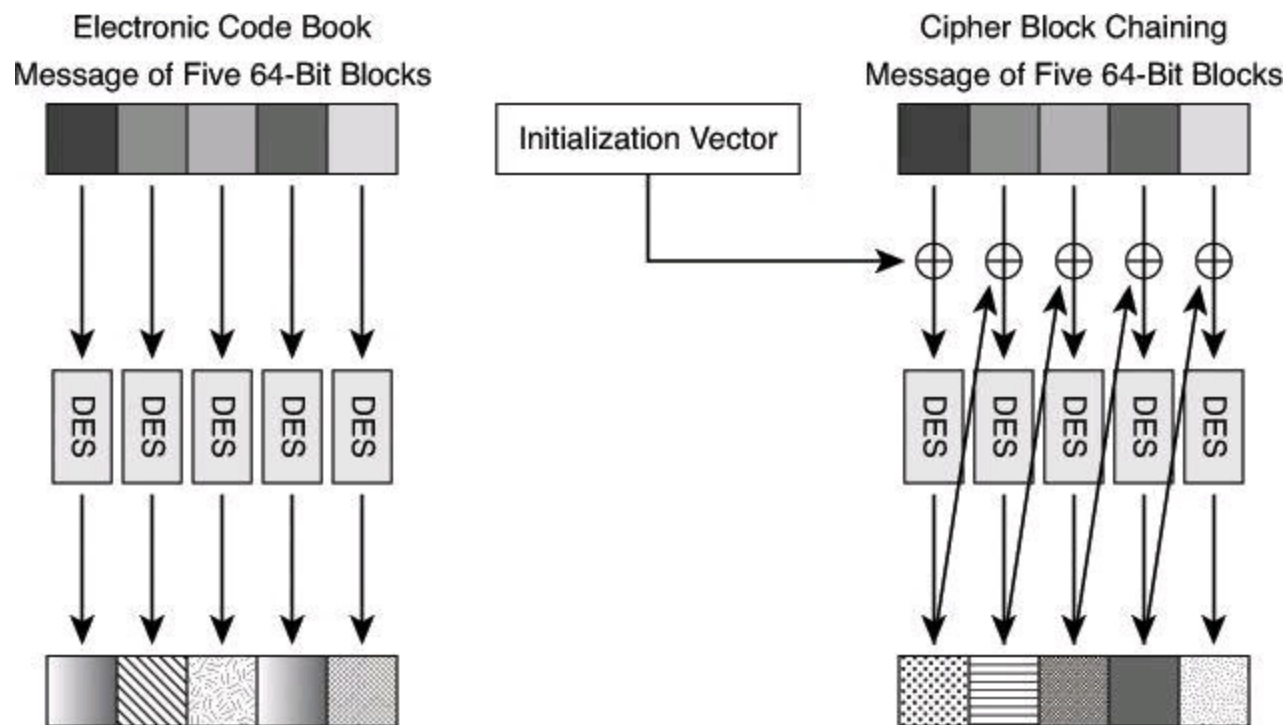


Figure 12-11. DES ECB Mode Versus DES CBC Mode

ECB mode serially encrypts each 64-bit plaintext block using the same 56-bit key. If two identical plaintext blocks are encrypted using the same key, their ciphertext blocks are the same. Therefore, an attacker could identify similar or identical traffic flowing through a communications channel and use that information to build a catalog of messages, which have a certain meaning, and replay them later without knowing their real meaning. For example, an attacker might capture a login sequence of someone with administrative privilege whose traffic is protected by ECB mode and then replay that sequence. That risk is undesirable, so CBC mode was invented to mitigate this risk.

In CBC mode, each 64-bit plaintext block is exclusive ORed (XORed) bitwise with the previous ciphertext block and then is encrypted using the DES key. Because of this process, the encryption of each block depends on previous blocks.

Encryption of the same 64-bit plaintext block can result in different ciphertext blocks. CBC mode can help guard against certain attacks, but it cannot help against sophisticated cryptanalysis or an extended brute-force attack.

Stream Ciphers

Unlike block ciphers, stream ciphers operate on smaller units of plaintext, typically bits. With a stream cipher, the transformation of these smaller plaintext units varies, depending on when they are encountered during the encryption process. Stream ciphers can be much faster than block ciphers, and generally do not increase the message size, because they can encrypt an arbitrary number of bits.

In stream cipher mode, the cipher uses previous ciphertext and the secret key to generate a pseudorandom stream of bits, which only the secret key can generate. To encrypt data, the data is XORed with the pseudorandom stream bit by bit, or sometimes byte by byte, to obtain the ciphertext. The decryption procedure is the same: The receiver generates the same random stream using the secret key, and XORs the ciphertext with the pseudorandom stream to obtain the plaintext.

Common stream ciphers include the following:

- DES and 3DES, running in output feedback (OFB) or cipher feedback (CFB) mode
- Rivest Cipher 4 (RC4)
- Software-optimized Encryption Algorithm (SEAL)

The Process of Encryption

Encryption is the process of disguising a message in such a way as to hide its original contents, as shown in [Figure 12-12](#). With encryption, the plaintext readable message is converted to ciphertext, which is the unreadable, “disguised” message. Decryption reverses the process. The purpose of encryption is to guarantee confidentiality so that only authorized entities can read the original message.

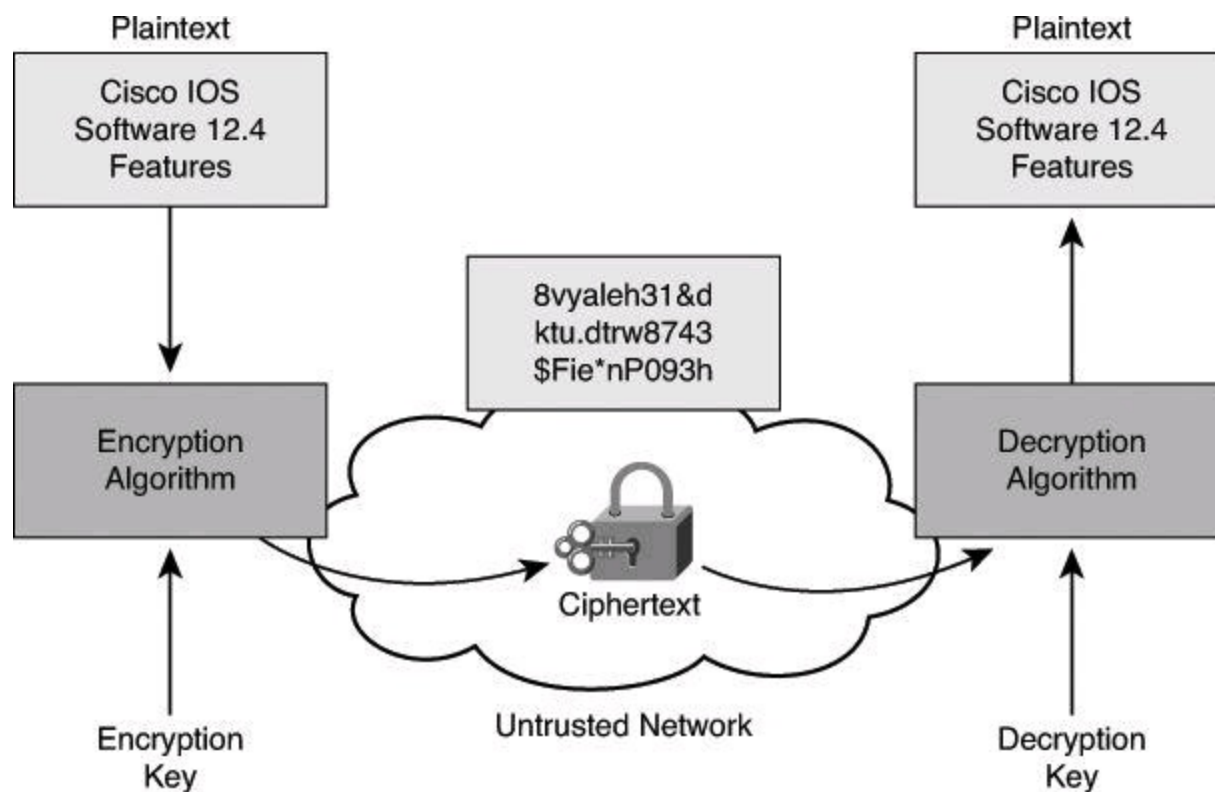


Figure 12-12. Transforming Plaintext into Ciphertext

Old encryption algorithms, such as the Caesar cipher or the Enigma machine, were based on the secrecy of the algorithm to achieve confidentiality. With modern technology, where reverse engineering is often simple, public-domain algorithms are often used. With most modern algorithms, successful decryption requires knowledge of the appropriate cryptographic keys; that is, the security of encryption lies in the secrecy of the keys, not the algorithm.

Encryption can provide confidentiality at an OSI layer, such as the following:

- Encrypt application layer data, such as secure email, secure database sessions (Oracle SQL*Net), and secure messaging (Lotus Notes sessions)
- Encrypt session layer data, using a protocol such as SSL or Transport Layer Security (TLS)
- Encrypt network layer data, using protocols such as those provided in the IPsec protocol suite
- Encrypt link layer data, using proprietary link-encrypting devices



Key
Topic

Nowadays, encryption algorithms such as Triple Data Encryption Standard (3DES) and Advanced Encryption Standard (AES) are readily distributed. So, because we all share the same algorithms, they have no need for protection (except for the fact that Western countries define cryptography as munitions, and therefore encryption algorithms are subject to the same export regulations as weapons).

Because we all share algorithms, what we protect are the cryptographic keys used with the algorithms. Cryptographic keys are sequences of bits that are input into an encryption algorithm together with the data to be encrypted.

Encryption Application Examples

The IPsec protocols can provide the encryption functionality for all the packets routed over an untrusted network. The encrypting IPsec peer takes a packet with the plaintext payload, encrypts the payload into ciphertext, and forwards the packet to the untrusted network (for example, the Internet). Its IPsec partner receives the ciphertext payload packet and decrypts the payload into the original plaintext. The two IPsec peers share the same encryption and decryption algorithm and proper keys.

The SSL protocol provides an encrypted channel on top of an existing TCP session. For example, HTTPS, which is HTTP encapsulated inside SSL, also known as HTTP over SSL or HTTP Secure, provides, among other services, confidentiality of the session between a web browser and a web server, using symmetric cryptography.

Both IPsec and SSL are used to set up a VPN. An IPsec VPN is application independent and requires a specialized IP stack on the end system or in the packet path that includes IPsec. Originally, SSL-based VPN supported only web-based applications with the SSL software included with all Internet browsers. Nowadays, SSL-based VPN also provides full tunnel support. Both SSL and IPsec are explained in more detail later in this book.

In contrast to IPsec and SSL, Layer 2 encryption, also known as data link layer encryption and depicted in [Figure 12-13](#), encrypts the whole frame, including the physical address fields located in the header of the frame, and therefore can be used only on point-to-point links where no network switching or routing equipment is required for path decision.

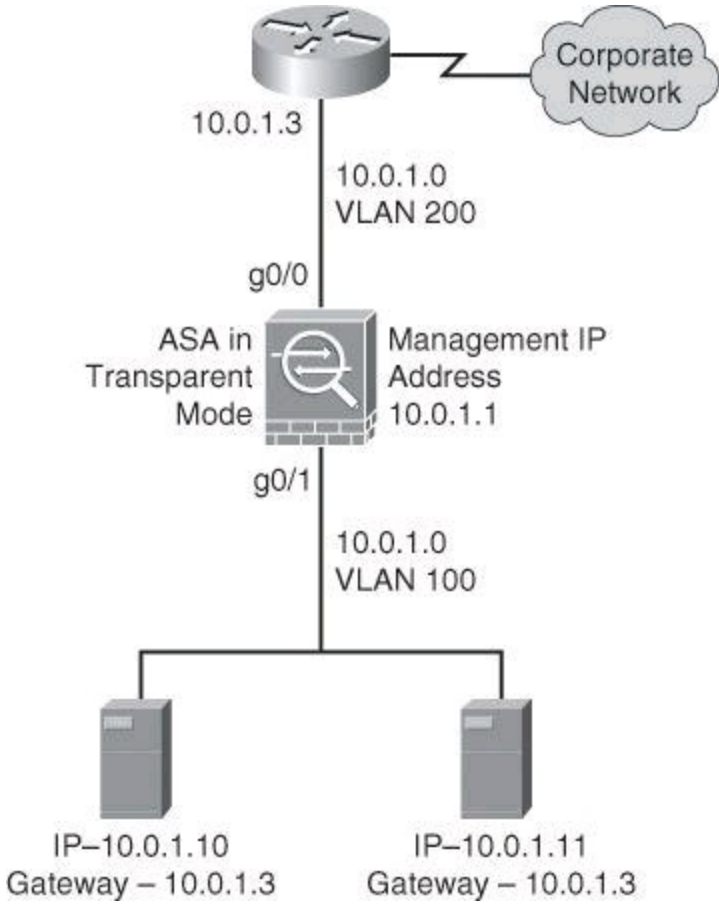


Figure 12-13. Cisco ASA Configured as a Layer 2 Firewall

Note

IPsec implementation is a mandatory part of IPv6.

Cryptanalysis

Cryptanalysis is the practice of breaking codes to obtain the meaning of encrypted data. An attacker who tries to break an algorithm or encrypted ciphertext might use one of the following attacks:

- Brute-force attack
- Ciphertext-only attack
- Known-plaintext (the usual brute-force) attack
- Chosen-plaintext attack
- Chosen-ciphertext attack
- Birthday attack
- Meet-in-the-middle attack

The sections that follow describe these attacks in greater detail.

Brute-Force Attack

In a brute-force attack, an attacker tries every possible key with the decryption algorithm, knowing that eventually one of them will work. All encryption algorithms are vulnerable to this attack. On average, a brute-force attack will succeed about 50 percent of the way through the keyspace (see “[Keyspaces](#),” later in this chapter, for more information). The objective of modern cryptographers is to have a sufficiently large keyspace so that it takes too much money and too much time to accomplish a brute-force attack.

Note

As reported by Distributed.net in 1999, a DES cracking machine was used to recover a 56-bit DES key in 22 hours using brute force. It is estimated that it would take 149 trillion years to crack AES using the same method.

Ciphertext-Only Attack

In a ciphertext-only attack, the attacker has the ciphertext of several messages, all of which have been encrypted using the same encryption algorithm, but the attacker has no knowledge of the underlying plaintext. The job of the attacker is to recover the ciphertext of as many messages as possible or, better yet, to deduce the key or keys used to encrypt the messages so as to decrypt other messages encrypted with the same keys. The attacker could use statistical analysis to achieve the result. Those attacks are no longer practical today because modern algorithms produce pseudorandom output that is resistant to statistical analysis.

Known-Plaintext Attack

In a known-plaintext attack, the attacker has access to the ciphertext of several messages, but also knows something about the plaintext underlying that ciphertext. With knowledge of the underlying protocol (HTTP, FTP, SMTP), file type (web page, file transferred, email message), and some characteristic strings that might appear in the plaintext, the attacker uses a brute-force attack to try keys until decryption with the correct key produces a meaningful result. This attack may be the most practical attack, because attackers can usually assume the type and some features of the underlying plaintext, if they can only capture ciphertext. However, modern algorithms with enormous keyspace make it unlikely for this attack to succeed, because on average an attacker would have to search through at least half of the keyspace to be successful.

Note

A known-plaintext attack was used in the cracking of the Enigma machine. During World War II, Alan Turing, a famous British crypto-mathematician, discovered that at around 6 a.m. each day the Germans were sending an encrypted weather report. Turing was sure that within the ciphertext captured around that time of day the word WETTER (weather in German) could be found. With the ciphertext equivalent to the plaintext WEATHER, Turing had a good start to continue reverse engineering the rest of the message.

In cryptanalysis, a sample of ciphertext suspected to be a resulting plaintext is called a *crib*.

Chosen-Plaintext Attack

In a chosen-plaintext attack, the attacker chooses what data the encryption device encrypts and observes the ciphertext output. A chosen-plaintext attack is more powerful than a known-plaintext attack because the attacker gets to choose the plaintext blocks to encrypt, allowing the attacker to choose plaintext that might yield more information about the key. This attack might not be very practical, because it is often difficult or impossible to capture both the ciphertext and plaintext, unless the trusted network has been broken into and the attacker already has access to confidential information.

Chosen-Ciphertext Attack

In a chosen-ciphertext attack, the attacker can choose different ciphertexts to be decrypted and has access to the decrypted plaintext. With the pair, the attacker can search through the keyspace and determine which key decrypts the chosen ciphertext in the captured plaintext. For example, the attacker has access to a tamperproof encryption device with an embedded key. His job is to deduce the embedded key by sending data through the box. This attack is analogous to the chosen-plaintext attack. This attack might not be very practical, because it is often difficult or impossible to capture both the ciphertext and plaintext, unless the trusted network has been broken into and the attacker already has access to confidential information.

Birthday Attack

The birthday attack gets its name from the amazing statistical probability involved in two individuals having the same birthday. According to statisticians, the probability that 2 people in a group of 23 people share the same birthday is greater than 50 percent.

This particular attack is a form of brute-force attack against hash functions. If some function, when supplied with a random input, returns one of k equally likely values, then by repeating the function with different inputs, the same output is expected after $1.2k^{1/2}$ number of times.

Tip

To test the birthday theory, input 365 in the place of k .

Meet-in-the-Middle

The meet-in-the-middle attack is a known-plaintext attack. Do not confuse this with the *man-in-the-middle attack*, which is discussed later. In a meet-in-the-middle attack, the attacker knows a portion of the plaintext and the corresponding ciphertext. The plaintext is encrypted with every possible key, and the results are stored. The ciphertext is then decrypted using every key until one of the results matches one of the stored values.

Desirable Encryption Algorithm Features

The following are features that a good encryption algorithm provides:

- Resists cryptographic attacks
- Supports variable and long key lengths and scalability
- Creates an avalanche effect

Note

When deciding on an algorithm for your organization, keep in mind the export regulations of cryptographic equipment, which might limit your choice.

A good cryptographic algorithm is designed in such a way that it resists common cryptographic attacks. The best way to break data protected by the algorithm is to try to decrypt the data using all the possible keys. The length of time that such an attack requires to succeed depends on the number of possible keys, but is generally very, very long. With appropriately long keys, such attacks are usually considered infeasible.

Variable key lengths and scalability are also desirable attributes of a good encryption algorithm. The longer the encryption key, the longer it takes an attacker to break it. For example, a 16-bit key would mean that there are 65,536 possible keys, but a 56-bit key would mean there are $7.2 * 10^{16}$ possible keys. Scalability provides flexible key lengths and allows you to select the strength and speed of encryption that you need.

When changing only a few bits of the plaintext message causes its ciphertext to change completely, this is known as an *avalanche effect*. The avalanche effect is a desired feature of an encryption algorithm because it allows very similar messages to be sent over an untrusted medium, with the encrypted (ciphertext) messages being completely different.

You must carefully consider export and import restrictions when you use encryption internationally. Some countries do not allow the export of encryption algorithms, or allow only the export of these algorithms with shorter keys, and some countries impose import restrictions on cryptographic algorithms.

In January 2000, and again in 2010, the restrictions that the U.S. Department of Commerce placed on export regulations were dramatically relaxed. As an example of the loosening of the rules, cryptographic products were made exportable under a license exception unless the end users were governments outside the United States or under embargo. For more information on the current U.S. Department of Commerce export regulations, see <http://www.commerce.gov>.

Key Management

Key management is often considered the most difficult part of designing a cryptosystem. Many cryptosystems have failed because of mistakes in their key management, and all modern cryptographic algorithms require the services of key management procedures. In practice, most attacks on cryptographic systems will be aimed at the key management level rather than at the cryptographic algorithm itself.

Key Management Components

Key management consists of the following components:

- **Key generation:** In a modern cryptographic system, key generation is usually automated and not left to the end user. The use of good random number generators is needed to ensure that all keys are likely to be equally generated so that the attacker cannot predict which keys are more likely to be used.

- **Key verification:** Almost all cryptographic algorithms have some weak keys that should not be used, and with the help of key verification procedures, you can regenerate these keys if they occur.
- **Key storage:** On a modern multiuser operating system that uses cryptography, a key can be stored in memory. This presents a possible problem when that memory is swapped to the disk, because a Trojan horse program, installed on the PC of a user, could then have access to the private keys of that user. A possible solution is to store the key on a USB stick and require a password to unlock that key.
- **Key exchange:** The key management procedures should also provide a secure key exchange mechanism, which allows secure agreement on the keying material with the other party, probably over an untrusted medium.
- **Key revocation and destruction:** Key revocation notifies all the interested parties that a certain key has been compromised and should no longer be used. Key destruction erases old keys in such a manner that malicious attackers cannot recover them.

Key
Topic

Key management deals with the secure generation, verification, exchange, storage, revocation, and destruction of keys.

Keyspaces

The keyspace of an algorithm is the set of all possible key values. A key that has n bits produces a keyspace that has 2^n possible key values. By adding 1 bit to the key, you effectively double the keyspace. For example, DES with its 56-bit keys has a keyspace of more than 72,000,000,000,000,000 (2^{56}) possible keys, but by adding 1 bit to the key length, the keyspace doubles, and an attacker needs twice the amount of time to search the keyspace.

As previously mentioned, almost every algorithm has some weak keys in its keyspace that enable an attacker to break the encryption via a shortcut. Weak keys show regularities in encryption or poor encryption. For instance, DES has four keys for which encryption is the same as decryption. This means that if one of these weak keys is encrypted twice, the original plaintext is recovered.

The weak keys of DES are those that produce 16 identical subkeys. This occurs when the key bits are

- Alternating 1s + 0s (0101010101010101)
- Alternating F + E (FEFEFEFEFEFEFEFE)
- E0E0E0E0F1F1F1F1
- 1F1F1F1F0E0E0E0E

It is unlikely that such keys would be chosen, but implementations should still verify all keys and prevent weak keys from being used. With manual key generation, you must take special care to avoid defining weak keys.

Key Length Issues

If the cryptographic system is trustworthy, the only way to break it is with a brute-force attack. A brute-force attack is a search through the entire keyspace, trying all the possible keys, to find a key that decrypts the data. If the keyspace is large enough, the search should require an enormous amount of time, making such an exhaustive effort infeasible. On average, an attacker has to search through half of the keyspace before the correct key is found. The time that is needed to accomplish this search depends on the computer power available to the attacker. However, current key lengths can easily make any attempt insignificant, because it would take many millions or billions of years to complete the search when a sufficiently long key is used.

With modern algorithms that are trusted, the strength of protection depends solely on the length of the key. Choose the key length so that it protects data confidentiality or integrity for an adequate period of time. Data that is more sensitive and needs to be kept secret longer must use longer keys.

The funding of the attacker should also affect your choice of key length. When you assess the risk of someone breaking the encryption algorithm, you must estimate the resources of the attacker and how long you must protect the data. For example, if the attacker has \$1 million of funding, and the data must be protected for one year, classic DES is not a good choice because it can be broken by a \$1 million machine in a couple of minutes. However, it would take an attacker some million years or more to crack 168-bit 3DES or 128-bit RC4, which makes either of these key length choices more than adequate.

Performance is another issue that can influence the choice of key length. You must find a good balance between the speed and protection strength of an algorithm, because some algorithms, such as RSA, run slower with larger key sizes. Strive for adequate protection while also enabling unhindered communication over untrusted networks.

Because of the rapid advances in technology and cryptanalytic methods, the key size needed for a particular application is constantly increasing. Go to the BlueKrypt website at <http://www.keylength.com/en/4/> to see updated NIST key length recommendations.

Example of the Impact of Key Length

Part of the strength of the RSA algorithm is the difficulty of factoring large numbers. If a 1024-bit number is hard to factor, then a 2048-bit number is going to be even harder to factor. Even with the fastest computers available today, it would take many lifetimes to factor a 1024-bit number that is a factor of two 512-bit prime numbers. Of course, this advantage is lost if an easy way to factor large numbers is found. However, cryptographers consider this possibility unlikely, and the rule “the longer the key, the better” is valid, except for possible performance reasons.

As of 2005, the best known attack on 3DES required around 2^{32} known plaintexts, 2^{113} steps, 2^{90} single DES encryptions, and 2^{88} memory operations. This is not currently practical. If the attacker seeks to discover any one of many cryptographic keys, there is a memory-efficient attack that will discover one of 2^{28} keys, given a handful of chosen plaintexts per key and around 2^{84} encryption operations. This attack is highly parallelizable and verges on the practical, given billion-dollar budgets and years to mount the attack, although the circumstances in which it would be useful are limited.

Symmetric and Asymmetric Encryption Overview

An encryption algorithm, which is also called a cipher, is a mathematical function that is used to encrypt and decrypt data. Generally, there are two functions, one to encrypt and one to decrypt. If the security of an encryption system is based on the secrecy of the algorithm itself, the algorithm code must be heavily guarded. If the algorithm is revealed, every party that is involved must change the algorithm.

Modern cryptography takes a different approach: all algorithms are public, and cryptographic keys are used to ensure the secrecy of data. There are two classes of encryption algorithms, which differ in their use of keys:

- **Symmetric encryption algorithms:** Use the same key to encrypt and decrypt data
- **Asymmetric encryption algorithms:** Use different keys to encrypt and decrypt data



A *key* is a required parameter for encryption algorithms. There are two concepts regarding keys:

- **Symmetric encryption:** The same key encrypts and decrypts.
- **Asymmetric encryption:** One key encrypts, and a different key decrypts.

Symmetric Encryption Algorithms

Symmetric encryption algorithms are algorithms in which the encryption and decryption keys are the same, as shown in [Figure 12-14](#). Therefore, the sender and the receiver must share the same secret key before communicating securely. The security of a symmetric algorithm rests in the secrecy of the symmetric key; by obtaining the key, anyone can encrypt and decrypt messages. Symmetric encryption is often called secret-key encryption. Symmetric, or secret-key, encryption is the more traditional form of cryptography. The typical key length range of symmetric encryption algorithms is 40 to 256 bits. Best practice considers that anything less than 128-bit key material is not considered strong enough for business use.

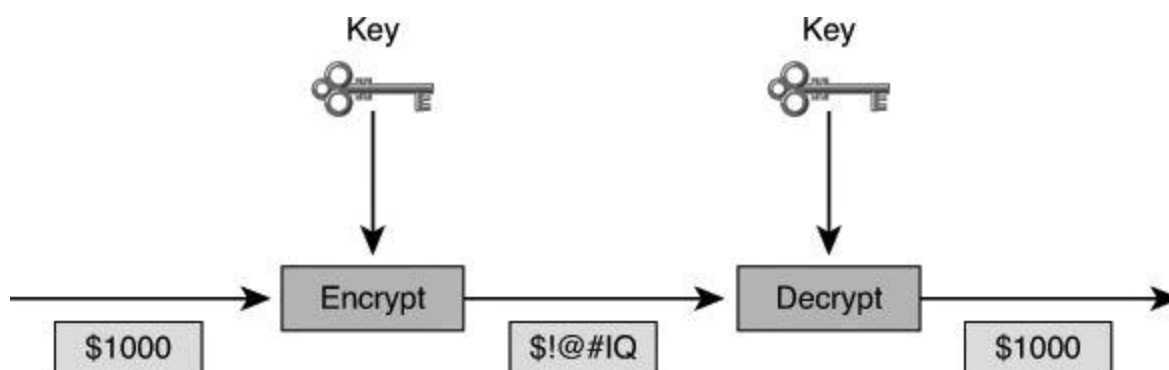


Figure 12-14. Symmetric Encryption at Work

The following are well-known encryption algorithms that use symmetric keys:

- **DES:** 56-bit keys
- **Triple DES (3DES):** 112- and 168-bit keys
- **AES:** 128-, 192-, and 256-bit keys
- **IDEA:** 128-bit keys

- **The RC series (RC2, RC4, RC5, and RC6):**
 - **RC2:** 40- and 64-bit keys
 - **RC4:** 1- to 256-bit keys
 - **RC5:** 0- to 2040-bit keys
 - **RC6:** 128-, 192-, and 256-bit keys
- **Blowfish:** 32- to 448-bit keys

The most commonly used techniques in symmetric encryption cryptography are block ciphers, stream ciphers, and message authentication codes (MAC).

Because symmetric algorithms are usually quite fast, they are often used for wire-speed encryption in data networks. Symmetric algorithms are based on simple mathematical operations and can easily be accelerated by hardware. Because of their speed, you can use symmetric algorithms for bulk encryption when data privacy is required (for example, to protect a VPN).

On the other hand, key management can be a challenge. The communicating parties must exchange the symmetric, secret key using a secure channel before any encryption can occur. Therefore, the security of any cryptographic system depends greatly on the security of the key exchange method.

Because of their speed, symmetric algorithms are frequently used for encryption services, with additional key management algorithms providing secure key exchange.

Some of the characteristics of symmetric algorithms are as follows:

- Faster than asymmetric algorithms.
- Much stronger than asymmetric algorithms.
- Much shorter key lengths than asymmetric algorithms.
- Simpler mathematics than asymmetric algorithms.
- One key is used for both encryption and decryption.
- Sometimes referred to as private-key encryption.

Modern symmetric algorithms use key lengths that range from 40 to 256 bits. This range gives symmetric algorithms keyspaces that range from 240 (1,099,511,627,776 possible keys) to 2^{256} ($1.5 * 10^{77}$) possible keys. (Keyspaces are described in depth later in this chapter.) This large range is the difference between whether or not the algorithm is vulnerable to a brute-force attack. If you use a key length of 40 bits, your encryption is likely to be broken easily using a brute-force attack. In contrast, if your key length is 256 bits, it is unlikely that a brute-force attack will be successful, because the keyspace is too large.

On average, a brute-force attack will succeed halfway through the keyspace. Key lengths that are too short can have the entire possible keyspace stored in RAM on a server cluster of a cracker, which makes it possible for the algorithm to be cracked in real time.

[Table 12-2](#) illustrates ongoing expectations for valid key lengths, assuming that the algorithms are mathematically and cryptographically sound. What is also assumed in such calculations is that computing power will continue to grow at its present rate and the ability to perform brute-force attacks will grow at the same rate.

Table 12-2. Acceptable Key Lengths in Bits

	Symmetric Key	Asymmetric Key	Digital Signature	Hash
Protection up to 3 years	80	1248	160	160
Protection up to 10 years	96	1776	192	192
Protection up to 20 years	112	2432	224	224
Protection up to 30 years	128	3248	256	256

Caution

If a method other than brute-force is discovered against a particular algorithm, the key lengths in [Table 12-2](#) become obsolete.

Note

Note the comparatively short symmetric key lengths, illustrating that symmetric algorithms are the strongest type of algorithm.

Comparing Symmetric Encryption Algorithms

Symmetric encryption algorithms operate under the same framework, but they present considerable differences. [Table 12-3](#) provides a summary analysis of their respective characteristics.

Table 12-3. Characteristics of Symmetric Encryption Algorithms

Algorithm	Description
DES	Block cipher, encrypts 64-bit data blocks. Fixed-length 64-bit key (only 56 bits used for encryption).
3DES	Applies DES three times in a row using three different keys. Key sizes of 168 and 112 bits.
AES	Iterated block cipher with variable block and key lengths. 128- or 192- or 256-bit keys.
Rivest ciphers	Widely deployed family of algorithms. Variable block and key sizes.

DES

Data Encryption Standard (DES) is a symmetric encryption algorithm, as shown in [Figure 12-14](#). It usually operates in block mode, in which it encrypts data in 64-bit blocks.

The DES algorithm is essentially a sequence of permutations and substitutions of data bits,

combined with the encryption key. The same algorithm and key are used for both encryption and decryption. Cryptography researchers have scrutinized DES for nearly 35 years and have found no significant flaws.

Because DES is based on simple mathematical functions, it can easily be implemented and accelerated in hardware.

DES has a fixed key length. The key is actually 64 bits long, but only 56 bits are used for encryption, with the remaining 8 bits used for parity; the least significant bit of each key byte is used to indicate odd parity.

A DES key is always 56 bits long. When you use DES with a weaker encryption of a 40-bit key, it actually means that the encryption key is 40 secret bits and 16 known bits, which makes the key length 56 bits. In this case, DES actually has a key strength of 40 bits.

DES Modes of Operation

To encrypt or decrypt more than 64 bits of data, DES uses two different types of ciphers:

- **Block ciphers:** Operate on fixed-length groups of bits, termed blocks, with an unvarying transformation
- **Stream ciphers:** Operate on individual digits one at a time with the transformation varying during the encryption

DES Security Guidelines

You should consider doing several things to protect the security of DES-encrypted data:

- Change keys frequently to help prevent brute-force attacks.
- Use a secure channel to communicate the DES key from the sender to the receiver.
- Consider using DES in CBC mode. With CBC, the encryption of each 64-bit block depends on previous blocks. CBC is the most widely used mode of DES.
- Verify that a key is not part of the weak or semi-weak key list before using it. DES has 4 weak keys and 12 semi-weak keys. Because there are 2^{56} possible DES keys, the chance of picking one of these keys is small. However, because testing the key has no significant impact on the encryption time, it is recommended that you test the key. The keys that should be avoided are listed in Section 3.6 of Publication 74 of the Federal Information Processing Standards, found at <http://www.itl.nist.gov/fipspubs/fip74.htm>.

Note

If possible, use 3DES rather than DES. You should use DES only for very short-term confidentiality.

3DES

With advances in computer processing power, the original 56-bit DES key became too short to withstand even medium-budget attackers. One way to increase the DES effective key length, without changing the well-analyzed algorithm itself, is to use the same algorithm with different keys several

times in a row.

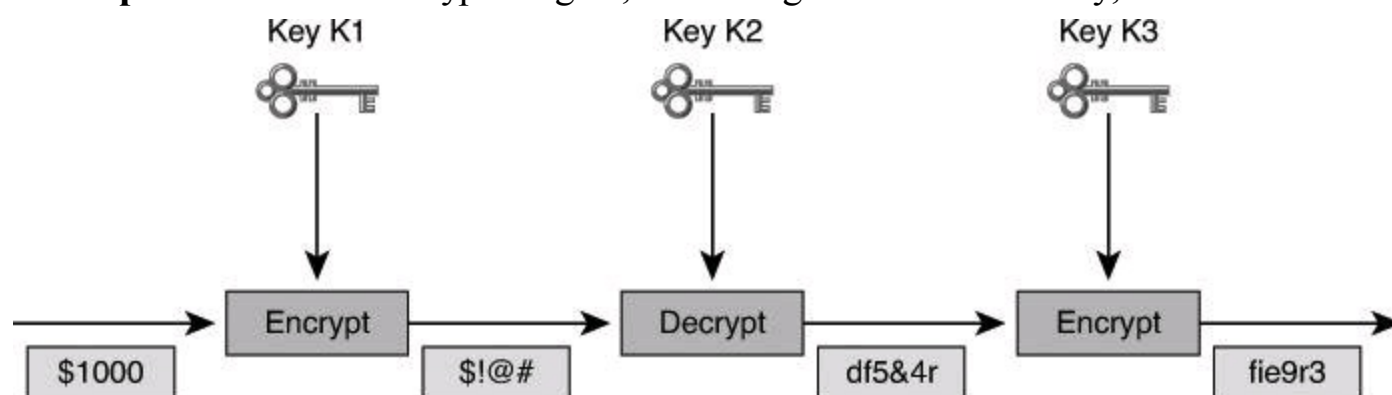
The technique of applying DES three times in a row to a plaintext block is called 3DES. Brute-force attacks on 3DES are considered infeasible today, and because the basic algorithm has been well tested in the field for more than 35 years, it is considered very trustworthy.

3DES uses a method called 3DES-Encrypt-Decrypt-Encrypt (3DES-EDE) to encrypt plaintext, as shown in [Figure 12-15](#). 3DES-EDE includes the following steps:

Step 1. The message is encrypted using the first 56-bit key, known as K1.

Step 2. The data is decrypted using the second 56-bit key, known as K2.

Step 3. The data is encrypted again, now using the third 56-bit key, known as K3.



• EDE (Encrypt-Decrypt-Encrypt) Method – 3DES-EDE Method:

- Data is encrypted using K1.
- Data is decrypted using K2.
- Data is encrypted using K3.
- If K1 = K3, Key Yields 112-Bit Key Length
- If K1 ≠ K3, Key Yields 168-Bit Key Length

Figure 12-15. 3DES Encryption Process

The 3DES-EDE procedure provides encryption with an effective key length of 168 bits. If keys K1 and K3 are equal, as in some implementations, a less-secure encryption of 112 bits is achieved.

The following procedure is used to decrypt a 3DES-EDE block:

Step 1. Decrypt the ciphertext using key K3.

Step 2. Encrypt the data using key K2.

Step 3. Decrypt the data using key K1.

Note

Encrypting the data three times with three different keys does not significantly increase security. The 3DES-EDE method must be used. For example, it can be shown that encrypting data three times in a row using different 56-bit keys equals an effective 58-bit key strength and not the full 168-bit key strength, as you would expect.

Also, here's why 3DES can be either 56-, 112-, or 168-bit strong: 3DES keying option 1, which provides 168-bit strength, uses three independent 56-bit values for K1, K2, and K3 (3×56 bits). With keying option 2, K1 and K3 use the same value and K2 is a different

value, so two different 56-bit keys produce a 112-bit strength (2×56). Keying option 1 uses only one 56-bit value for K1, K2, and K3, offering only 56-bit strength.

AES

For a number of years, it was recognized that DES would eventually reach the end of its usefulness. In 1997, the AES initiative was announced, and the public was invited to propose candidate encryption schemes, one of which could be chosen as the encryption standard to replace DES. There were rigorous reviews of 15 original candidates. Rijndael, Twofish, and RC6 were among the finalists. Rijndael was ultimately selected.

The Rijndael Cipher

On October 2, 2000, the U.S. National Institute of Standards and Technology (NIST) announced the selection of the Rijndael cipher as the AES algorithm. The Rijndael cipher, developed by Joan Daemen and Vincent Rijmen, has a variable block length and key length. The algorithm currently specifies how to use keys with a length of 128, 192, or 256 bits to encrypt blocks with a length of 128, 192, or 256 bits, which provides nine different combinations of key length and block length. Both block length and key length can be extended very easily in multiples of 32 bits.

The U.S. Secretary of Commerce approved the adoption of AES as an official U.S. government standard, effective May 26, 2002. AES is listed in Annex A of FIPS Publication 140-2 as an approved security function.

Note

Go to <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf> for more information about AES.

Rijndael is an iterated block cipher, which means that the initial input block and cipher key undergo multiple transformation cycles before producing output. The algorithm can operate over a variable-length block using variable-length keys; a 128-, 192-, or 256-bit key can be used to encrypt data blocks that are 128, 192, or 256 bits long, and all nine combinations of key and block length are possible. The accepted AES implementation of Rijndael contains only some of the capabilities of the Rijndael algorithm. The algorithm is written so that the block length or the key length or both can easily be extended in multiples of 32 bits, and the system is specifically designed for efficient implementation in hardware or software on a range of processors.

AES Versus 3DES

AES was chosen to replace DES and 3DES because the key length of AES makes it stronger than DES and AES runs faster than 3DES on comparable hardware. AES is more efficient than DES and 3DES on comparable hardware, usually by a factor of five when it is compared with DES. Also, AES is more suitable for high-throughput, low-latency environments, especially if pure software encryption is used. However, AES is a relatively young algorithm, and, as the golden rule of cryptography states, a more mature algorithm is always more trusted. 3DES is therefore a more conservative and more trusted choice in terms of strength, because it has been analyzed for around 35 years.

Rivest Ciphers

The RC family of algorithms is widely deployed in many networking applications because of their favorable speed and variable key length capabilities.

The RC algorithms were designed all or in part by Ronald Rivest. Some of the most widely used RC algorithms are as follows:

- **RC2:** A variable-key-size block cipher that was designed as a “drop-in” replacement for DES.
- **RC4:** A variable-key-size Vernam stream cipher often used in file encryption products and for secure communications, such as within SSL. It is not considered a one-time pad because its key is not random. The cipher can be expected to run very quickly in software and is considered secure, although it can be implemented insecurely, as in WEP.
- **RC5:** A fast block cipher that has variable block size and variable key size. With a 64-bit block size, RC5 can be used as a drop-in replacement for DES.
- **RC6:** A block cipher that was designed by Rivest, Sidney, and Yin and is based on RC5. Its main design goal was to meet the requirement of AES.

SEAL

Although it is beyond the scope of this book, it is good to know that Cisco also provides support for SEAL encryption. The Software-optimized Encryption Algorithm is an alternative algorithm to software-based DES, 3DES, and AES. SEAL encryption uses a 160-bit encryption key and has a lower impact on the CPU compared to other software-based algorithms. The SEAL encryption feature provides support for the SEAL algorithm in Cisco IOS IPsec implementations. SEAL support was added to Cisco IOS Software Release 12.3(7)T.

Asymmetric Encryption Algorithms

Asymmetric algorithms, also sometimes called *public-key algorithms*, are designed in such a way that the key used for encryption differs from the key used for decryption, as shown in [Figure 12-16](#).

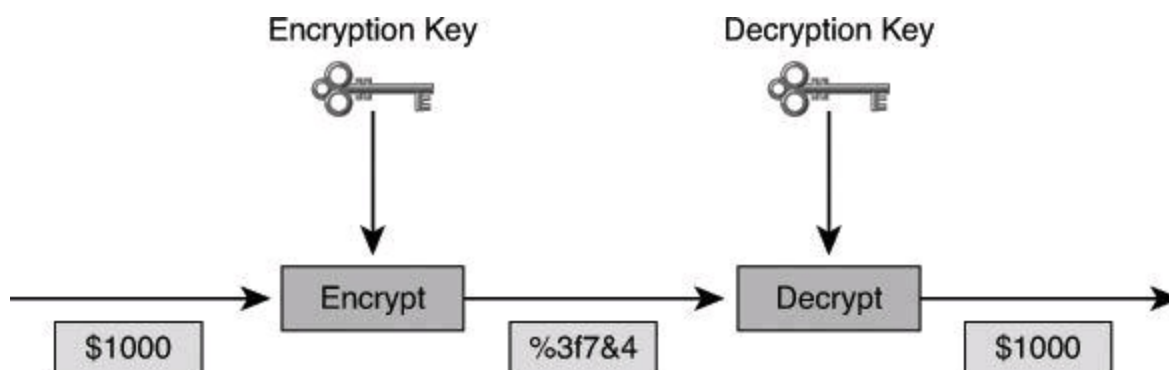


Figure 12-16. Asymmetric Encryption at Work

The decryption key cannot, in any reasonable amount of time, be calculated from the encryption key, and vice versa. The typical key length range for asymmetric algorithms is 512 to 4096 bits. You cannot directly compare the key length of asymmetric and symmetric algorithms because the

underlying design of the two algorithm families differs greatly.

Note

For further reference, consult the book *Applied Cryptography*, by Bruce Schneier. Mr. Schneier also maintains an informative and entertaining blog/newsletter, which you can subscribe to at <http://www.schneier.com/crypto-gram.html>.

The best-known asymmetric cryptographic algorithms are RSA, ElGamal, Digital Signature Algorithm (DSA), and elliptic curve algorithms.

Note

Rivest, Shamir, and Aldeman, who met while at the Massachusetts Institute of Technology (MIT), released the RSA algorithm in 1977.

Asymmetric algorithms are substantially slower than symmetric algorithms. Their design is based on computational problems, such as factoring extremely large numbers or computing discrete logarithms of extremely large numbers. Because they lack speed, asymmetric algorithms are typically used in low-volume cryptographic mechanisms, such as digital signatures and key exchange. However, the key management of asymmetric algorithms tends to be simpler than that of symmetric algorithms, because usually one of the two encryption or decryption keys can be made public.



Because symmetric ciphers are faster than asymmetric algorithms, they are used for bulk data encryption.

Asymmetric algorithms are explained in greater detail later in this book.

Public Key Confidentiality

The confidentiality objective of asymmetric algorithms is achieved when the encryption process is started with the public key. When the public key is used to encrypt the data, the private key must be used to decrypt the data. Because only one host has the private key, confidentiality is achieved.

Public key (encrypt) + Private key (decrypt) = Confidentiality



In [Figure 12-17](#), Alice and Bob exchange data with the goal of confidentiality. They follow these steps:

Step 1. Alice acquires Bob's public key (how the public key is acquired will be discussed later in this chapter when we discuss certificate authorities).

Step 2. Alice uses Bob's public key to encrypt a message, which is often a symmetric key, using an agreed upon algorithm.

Step 3. Alice transmits the encrypted message.

Step 4. Bob uses his private key to decrypt, and reveal, the message.




Figure 12-17. Asymmetric Confidentiality Process

We just saw how asymmetric encryption can be used to provide confidentiality. In an upcoming section, we will see how asymmetrical encryption is used to provide proof of authenticity and integrity by using digital signatures.

Encryption Algorithm Selection

Choosing an encryption algorithm is one of the most important steps of building a cryptography-based solution. You should consider two main criteria when selecting an encryption algorithm for your organization:

- **Trust in the algorithm by the cryptographic community:** Most new algorithms are broken quickly, so algorithms that have been resisting attacks for a number of years are preferred. Inventors and promoters often oversell the benefits of new algorithms. The truth is that there are few or no revolutions in cryptography.
- **Protection against brute-force attacks:** If the algorithm is considered trusted, there is no shortcut to break it, and the attacker must search through the keyspace to guess the correct key. The algorithm must allow key lengths that satisfy the confidentiality requirements of an organization. For example, DES no longer provides enough protection for most modern needs because of its short key length.

The following symmetric encryption algorithms are considered trustworthy:

- 3DES
- IDEA
- AES

The following asymmetric encryption algorithms, all covered in more detail later in this book, are considered trustworthy:

- RSA
- Elliptic Curve Digital Signature Algorithm (ECDSA)
- Diffie-Hellman (DH)

For symmetric algorithms, each bit in a key doubles the difficulty of finding the key. But with asymmetric algorithms, such as RSA, each additional bit only nominally increases the difficulty of

factoring the composite number that is used by the algorithm. Therefore, comparatively shorter key lengths for symmetric algorithms are similar to larger key lengths using asymmetric algorithms, when both types of algorithm are used for similar objectives (for instance, when comparing the two types of algorithm in terms of confidentiality and payload encryption). Symmetric and asymmetric keys compare as follow:

- An 80-bit symmetric key is considered equal to a 1024-bit key using the RSA algorithm.
- A 112-bit symmetric key is considered equal to a 2048-bit key using the RSA algorithm.
- A 128-bit symmetric key is considered equal to a 3072-bit key using the RSA algorithm.

For more information about the comparison of the key strengths of symmetric algorithms to RSA, refer to <http://www.rsasecurity.com/rsalabs/node.asp?id=2004>.

Cryptographic Hashes and Digital Signatures

Cryptographic hashes and digital signatures play an important part in modern cryptosystems. They provide verification and authentication and play an important role in non-repudiation. It is important to understand the basic mechanisms of these algorithms and some of the issues involved in choosing a particular hashing algorithm or digital signature method.

Hashing is a mechanism used for data integrity assurance. The hashing process uses a hash function, which is a one-way function of input data that produces a fixed-length digest of output data, also known as a fingerprint. The digest is cryptographically very strong. Hashes are functions that are relatively easy to compute, but *significantly* harder to reverse. Grinding coffee is a good example of a one-way function: it is easy to grind coffee beans, but it is impractical (not to say impossible) to put all the tiny pieces back together to rebuild the original beans.

If the input data changes just a little bit, the digest changes substantially. This is known as the *avalanche effect*. Essentially, the fingerprint that results from hashing data uniquely identifies that data. If you are given only a fingerprint, it is computationally infeasible to generate data that would result in that fingerprint.

[Figure 12-18](#) illustrates how hashing is performed. Data of arbitrary length is input into the hash function. Hashing is similar to the calculation of cyclic redundancy check (CRC) checksums, but it is much stronger cryptographically. That is, given a CRC value, it is easy to generate data with the same CRC. However, with hash functions, it is computationally infeasible for an attacker to have two different sets of data that would come up with the same fingerprint.

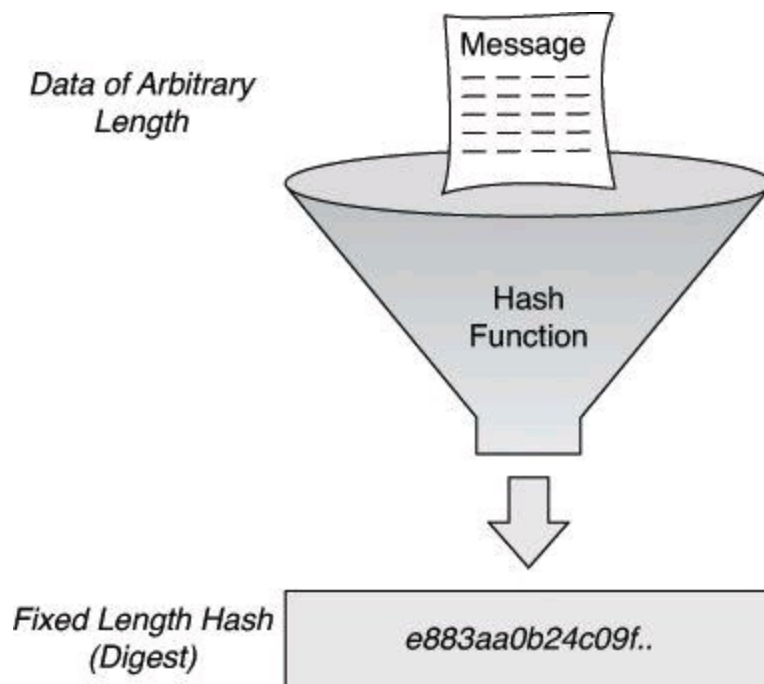


Figure 12-18. Cryptographic Hashes

Note

Extrapolating from the explanation above, it is easy to understand the possibility that two different pieces of data could produce the same hash result. This is called a collision. As an example, a few years ago I transferred my Microsoft Outlook PST file to a new computer. Upon launching MS Outlook from my new computer, I was asked for the password of the PST file. I couldn't remember it because it had been cached in the settings of my previous laptop for years and, therefore, I never had to type it in. Knowing that MS Outlook PST passwords are kept as a CRC32 hash in the PST file itself, I used freeware designed for this purpose and was able to get the hash result of my password. I then proceeded to find a CRC32 hash-guessing freeware tool and get the hash result found in the PST file. The hash-guessing tool provided me with a long alphanumeric code which, when hashed, produced the same result as the hashed password for my PST file. This code was obviously not my real password, but at least it provided the same hash result, so I proceeded to use that code as the password to open my PST file. This is how I cracked my own MS Outlook file. And, by the way, I immediately proceeded to change the Outlook password to a string I could remember.

This is an example of having two separate seed values producing the same hash results is an example of collision.

Key
Topic

The hashing process is not reversible; it is a *one-way function* with a fixed-length output. If you hash the word "Hello" with MD5 (covered later), the output, called the *message digest*, will be 128 bits long. If you process the *Oxford English Dictionary* through MD5, the message digest will be 128 bits long. It is impossible to take a message digest of 128 bits long and try to reverse-engineer it into

a 1200-page dictionary; thus the expression that hashing is a one-way function.

Hashing is often applied in the following situations:

- To generate one-time and one-way responses to challenges in authentication protocols such as PPP Challenge Handshake Authentication Protocol (PPP CHAP), Microsoft NT Domain, and Extensible Authentication Protocol-Message Digest 5 (EAP-MD5)
- To provide proof of the integrity of data, such as that provided with file integrity checkers, digitally signed contracts, and PKI certificates
- To provide proof of authenticity when it is used with a symmetric secret authentication key, such as IPsec or routing protocol authentication



Hashing algorithms are one-way processes.

A hash function, H , is a transformation that takes an input, (x) , and returns a fixed-size string, which is called the hash value, h . The formula for the calculation is $h = H(x)$.

A cryptographic hash function should have the following general properties:

- The input can be any length.
- The output has a fixed length.
- $H(x)$ is relatively easy to compute for any given x .
- $H(x)$ is one way and not reversible.
- $H(x)$ is collision free.

If a hash function is hard to invert, it is considered a one-way hash. “Hard to invert” means that given a hash value h , it is computationally infeasible to find some input, (x) , such that $H(x) = h$. H is said to be a weakly collision-free hash function if, given a message x , it is computationally infeasible to find a message y not equal to x such that $H(x) = H(y)$. A strongly collision-free hash function, H , is one for which it is computationally infeasible to find any two messages x and y such that $H(x) = H(y)$.

[Figure 12-18](#) illustrates the hashing process. Data of arbitrary length is input into the hash function, and the result of the hash function is the fixed-length digest, or fingerprint. Hashing is similar to the calculation of CRC checksums, but is cryptographically stronger. That is, given a CRC value, it is easy to generate data with the same CRC. However, as mentioned before, with hash functions, it is computationally infeasible for an attacker, given a hash value h , to find some input, (x) , such that $H(x) = h$.

[Figure 12-19](#) illustrates hashing in action. The sender wants to ensure that the message is not altered on its way to the receiver. The sending device inputs the message into a hashing algorithm and computes its fixed-length digest. This fingerprint is then attached to the message, where the message and the hash are in plaintext. The message, and its fingerprint attached, is sent to the receiver. The receiving device removes the fingerprint from the message and inputs the message into the same hashing algorithm. If the hash that is computed by the receiving device is equal to the one that is attached to the message, the message has not been altered during transit.

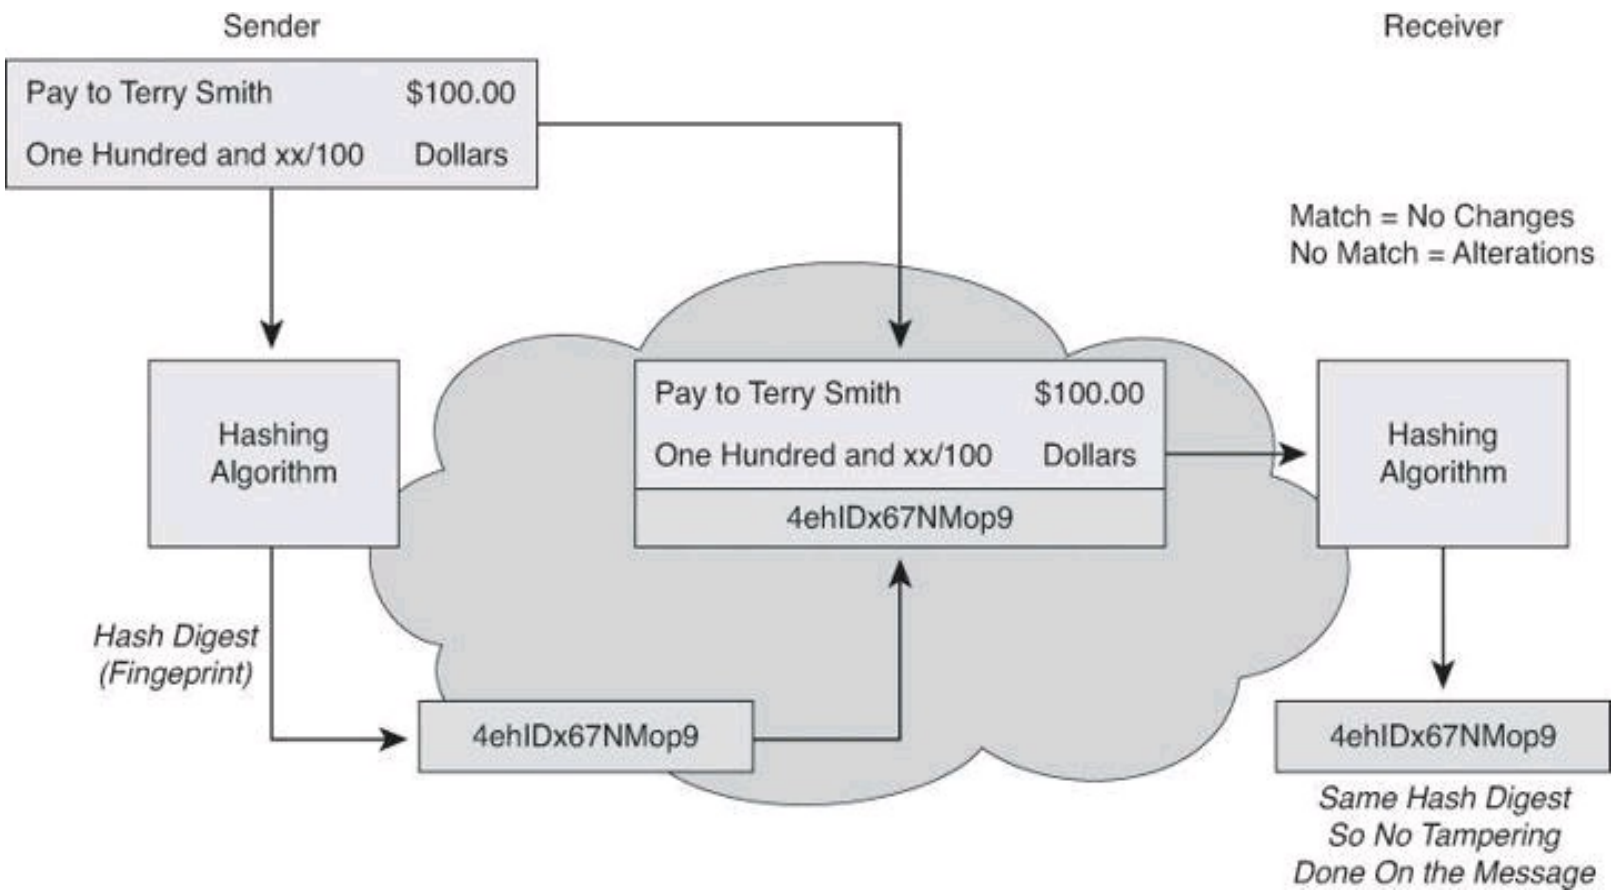


Figure 12-19. Data Integrity: Hashing in Action

Hashing does not add security to the message. When the message traverses the network, a potential attacker could intercept the message, change it, recalculate the hash, and append it to the message. Hashing only prevents the message from being changed accidentally, such as by a communication error. There is nothing unique to the sender in the hashing procedure; therefore, anyone can compute a hash for any data, as long as they have the correct hash function.

Thus, hash functions are helpful to ensure that data did not change accidentally, but it cannot ensure that data was not deliberately changed.

These are two well-known hash functions:

- Message Digest 5 (MD5) with 128-bit digests
- Secure Hash Algorithm 1 (SHA-1) with 160-bit digests

Hashing Algorithms

The two most commonly used cryptographic hash functions are MD5 and SHA-1. SHA-2 is an additional and newer option. [Table 12-4](#) compares the three algorithms in terms of strength, key size, and vulnerability to cryptanalytic attacks.

Table 12-4. Comparing Hashing Algorithms

Algorithm	Description
MD5	Ubiquitous hashing algorithm 128-bit message digest Not recommended for new applications
SHA-1	160-bit message digest Preferred to MD5 because, although slower, is the stronger algorithm
SHA-2	Similar to SHA-1, with message digest of 224, 256, 384, or 512 bits Adopted by U.S. federal government as secure hash standard in 2008

Note

Although originally thought to be collision resistant, both MD5 and SHA were later found to be vulnerable to such forms of attack.

MD5

The MD5 algorithm is a ubiquitous hashing algorithm that was developed by Ron Rivest and is used in a variety of Internet applications today.

MD5 is a one-way function, which makes it easy to compute a hash from the given input data but makes it infeasible to compute input data given only a hash. MD5 is also collision resistant, which means that two messages with the same hash are very unlikely to occur. MD5 is essentially a complex sequence of simple binary operations, such as XORs and rotations, that are performed on input data and produce a 128-bit digest.

The main algorithm itself is based on a compression function, which operates on blocks. The input is a data block plus a feedback of previous blocks. Each 512-bit block is divided into sixteen 32-bit subblocks. These blocks are then rearranged with simple operations in a main loop, which consists of four rounds. The output of the algorithm is a set of four 32-bit blocks, which concatenate to form a single 128-bit hash value. The message length is also encoded into the digest.

MD5 is based on MD4, an earlier algorithm. MD4 has been broken, and currently MD5 is considered less secure than SHA-1 by many authorities on cryptography, such as ICSA Labs (<http://www.icsalabs.com>). These authorities consider MD5 less secure than SHA-1 because some noncritical weaknesses have been found in one of the MD5 building blocks, causing uneasy feelings inside the cryptographic community. The availability of the SHA-1 and RACE Integrity Primitives Evaluation Message Digest (RIPEMD)-160 HMAC functions, which do not show such weaknesses and use a stronger (160-bit) digest, makes MD5 a second choice as far as hash methods are concerned.

SHA-1

The U.S. National Institute of Standards and Technology (NIST) developed the Secure Hash Algorithm (SHA), the algorithm that is specified in the Secure Hash Standard (SHS). SHA-1 is a revision to the SHA that was published in 1994; the revision corrected an unpublished flaw in SHA.

Its design is similar to the MD4 family of hash functions that Ron Rivest developed.

The SHA-1 algorithm takes a message of less than 264 bits in length and produces a 160-bit message digest. The algorithm is slightly slower than MD5, but the larger message digest makes it more secure against brute-force collision and inversion attacks. More details on SHA-1 collision attacks can be found at <http://www.rsa.com/rsalabs/node.asp?id=2927>.

You can find the official text for the standard at <http://www.itl.nist.gov/fipspubs/fip180-1.htm>.

Note

SHA also has 224-, 256-, 384-, and 512-bit versions.

SHA-2

Secure Hash Algorithm 2 (SHA-2) specifies five SHAs—SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512—for computing a condensed representation of electronic data. When a message of any length less than 264 bits (for SHA-224 and SHA-256) or less than 2128 bits (for SHA-384 and SHA-512) is input to a hash algorithm, the result is a message digest that ranges in length from 224 to 512 bits, depending on the algorithm.

The SHA-2 family of hash functions was approved by NIST for use by federal agencies in 2006, for all applications using SHAs. The publication encouraged all federal agencies to stop using SHA-1 for digital signatures, digital time stamping, and other applications that require collision resistance as soon as practical, and it mandated the use of the SHA-2 family of hash functions for these applications after 2010. After 2010, federal agencies used SHA-1 only for the following applications: HMACs, key derivation functions (KDF), and random number generators (RNG). This change was triggered in 2005 when security flaws were identified for SHA-1 in theoretical exploits that exposed weaknesses to collision attacks.

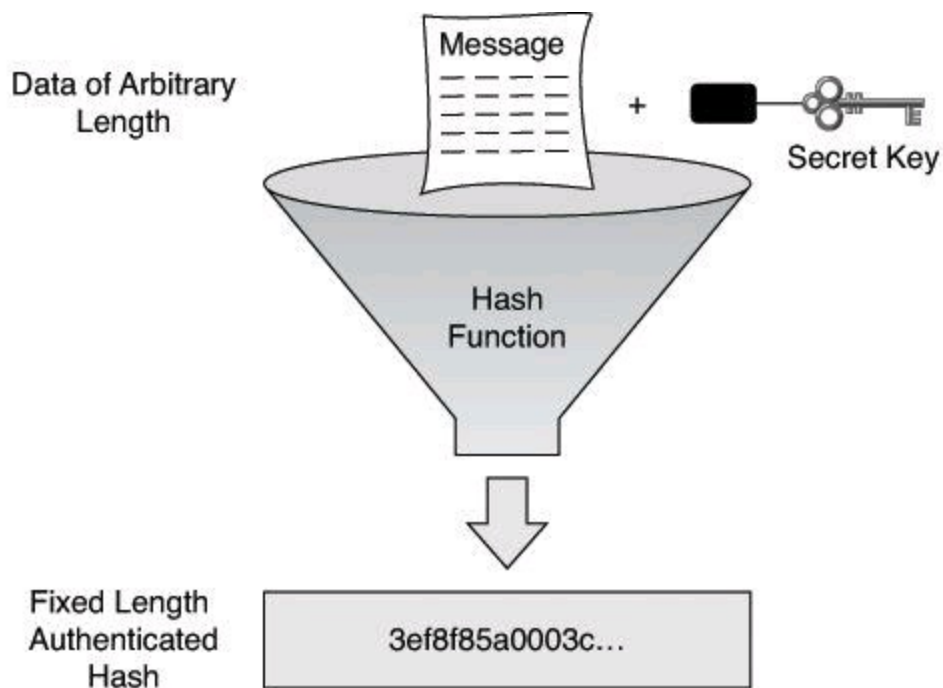
Hashed Message Authentication Codes

Hash functions are the basis of the protection mechanism of Hashed Message Authentication Codes (HMAC). HMACs use existing hash functions, but with a significant difference; HMACs add a secret key as input to the hash function. Only the sender and the receiver know the secret key, and the output of the hash function now depends on the input data and the secret key. Therefore, only parties who have access to that secret key can compute the digest of an HMAC function. This behavior defeats man-in-the-middle attacks and provides authentication of the data origin. If two parties share a secret key and use HMAC functions for authentication, a properly constructed HMAC digest of a message that a party has received indicates that the other party was the originator of the message, because it is the only other entity possessing the secret key.

Cisco technologies use two well-known HMAC functions:

- Keyed MD5, based on the MD5 hashing algorithm
- Keyed SHA-1, based on the SHA-1 hashing algorithm

[Figure 12-20](#) illustrates how an HMAC digest is created. Data of an arbitrary length is input into the hash function, together with a secret key. The result is the fixed-length hash that depends on the data and the secret key.



- Same procedure is used for generation and verification of secure fingerprints.

Figure 12-20. HMAC Digest Creation

[Figure 12-21](#) illustrates an HMAC in action. The sender wants to ensure that the message is not altered in transit and wants to provide a way for the receiver to authenticate the origin of the message.

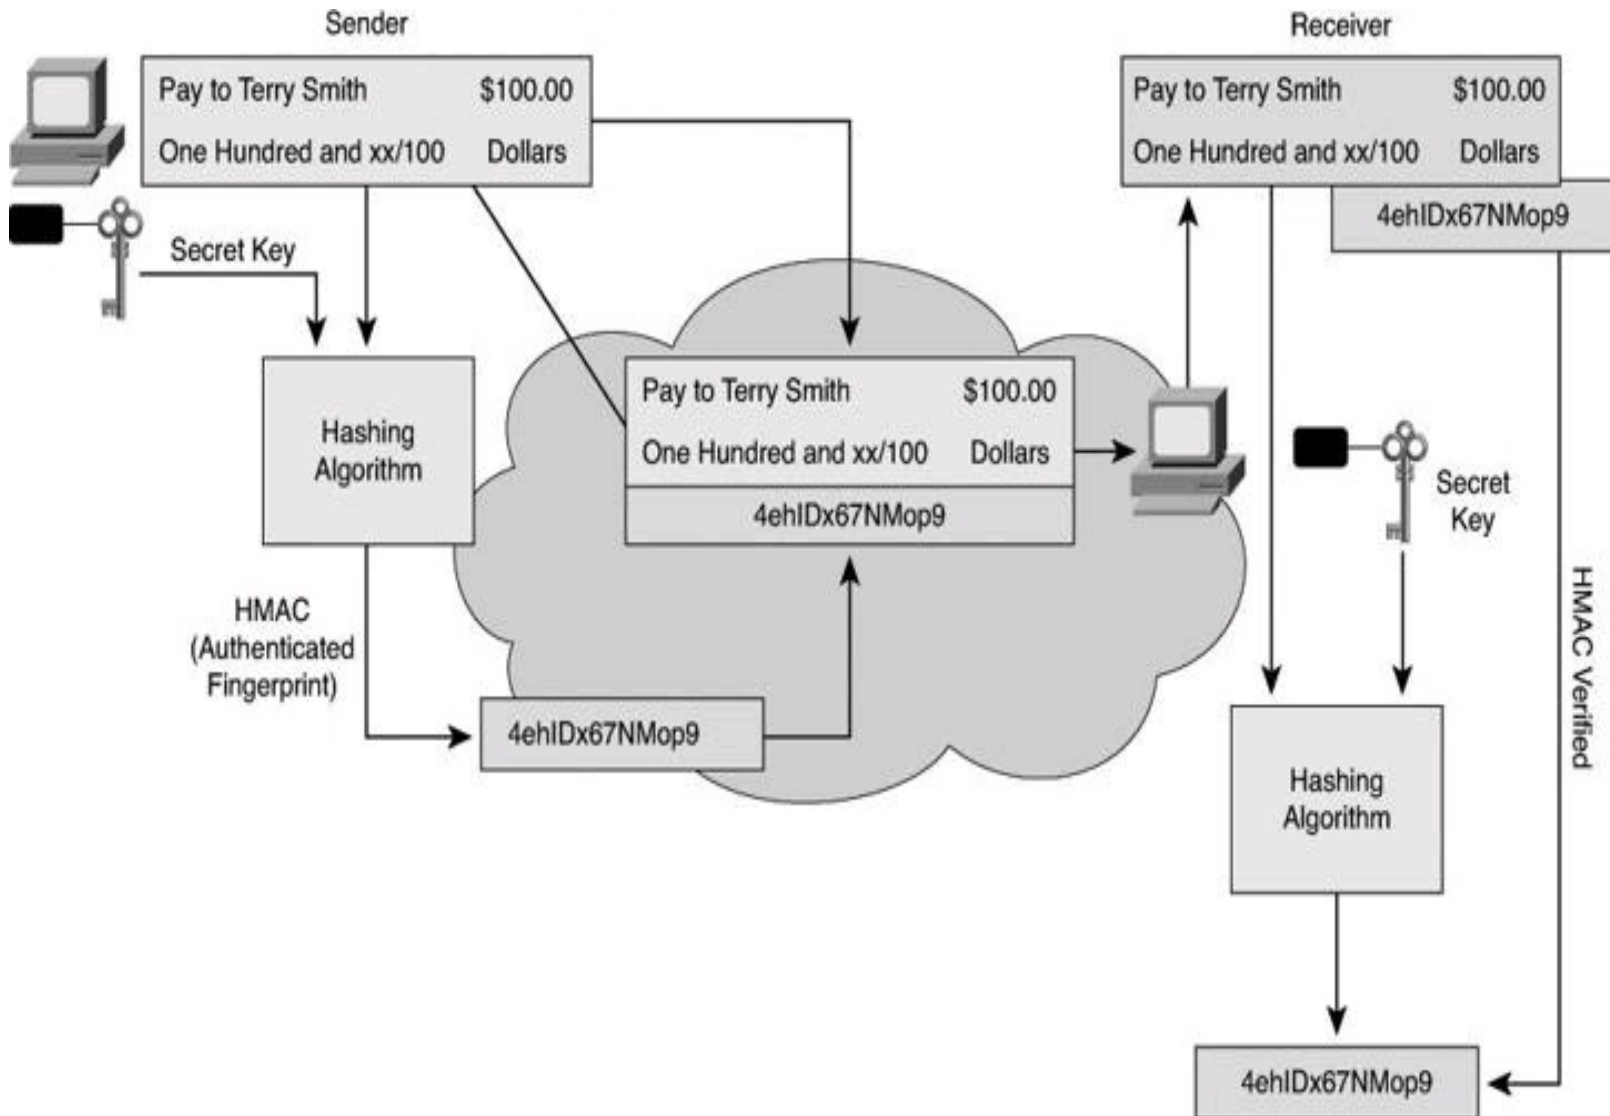


Figure 12-21. HMAC in Action

The sending device inputs data and the secret key into the hashing algorithm and calculates the fixed-length HMAC digest. This authenticated fingerprint is then attached to the message and sent to the receiver. The receiving device removes the fingerprint from the message and uses the plaintext message with the secret key as input to the same hashing function. If the fingerprint calculated by the receiving device is equal to the fingerprint that was sent, the message has not been altered. In addition, the origin of the message is authenticated, because only the sender possesses a copy of the shared secret key. The HMAC function has ensured the authenticity of the message.

Note

IPsec VPNs rely on HMAC functions to authenticate the origin and provide data integrity checking of every packet.

Cisco products use hashing for entity-authentication, data-integrity, and data-authenticity purposes:

- IPsec gateways and clients use hashing algorithms, such as MD5 and SHA-1 in HMAC mode, to provide packet integrity and authenticity.
- Cisco IOS routers use hashing with secret keys in an HMAC-like manner, to add authentication information to routing protocol updates.
- Cisco software images that you can download from Cisco.com have an MD5-based checksum available so that customers can check the integrity of downloaded images.
- Hashing can also be used in a feedback-like mode to encrypt data; for example, TACACS+ uses MD5 to encrypt its session.

Overview of Digital Signatures

When data is exchanged over untrusted networks, several major security issues must be determined:

- **Whether data has changed in transit:** Hashing and HMAC functions rely on a cumbersome exchange of secret keys between parties to provide the guarantee of integrity.
- **Whether a document is authentic:** Hashing and HMAC can provide some guarantee of authenticity, but only by using secret keys between two parties. Hashing and HMAC cannot guarantee authenticity of a transaction or a document to a third party.

Digital signatures are often used in the following situations:

- To provide a unique proof of data source, which can be generated only by a single party, such as with contract signing in e-commerce environments
- To authenticate a user by using the private key of that user and the signature it generates
- To prove the authenticity and integrity of PKI certificates
- To provide a secure time stamp, such as with a central trusted time source

Suppose a customer sends transaction instructions via an email to a stockbroker, and the transaction turns out badly for the customer. It is conceivable that the customer could claim never to have sent the transaction order, or that someone forged the email. The brokerage could protect itself

by requiring the use of digital signatures before accepting instructions via email.

Handwritten signatures have long been used as a proof of authorship of, or at least agreement with, the contents of a document. Digital signatures can provide the same functionality as handwritten signatures, and much more.

Digital signatures provide three basic security services in secure communications:

- **Authenticity of digitally signed data:** Digital signatures authenticate a source, proving that a certain party has seen and has signed the data in question.
- **Integrity of digitally signed data:** Digital signatures guarantee that the data has not changed from the time it was signed.
- **Nonrepudiation of the transaction:** The recipient can take the data to a third party, which accepts the digital signature as a proof that this data exchange did take place. The signing party cannot repudiate that it has signed the data.

Note

To better understand nonrepudiation, consider the use of HMAC functions, which also provide authenticity and integrity guarantees. With HMAC functions, two or more parties share the same authentication key and can compute the HMAC fingerprint. Therefore, taking received data and its HMAC fingerprint to a third party does not prove that the other party sent this data; you could have generated the same HMAC fingerprint yourself, because you have a copy of the HMAC authentication key. With digital signatures, each party has a unique, secret signature key, which is not shared with any other party, making nonrepudiation possible.

To achieve the preceding goals, digital signatures have the following properties:

- **The signature is authentic:** The signature convinces the recipient of the document that the signer signed the document.
- **The signature is not forgeable:** The signature is proof that the signer, and no one else, signed the document.
- **The signature is not reusable:** The signature is a part of the document and cannot be moved to a different document.
- **The signature is unalterable:** After a document is signed, the document cannot be altered without detection.
- **The signature cannot be repudiated:** The signature and the document are physical things. The signer cannot claim later that they did not sign it.

Well-known asymmetric algorithms, such as RSA or DSA, are typically used to perform digital signing.

In some countries, including the United States, digital signatures are considered equivalent to handwritten signatures, if they meet certain provisions. Some of these provisions include the proper protection of the certificate authority, the trusted signer of all other public keys, and the proper protection of the private keys of the users. In such a scenario, users are responsible for keeping their

private keys private, because a stolen private key can be used to “steal” their identity.

Later in this chapter you will have the opportunity to delve deeper into the mechanics of digital signatures, but for now [Figure 12-22](#) illustrates the basic functioning of digital signatures:

Step 1. When someone wants to sign some data, they use a signature algorithm with their signature key. Only the signer knows this signature key. Therefore, the signer must keep the signature key secret.

Step 2. Based on the input data and a signature key, the signature algorithm generates its output, which is called a digital signature.

Step 3. The sending device attaches the digital signature to the message and sends the message to the receiver.

Step 4. The receiving device verifies the signature with the verification key, which is usually public.

Step 5. The receiving device inputs the message, the digital signature, and the verification key into the verification algorithm, which checks the validity of the digital signature.

Step 6. If the check is successful, the document has not been changed after signing and the document was originated by the signer of the document.

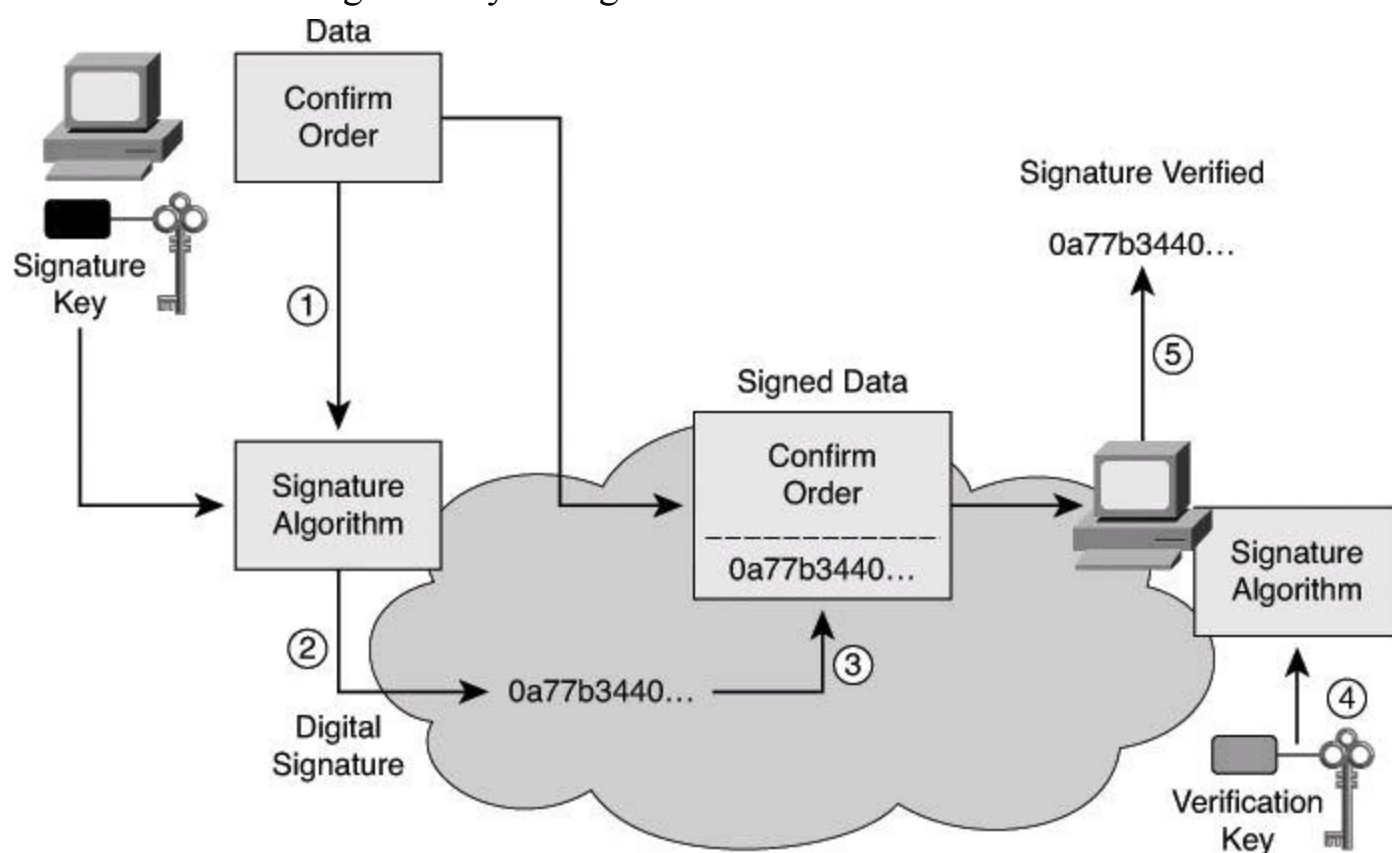


Figure 12-22. Digital Signatures in Action

Digital Signatures = Encrypted Message Digest

A digital signature is the result of encrypting, with the user's private key, the digest and appending that encrypted digest to the plaintext or encrypted message to verify the identity of the sender. The digest will be decrypted with the corresponding public key.

A digital signature provides authentication and integrity. If the recipient is successful at decrypting

the digest using the public key of the sender, the recipient has proof of the origin of the message. Also, if both hashes have the same value—the hash calculated by the recipient upon receiving the message and the decrypted hash that was appended to the message in the first place—the recipient has proof that the message wasn't tampered with during transmission, and thus proof of its integrity.

Digital signatures are commonly used to provide assurance of the code authenticity and integrity of both mobile and classic software. As an example, you might have noticed that some executable files (or possibly the whole installation package of a program) are wrapped with a digitally signed envelope, which allows the end user to verify the signature before installing the software.

Digitally signing code provides several assurances about the code:

- The code has not been modified since it left the software publisher.
- The code is authentic and is actually sourced by the publisher.
- The publisher undeniably publishes the code. This provides nonrepudiation of the act of publishing.

The digital signature could be forged only if someone obtained the private key of the publisher. The assurance level of digital signatures is extremely high if the private key is protected properly.

The user of the software must also obtain the public key, which is used to verify the signature. The user can obtain the key in a secure fashion; for example, the key could be included with the installation of the operating system, or transferred securely over the network (for example, using the PKI and certificate authorities).

Diffie-Hellman

Earlier, we discussed how symmetric encryption is faster than asymmetric encryption. However, a significant drawback of symmetric encryption is key generation and distribution. The Diffie-Hellman algorithm is one answer to this problem. The topic of key management, such as key generation and distribution, is covered in more detail in the next section. For now, let's only focus on the mathematical function that could be used to solve these problems.

The goal of DH is that each peer deduces the same secret number following the exchange of some values, in cleartext. That secret number that both peers deduce could then be used as the symmetric encryption key.

The DH algorithm is the basis of most modern automatic key exchange methods. The Internet Key Exchange (IKE) protocol in IPsec VPNs uses DH algorithms extensively to provide a reliable and trusted method for key exchange over untrusted channels.

Whitfield Diffie and Martin Hellman invented the DH algorithm in 1976. Its security stems from the difficulty of calculating the discrete logarithms of very large numbers. The DH algorithm, depicted in [Figure 12-23](#), provides secure key exchange over unsecure channels and is frequently used in modern key management to provide keying material for other symmetric algorithms, such as DES, 3DES, and AES. Don't be too intimidated by [Figure 12-23](#) and the explanation that follows. Later we will demonstrate an example of DH and you'll see it's not too complicated.

First, Alice and Bob need to agree on g (generator) and p (modulus ~ prime).

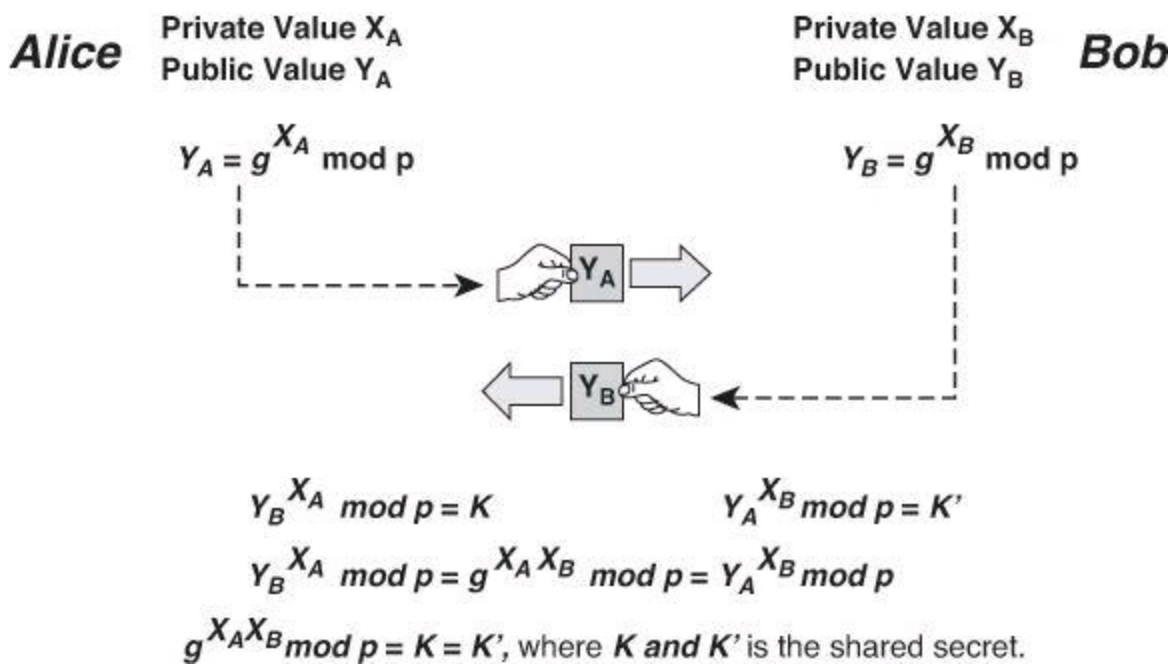


Figure 12-23. Diffie-Hellman Key Exchange Algorithm

To start a DH exchange, the two parties must agree on two nonsecret numbers. The first number is g , the generator, and the second number is p , the modulus. These numbers can be made public and are usually chosen from a table of known values. g is usually a very small number, such as 2, 3, or 4, and p is a very large prime number.

Next, every party generates its own secret value. Then, based on g , p , and the secret value of each party, each party calculates its public value. The public value is computed according to the following formula:

$$Y = g^x \bmod p$$

In this formula, x is the secret value of the entity, and Y is the public value of the entity.

After computing the public values, the two parties exchange their public values. Each party then exponentiates the received public value with its secret value to compute a common shared-secret value, represented by K and K' in [Figure 12-23](#). When the algorithm completes, both parties have the same shared secret, which they have computed from their secret value and the public value of the other party.

No one listening on the channel can compute the secret value, because only g , p , Y_A , and Y_B are known; at least one secret value is needed to calculate the shared secret. Unless attackers can compute the discrete algorithm of the preceding equation to recover X_A or X_B , they cannot obtain the shared secret.

The following steps describe a DH exchange:

Step 1. Alice and Bob agree on generator g and modulus p .

Step 2. Alice chooses a random large integer, X_A , and sends Bob its public value, Y_A ,

where $Y_A = g^{x_A} \bmod p$.

Step 3. Bob chooses a random large integer, X_B , and sends Alice his public value, Y_B , where $Y_B = g^{x_B} \bmod p$.

Step 4. Alice computes $K = Y_B^{x_A} \bmod p$.

Step 5. Bob computes $K' = Y_A^{x_B} \bmod p$.

Step 6. Both K and K' are equal to $g^{x_A x_B} \bmod p$.

Alice and Bob now have a shared secret ($K = K'$), and even if someone has listened on the untrusted channel, there is no way the listener could compute the secret from the captured information, assuming that computing a discrete logarithm of Y_A or Y_B is practically infeasible.

Note

RFC 2409 (<http://www.ietf.org/rfc/rfc2409>) and RFC 3526 (<http://www.ietf.org/rfc/rfc3526>) provide more details about the values of g and p .

Diffie-Hellman Example

As you just read, a modulus operation is at the heart of DH. In simple terms, a modulus gives you the remainder of a division when the result of that division is an integer; you don't want a fraction here. [Figure 12-24](#) shows a simple modulus operation.

27 mod 6 = ?

Remember the goal of the modulus: to find the remainder when a division produces an integer.

How many times does 6 go into 27? Four times.

$$\begin{array}{r} 27 \quad | \quad 6 \\ -24 \quad | \quad 4 \\ \hline 3 \end{array}$$

What is the remainder? 3

So, 27 mod 6 = 3

Figure 12-24. Simple Modulus Calculation

So, let's try a Diffie-Hellman together. Follow along with [Figure 12-25](#).

Step 1 In the clear, Alice and Bob agree on: generator = 7

Prime number for modulus = 13

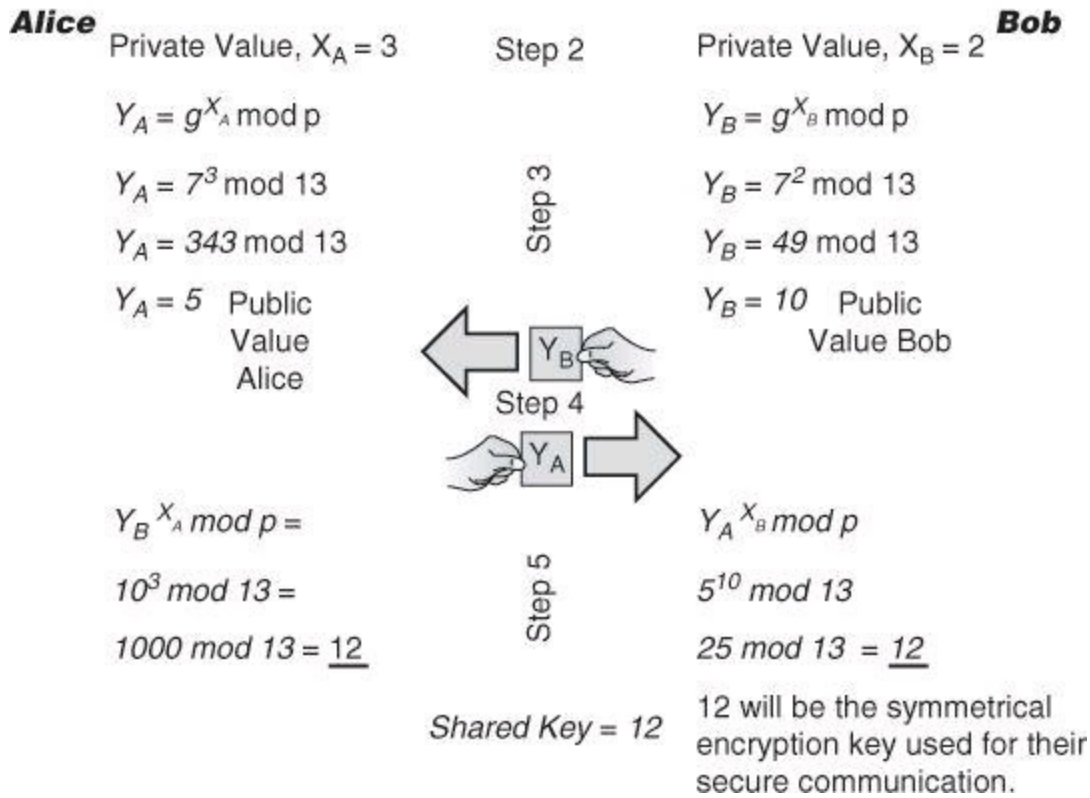


Figure 12-25. Diffie-Hellman Example

Step 1. In cleartext, Alice and Bob agree on a generator and on a modulus. The generator is usually a small number. In our example, shown in [Figure 12-25](#), the peers, Alice and Bob, agree that the generator will be 7. Typically, the modulus is a very large prime number. To keep our calculations simple, we are using a very small modulus number, 13.

Step 2. Alice and Bob independently come up with a random private value that they will keep secret. Alice randomly generates as Alice's private value 3. Bob randomly generates as Bob's private value 2.

Step 3. Alice and Bob proceed to derive their public value, using the generator and primer they agreed on in Step 1 and their own private value they generated in Step 2. Review [Figure 12-24](#) to confirm the procedures of modulus calculations of Step 3. Through the modulus operation, Alice derives that her public value is 5 and Bob derives that his public value is 10.

Step 4. They proceed at sharing their public value with each other, in the clear. Thus, Alice hears that Bob's public value is 10 and Bob hears that Alice's public value is 5.

Step 5. They use their peer's public value, their own private value, and the generator and the prime number of Step 1 to finally calculate a key of 12. The result in Step 5, though independently calculated, will be the same. This will be the shared key that will be used for symmetric encryption. In our example, the key is decimal 12, which is very small. It's only 4 bits long. In real life, with AES as an example, the symmetric keys are 128, 192, or 256 bits long.

If you had sniffed the communication, what would you have learned? The generator (7), the prime number (13), Alice's public value (5), and Bob's public value (10). But you would not know the

private values used by Alice and Bob, so it would be impossible for you to proceed with Step 5. DH hasn't been broken yet.



The result of a DH calculation, which will be the same result on both peers, could be used as the encryption key for symmetric algorithms, such as when encrypting with DES, 3DES, or AES.

Additional information can be found at <http://www.rsa.com/rsalabs/node.asp?id=2248>. There is also a great and colorful explanation of Diffie-Hellman at http://en.wikipedia.org/wiki/Diffie%E2%80%93Hellman_key_exchange.

Cryptographic Processes in VPNs

In the previous pages, you learned about encryption algorithms such as AES, hashing algorithms such as SHA, and key management algorithms such as DH. The combination of symmetric algorithms, hashing functions, and key management algorithms is what comprises IPsec. In [Chapter 14](#), you will see how all these pieces fit together.

Similarly, the SSL and TLS protocols support the use of a variety of different cryptographic algorithms, or ciphers, for use in operations. In [Chapter 15](#), you will see in more detail how SSL and TLS work.

Different protocols support different cryptographic algorithms to accomplish these goals. The selection of the protocol is part of the design phase of VPN implementations, and is directly tied to the strength of the protocol itself, as well as the strength of the keys. The strength of the keys, as you know, is directly related to the key size.

The prevalent algorithms today are 160-bit hashes, provided by SHA-1, and 1024-bit DH keys, among others.

Federal Information Processing Standard (FIPS) 140-2 is the current version of the FIPS 140 publication that specifies requirements for cryptography modules in the United States. NIST issued the FIPS 140 series to uphold the standards that describe the U.S. federal government requirements that IT products should meet. The 2011 specification calls for the following cryptographic algorithms: encryption with AES-128, key generation with Diffie-Hellman 2048 bit, digital signature with RSA 2048 bit, and hashing with SHA 256-bit. These stricter requirements are an indication of the ever-evolving and dynamic nature of threats and acceptable levels of risk in information security. Note that FIPS 140-3 is currently available in its draft format.

Asymmetric Encryption: Digital Signatures

Asymmetric encryption algorithms accomplish two primary objectives: confidentiality and authentication. Asymmetric algorithms are slower than symmetric algorithms because they use more complex mathematics. Because asymmetric algorithms are slower, they are usually used as key exchange protocols and are rarely used for bulk encryption. The sections that follow cover the principles behind asymmetric encryption for digital signatures. Later in the chapter, we will also see how asymmetric encryption is used for PKI environments.

Asymmetric Encryption Overview

To provide the two main objectives of confidentiality and authentication, asymmetric algorithms are based on mathematical formulas that are considerably more complex than the formulas symmetric algorithms are based on. As a result, computation takes more time for asymmetric algorithms. Despite this slower computation trait, asymmetric algorithms are often used as key exchange protocols for symmetric algorithms, which have no inherent key exchange technology.

Also known as *public-key encryption*, asymmetric algorithms have two keys: a public key and a private key. Both keys are capable of the encryption process. However, the complementary matched key is required for decryption. For example, if a public key encrypts the data, the matching private key decrypts the data. The opposite is also true. If a private key encrypts the data, the corresponding public key decrypts the data.

Examples of public-key encryption algorithms are RSA, DSA, DH, ElGamal, and elliptic curve cryptography (ECC). The mathematical operations differ with each algorithm, but the algorithms all share one trait: the mathematics are complicated.

Asymmetric algorithms are designed in such a way that the key that is used for encryption is different from the key that is used for decryption ([Figure 12-16](#), shown earlier, illustrates the mechanics of asymmetric encryption). The decryption key cannot, in any reasonable amount of time, be calculated from the encryption key, and vice versa. The usual key length for asymmetric algorithms ranges from 512 to 4096 bits. Asymmetric algorithm key lengths cannot be directly compared to symmetric algorithm key lengths because the two algorithm families differ greatly in their underlying design.

To illustrate this point, it is generally thought that an encryption key of RSA that is 2048 bits is roughly equivalent to a 128-bit key of RC4 in terms of resistance against brute-force attacks.

Public Key Authentication

The authentication objective of asymmetric algorithms is achieved when the encryption process is started with the private key. When the private key is used to encrypt the data, the public key must be used to decrypt the data. Because only one host has the private key, only that host could have encrypted the message, therefore providing authentication of the sender.

Private key (encrypt) + Public key (decrypt) = Authentication.



Typically, no attempt is made to preserve the secrecy of the public key, so any number of hosts can decrypt the message. When a host successfully decrypts a message using a public key, it is trusted that the private key encrypted the message, which verifies who the sender is; this is a form of authentication.

In [Figure 12-26](#), Alice and Bob exchange data with the goal of authentication. They follow these steps:

Step 1. Alice, using her private key, creates a digital signature and appends it to the message.

Step 2. Alice transmits the signed message to Bob.

Step 3. Bob acquires Alice's public key.

Step 4. Bob uses Alice's public key to verify the signature.

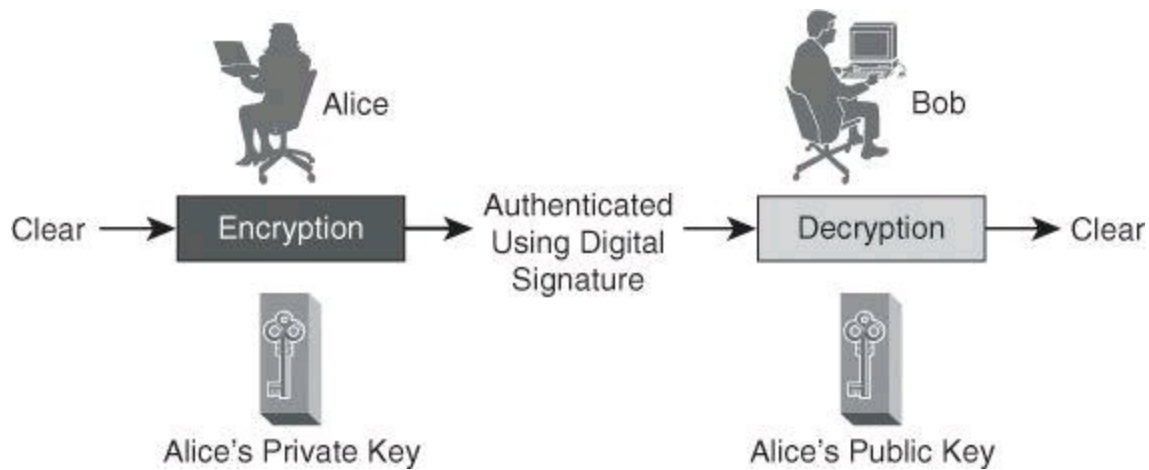


Figure 12-26. Asymmetric Authentication Process = Digital Signature

If Bob wishes to also sign a document, he will need to go through Steps 1 and 2 using his private key, and Alice will need to acquire Bob's public key to check his signature. Therefore, the authentication is unidirectional. Each party is responsible for creating his or her own signature.

Safekeeping the Private Key

As you can imagine, the private key must be kept secret. As an example, in some organizations, the private key is kept on a USB thumb drive. When the owner of the private key wishes to use it, he inserts it into the USB port and is typically asked for the PIN to unlock the private key. Using a private key provides strong authentication. Strong authentication requires using two separate authentication methods. To unlock the private key, we are therefore using strong authentication: something you have (the USB thumb drive) and something you know (the PIN).

Caution

If the private key is compromised, another key pair must be generated to replace the compromised key.

RSA and Digital Signatures

RSA is one of the most common asymmetric algorithms. Ron Rivest (discussed earlier in this chapter), Adi Shamir, and Len Adleman invented the patented public-key RSA algorithm in 1977. The patent expired in September 2000, and the algorithm is now in the public domain. Of all the public-key algorithms that were proposed over the years, RSA is by far the easiest to understand and implement.

The RSA algorithm is very flexible because it has a variable key length that allows speed to be traded for the security of the algorithm if necessary.

The RSA keys are usually 512 to 2048 bits long. RSA has withstood years of extensive cryptanalysis, and although the security of RSA has been neither proved nor disproved, it does

suggest a confidence level in the algorithm. The security of RSA is based on the difficulty of *factoring* very large numbers, which means breaking large numbers into multiplicative factors. If an easy method of factoring these large numbers were discovered, the effectiveness of RSA would be destroyed.

The RSA algorithm is based on the fact that each entity has two keys: a public key and a private key. The public key can be published and given away, but the private key must be kept very secret. It is not possible to determine, using any computationally feasible algorithm, the private key from the public key, and vice versa. What one of the keys encrypts, the other key decrypts, and vice versa.

RSA keys are long term and are usually changed or renewed after some months or even years.

The current procedures for signing digital signatures are not simply implemented by public-key operations. In fact, a modern digital signature is based on a hash function and a public-key algorithm. [Figure 12-27](#) illustrates this procedure.

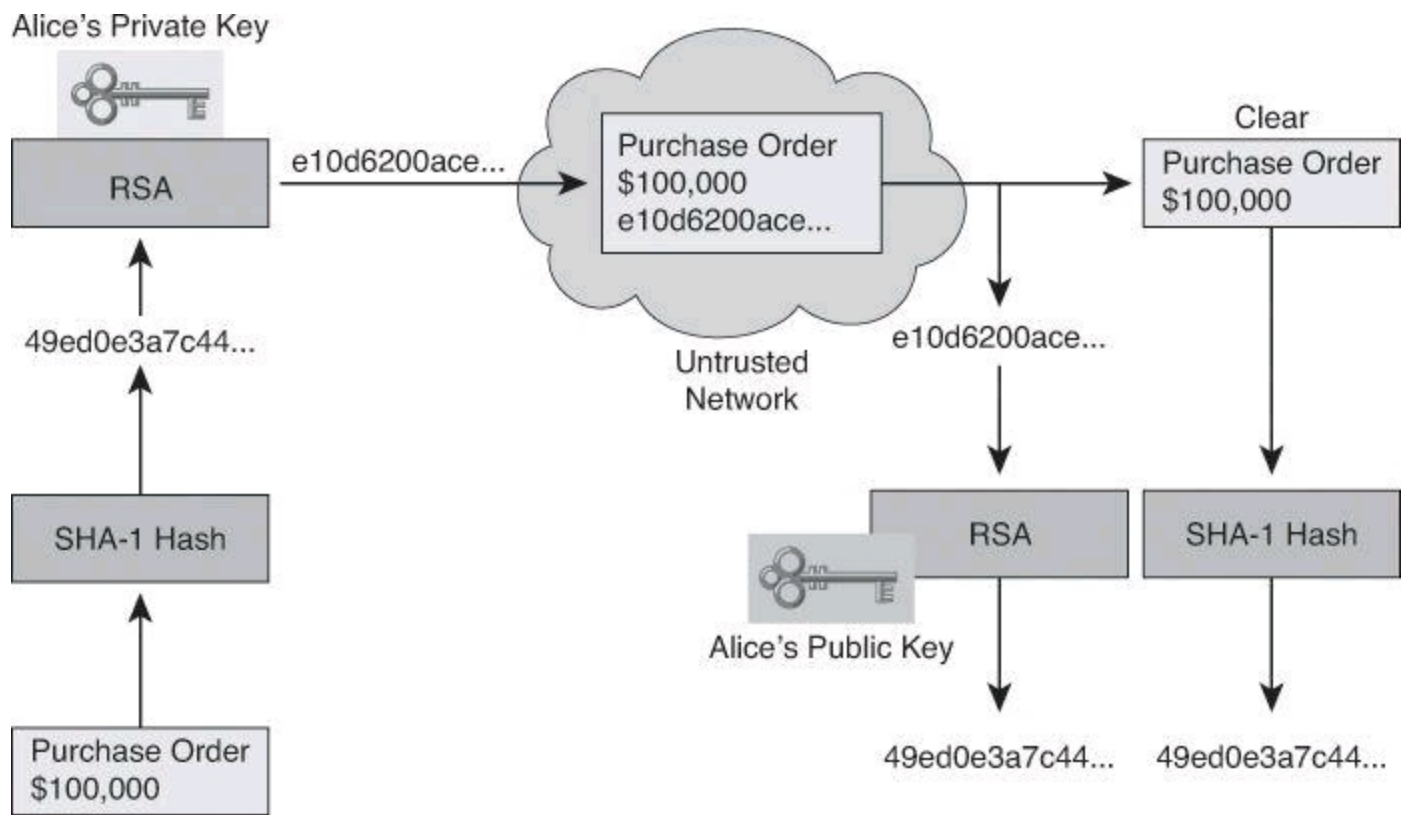


Figure 12-27. RSA Digital Signatures

The signature process shown in [Figure 12-27](#) is as follows:

- Step 1.** The signer (Alice) makes a hash, or fingerprint, of the document, which uniquely identifies the document and all of its contents.
- Step 2.** Alice encrypts the hash only with her private key.
- Step 3.** The encrypted hash, known as the signature, is appended to the document.

The verification process works as follows:

- Step 1.** The verifier (Bob) obtains Alice's public key.
- Step 2.** Bob decrypts the signature using Alice's public key. This step unveils the assumed hash value of the signer.

Step 3. Bob makes a hash of the received document, without its signature, and compares this hash to the decrypted signature hash. If the hashes match, the document is authentic; that is, it has been signed by Alice, and has not changed since it was signed.

This example illustrates how the authenticity and integrity of the message are ensured even though the actual text is public. Both encryption and digital signatures are required to ensure that the message is private and has not changed.

Note

The RSA algorithm is currently the most common method for signature generation and is used widely in e-commerce systems and Internet protocols in that role.

RSA is substantially slower than DES in both hardware and software implementations. This performance problem is the main reason that RSA is typically used only to protect small amounts of data. RSA is mainly used for two services:

- To ensure confidentiality of data by performing encryption
 - To perform authentication of data, nonrepudiation of data, or both by generating digital signatures
-

Note

RSA encryption is faster than decryption, and verification is faster than signing.

Public Key Infrastructure

The example in the previous section of authentication using asymmetric encryption and digital signatures suffers from one critical drawback: scalability. If applied to two parties, such as the Bob and Alice example, you need two sets of private/public keys to authenticate the two parties. As the number of parties increases, and if you want to maintain authentication separate for each pair of entities, the number of private/public keys increases exponentially, and the number of validations to verify the signatures does as well.

For instance, if 10 individuals need to validate each other, 95 validations would need to be performed before everyone would have validated everyone else. Adding a single individual to the group would require an additional 20 validations because each of the original 10 individuals would need to authenticate the new individual, and the new individual would need to authenticate the original 10. This method does not scale well.

The impact is more tangible in real-life environments involving large organizations. In those scenarios, it is impractical for all parties to continually exchange identification documents. For example, Cisco goes to reasonable measures to identify employees and contractors, and then issues ID badges. The badge is relatively difficult to forge. Measures are in place to protect the integrity of the badge and the badge issuance. Because of these measures, all Cisco personnel accept this badge as authoritative as to the identity of any individual.

The badge scenario is an example of a trusted third-party protocol. This concept involves the use

of a trusted introducer to all the parties trying to communicate. All individuals agree to accept the word of this neutral third party. In this way, parties that need to validate each other rely on the in-depth authentication of an agreed-upon third party instead of performing their own authentication. Presumably, the third party does an in-depth investigation before the issuance of credentials. After this in-depth investigation, the third-party issues credentials that are difficult to forge. From that point forward, all individuals who trust the third party simply accept the credentials that the third party issues.

Passports and driver's licenses are real-life examples of a trusted third-party environment that uses the concept of a trusted introducer. Certificate authority (CA) servers are an example of this concept in PKI environments.

Note

Certificate servers are an example of a trusted third party.

In [Figure 12-28](#), Alice applies to the license bureau for a driver's license. In this process, she submits evidence of her identity and her qualifications to drive. Her application is approved and a license is issued.

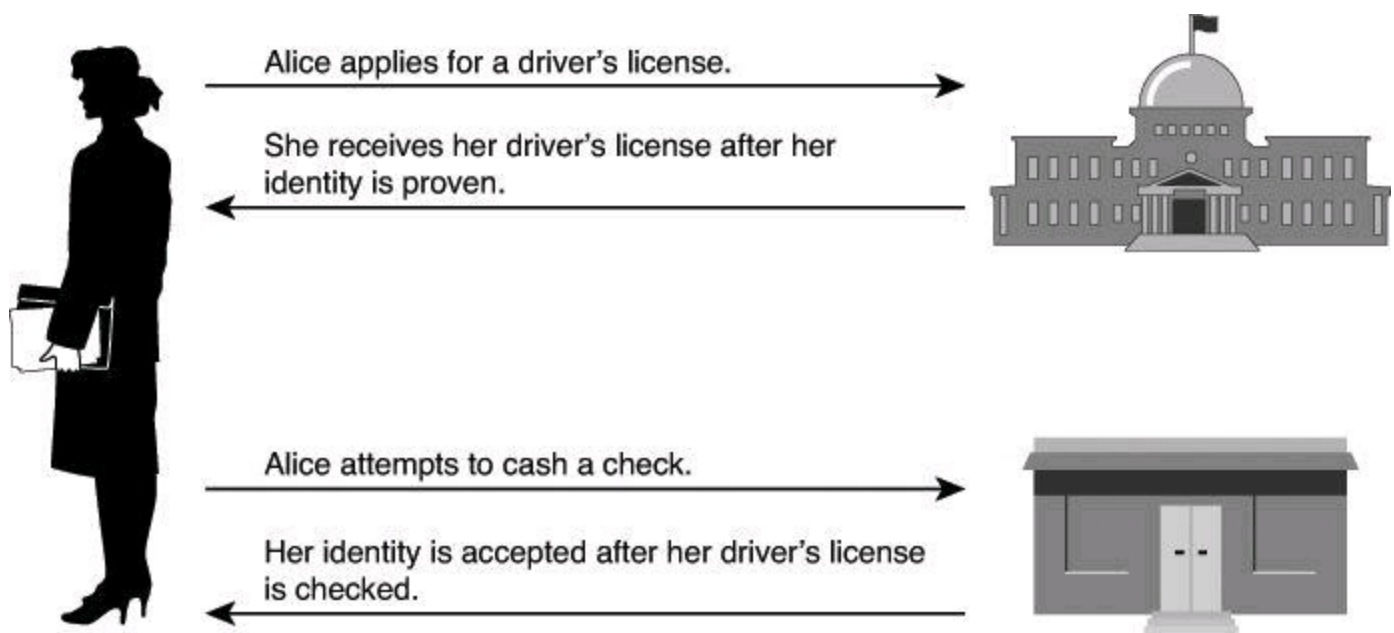


Figure 12-28. Trusted Third Party Example

Later, Alice needs to cash a check at the bank. Upon presenting the check to the bank teller, the bank teller asks her for ID. Alice presents her driver's license. The bank, because it trusts the government agency that issued the driver's license, verifies her identity, and cashes her check.

Note

Certificate servers function like the license bureau in this example. The driver's license is analogous to a certificate in a PKI or a technology that supports certificates.

A PKI provides a framework upon which you can base security services, such as encryption, authentication, and nonrepudiation. A PKI allows for very scalable solutions, and is becoming an

extremely important authentication solution for VPNs. A PKI uses specific terminology to name its components.

PKI Terminology and Components

When these concepts are applied in practice, it is important to understand the supporting framework. A PKI is the service framework that is needed to support large-scale, public-key-based technologies. PKI is a set of technical, organizational, and legal components that are needed to establish a system that enables large-scale use of public-key cryptography to provide authenticity, confidentiality, integrity, and nonrepudiation services.

Three very important terms must be defined when talking about a PKI:

- **PKI:** A service framework needed to support large-scale PK-based technologies
- **Certificate authority (CA):** The trusted third party that signs the public keys of entities in a PKI-based system
- **Certificate:** A document that in essence binds together the name of the entity and its public key and that has been signed by the CA

Note

The certificate of a user is always signed by a CA. Moreover, every CA has its own certificate, containing its public key, signed by itself. This is called a *CA certificate*, or more properly, a *self-signed CA certificate*, or more commonly, a *root certificate*.

A PKI is more than just a CA and its users. And implementing the enabling technology and building a large PKI involves a huge amount of organizational and legal work. There are five main areas of a PKI:



- CAs for key management
- PKI users, such as people, devices, servers, and so on
- Storage and protocols
- Supporting organizational framework, known as practices and user authentication using local registration authorities (LRA)
- Supporting legal framework

Many vendors offer CA servers as a managed service or as an end-user product:

- VeriSign
- Entrust Technologies
- RSA
- Cybertrust
- Microsoft
- Novell

Certificate Classes

CAs, especially outsourced ones, can issue certificates of a number of classes, which determine how trusted a certificate is. A single outsourcing vendor (for example, VeriSign) might run a single CA, issuing certificates of different classes, and its customers will use the CA they need depending on the desired level of trust.

A certificate class is usually a number; the higher the number, the more trusted the certificate is considered. The trust in the certificate is usually determined by how rigorous the procedure was that verified the identity of the holder when the certificate was issued. For example, a class 0 certificate might be issued without any checks, such as for testing purposes. A class 1 certificate might require an email reply from the future certificate holder to confirm his wish to enroll. This confirmation is a weak authentication of the holder. For a class 3 or 4 certificate, the future holder must prove his identity and authenticate his public key by showing up in person with at least two official ID documents.

Certificate Authorities

PKIs form different topologies of trust. In the simple model shown in [Figure 12-29](#), a single CA, which is also known as the root CA, issues all the certificates to the end users. The benefit in such a setup is simplicity, but there are several pitfalls:

- It is difficult to scale this topology to a large environment.
- This topology needs a strictly centralized administration.
- There is a critical vulnerability in using a single-signing private key; if this key is stolen, the whole PKI falls apart because the CA can no longer be trusted as a unique signer.

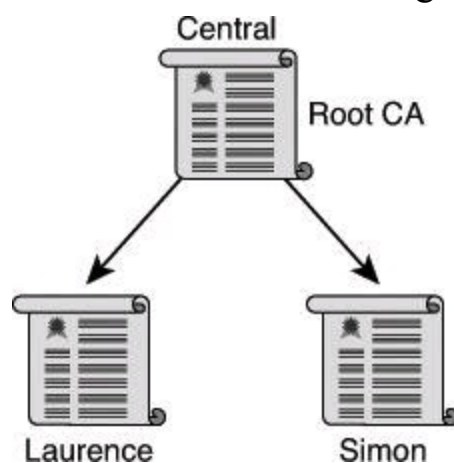


Figure 12-29. PKI Topology Using a Single-Root CA

Because of its simplicity, VPNs that are managed by a single organization often use this topology.

Going beyond the single-root CA, topologies that are more complex can be devised that involve multiple CAs within the same organization. One such topology is the hierarchical CA system, shown in [Figure 12-30](#). With the hierarchical topology, CAs can issue certificates to end users and to subordinate CAs, which in turn issue their certificates to end users, other CAs, or both. Therefore, a tree of CAs and end users is built in which every CA can issue certificates to lower-level CAs and end users.

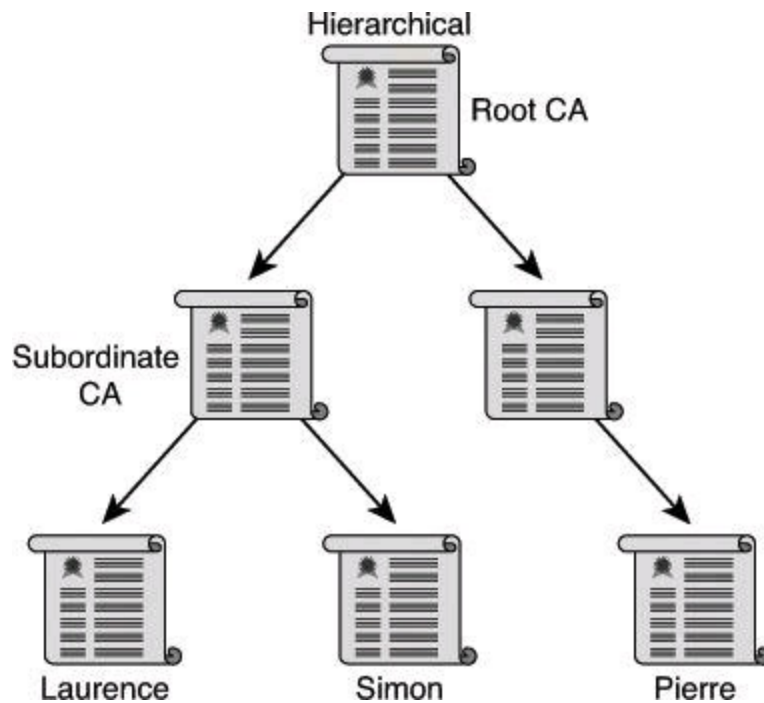


Figure 12-30. PKI Topology Using Hierarchical CAs

The main benefits of a hierarchical PKI topology are increased scalability and manageability; trust decisions can now be hierarchically distributed to smaller branches. This distribution works well in most large organizations. For example, a large company may have a root CA that issues certificates to level-2 CAs. These level-2 CAs issue the certificates to the end users. Because the root-signing key is seldom used after the subordinate CA certificates are issued, it is less exposed and therefore much more trusted. Also, if a subordinate CA has its private key stolen, only a branch of the PKI is rendered untrusted. All other users can consider this by no longer trusting that particular CA.

One issue with hierarchical PKI topologies lies in finding the certification path for a certificate (in other words, determining the chain of the signing process). The more CAs involved in establishing the trust between the root CA and the end user, the more difficult the task.

Another approach to hierarchical PKIs is called cross-certifying, as shown in [Figure 12-31](#). In this scenario, multiple flat single-root CAs establish trust relationships horizontally, by cross-certifying their own CA certificates.

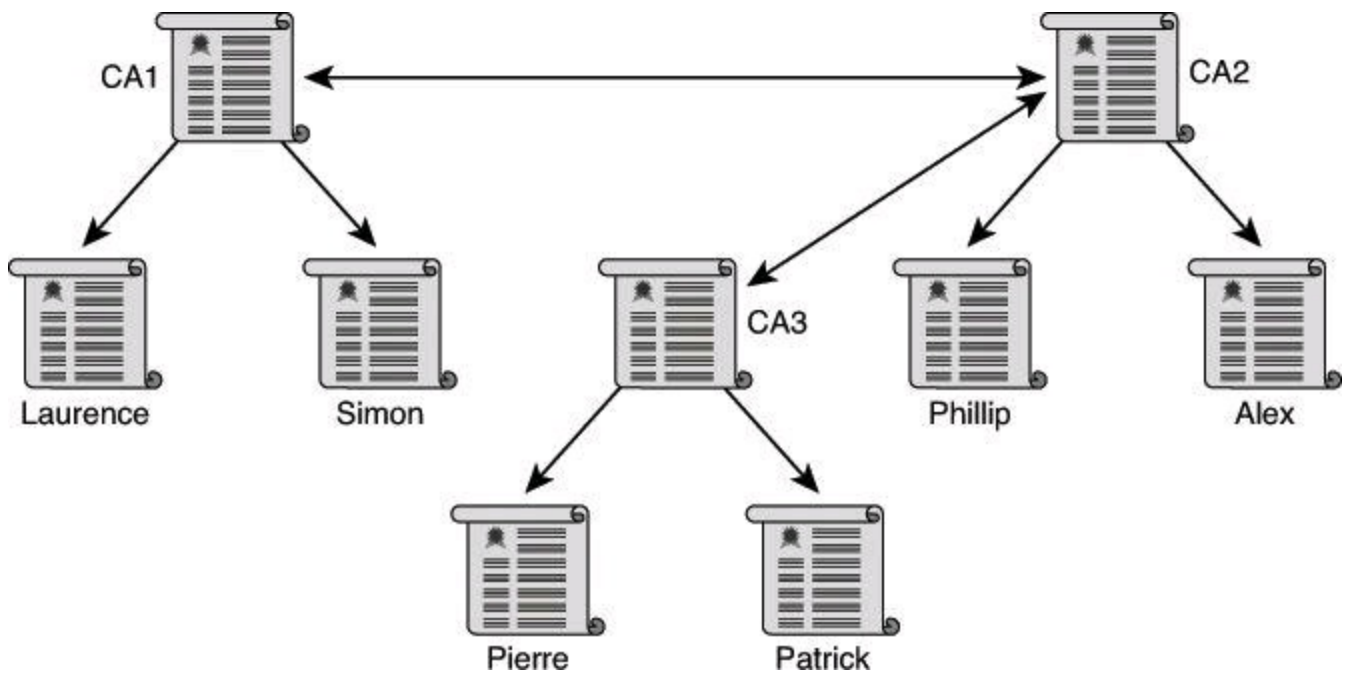


Figure 12-31. PKI Topology Using Cross-Certifying CAs

Some PKIs might offer the possibility or even require the use of two key pairs per entity:

- One public and private key pair is intended only for encryption operations. The public key encrypts, and the private key decrypts.
- The other public and private key pair is intended only for signing operations. The private key signs, and the public key verifies the signature.

These keys are sometimes called “usage” or “special” keys. They may differ in key length and even in the choice of the public-key algorithm. If the PKI requires two key pairs per entity, a user has two certificates:

- An encryption certificate containing the public key of the user who encrypts the data
- A signature certificate containing the public key of the user who verifies the digital signature of the user

The following scenarios typically use usage keys:

- When encryption is used much more frequently than signing, a certain public and private key pair is more exposed due to its frequent usage. In this case, it might be a good idea to shorten its lifetime and change it more often, while having a separate signing private and public key pair with a longer lifetime.
- When different levels of encryption and digital signing are required, because of legal, export, or performance issues, usage keys allow you to assign different key lengths to the two pairs.
- When key recovery is desired (for example, a copy of a user’s private key is kept in a central repository for various backup reasons), usage keys allow you to back up only the private key of the encrypting pair; the signing private key remains with the user, enabling true nonrepudiation.

The CA, with its private key, is the security-critical component in a PKI system. To make the operation of a CA simpler, and therefore more secure, many key management tasks are often

offloaded to registration authorities (RA). RAs are PKI servers that perform management tasks on behalf of the CA, so that the CA can focus on the signing process.

Usually, the following tasks are offloaded to the RA:

- Authentication of users when they enroll with the PKI
- Key generation for users who cannot generate their own keys
- Distribution of certificates after enrollment

PKI Standards

Standardization and interoperability of different PKI vendors are still issues when interconnecting PKIs. The X.509 standards and the Internet Engineering Task Force (IETF) Public-Key Infrastructure X.509 (PKIX) workgroup have made progress toward publishing a common set of standards for PKI protocols and data formats.

A PKI also uses a number of supporting services, such as Lightweight Directory Access Protocol (LDAP)-accessible X.500 directories.

Interoperability between a PKI and its supporting services is a concern because many vendors have proposed and implemented proprietary solutions instead of waiting for standards to develop. The state of interoperability can still be described as basic, even after 10 years of PKI software development.

Note

The IETF has formed a working group that is dedicated to promoting and standardizing PKI in the Internet. The working group has published a draft set of standards detailing common data formats and PKI-related protocols in a network. The goals and milestones of the PKI working group can be tracked at <http://www.ietf.org/html.charters/pkix-charter.html>.

The X.509 standard describes an identity and how to store an authentication key, more precisely defining basic PKI formats, such as certificate and certificate revocation list (CRL) formats, to enable basic interoperability. Abstract Syntax Notation One (ASN.1) provides information about the format of the X.509 certificate and the syntax of the fields in the certificate. The X.509 standard has been widely used for years with many Internet applications, such as SSL and IPsec.

The X.509 Version 3 (X.509v3) standard defines the format of a digital certificate. This format is already extensively used in the infrastructure of the Internet, in the following ways:

- Secure web servers use X.509v3 for website authentication in the SSL and TLS protocols.
- Web browsers use X.509v3 for services that implement client certificates in the SSL protocol.
- User mail agents that support mail protection using the Secure/Multipurpose Internet Mail Extensions (S/MIME) protocol use X.509.
- IPsec VPNs where certificates can be used as a public-key distribution mechanism for IKE RSA-based authentication use X.509.

Certificates are public information. They contain the binding between an entity's names and public keys and are usually published in a centralized directory so that other PKI users can easily access them.

In the CA authentication procedure, the first step of the user, when contacting the PKI, is to securely obtain a copy of the public key of the CA. The public key of the CA verifies all the certificates issued by the CA and is vital for the proper operation of the PKI.

The public key of the CA is also distributed in the form of a certificate issued by the CA itself. This certificate is also called a *self-signed certificate*, because the signer and the holder are the same entity. Only a root CA issues self-signed certificates to itself.

Public-Key Cryptography Standards (PKCS) provide basic interoperability of applications, which use public-key cryptography. The PKCS define the low-level standardized formats for the secure exchange of arbitrary data, such as a standard format for an encrypted piece of data, a signed piece of data, and so on. The origin and purpose of PKCS are described on the RSA Laboratories website: "The Public-Key Cryptography Standards are specifications produced by RSA Laboratories in cooperation with secure systems developers worldwide for the purpose of accelerating the deployment of public-key cryptography."

There are many defined PKCS standards:

- **PKCS #1:** RSA Cryptography Standard
- **PKCS #3:** Diffie-Hellman Key Agreement Standard
- **PKCS #5:** Password-Based Cryptography Standard
- **PKCS #6:** Extended-Certificate Syntax Standard
- **PKCS #7:** Cryptographic Message Syntax Standard
- **PKCS #8:** Private-Key Information Syntax Standard
- **PKCS #9:** Selected Attribute Types
- **PKCS #10:** Certification Request Syntax Standard
- **PKCS #11:** Cryptographic Token Interface Standard
- **PKCS #12:** Personal Information Exchange Syntax Standard
- **PKCS #13:** Elliptic Curve Cryptography Standard
- **PKCS #15:** Cryptographic Token Information Format Standard

Note

For more information about these standards, visit <http://www.rsa.com/rsalabs/node.asp?id=2124>.

Public-key technology is becoming more widely deployed and is becoming the basis for standards-based security, such as the IPsec and IKE protocols. With the use of public-key certificates in network security protocols comes the need for a certificate management protocol that PKI clients and CA servers can use to support certificate lifecycle operations, such as certificate enrollment and revocation, and certificate and CRL access. The goal of the Simple Certificate Enrollment Protocol

(SCEP) is to support the secure issuance of certificates to network devices in a scalable manner, using existing technology wherever possible.

As shown in [Figure 12-32](#), an end entity starts an enrollment transaction by creating a certificate signing request using PKCS#10. The PKCS#10 signing request is encapsulated in a PKCS#7 and sent to the CA or RA. PKCS#7 is a syntax standard for cryptographic messages. After the CA or RA receives the request, it either automatically approves the request and sends the certificate back, or it compels the end entity to wait until the operator can manually authenticate the identity of the requesting end entity.

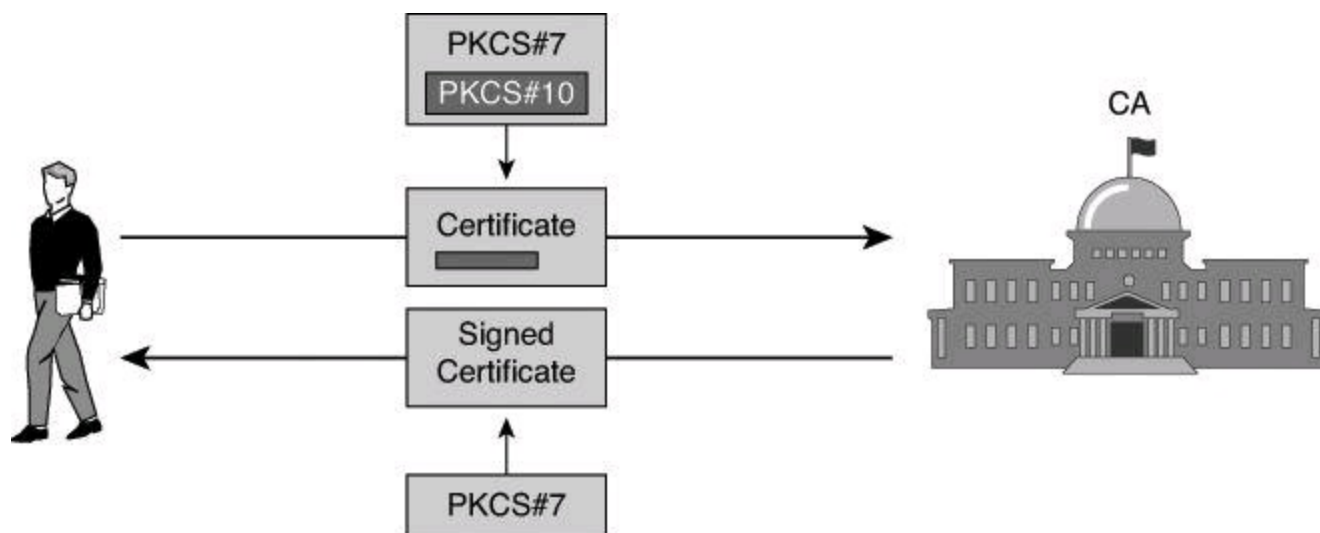


Figure 12-32. Certificate Signing Request

Distinguished names (DN) provide a way to identify an entity by using multiple fields to provide hierarchical identification. DNs are used on a certificate in the Subject field.

Note

An example of a distinguished name used as the Subject Name field in an X.509 user certificate would appear as

CN=Harry Wales,OU=Sales,O=My Computer,L=My Company,L=Chicago,S=Ohio,C=US

where

- CN= commonName
- OU= organizationalUnitName
- O= Organization
- L= Locality (City)
- S= State
- C= US

The merging of the X.509 standard with public-key encryption allows the introduction of a trusted third party: the CA. The CA has a pair of asymmetric keys, a private key, and a public key. An X.509 certificate is created to identify the CA. The certificate of the CA contains the following information:

- The identity of the CA (for example, a subject containing the identity in the DN format)

- Other parameters (such as serial number, algorithms used, and validity period)
- The public key of the CA (for example, an RSA public key)
- The signature using the private key of the CA (for example, self-signing using the private key of the CA with RSA encryption and the SHA-1 hash algorithm)

Caution

The certificate is freely distributed. The receiver of the certificate should verify the authenticity of the certificate of the CA out-of-band.

Also, browsers such as Microsoft Internet Explorer come with certificates of large CAs, such as VeriSign, GoDaddy, and so on, already provisioned in the installation. That's how our personal computers are able to establish HTTPS sessions with banks and Internet retailers without having to install additional software.

In [Figure 12-33](#), the following steps occur to retrieve the CA certificate:

Step 1. Alice and Bob request the CA certificate that contains the CA public key.

Step 2. Upon receipt of the CA certificate, Alice's and Bob's systems verify the validity of the certificate using public-key cryptography.

Step 3. Alice and Bob follow up the technical verification done by their systems by telephoning the CA administrator and verifying the public key and serial number of the certificate.

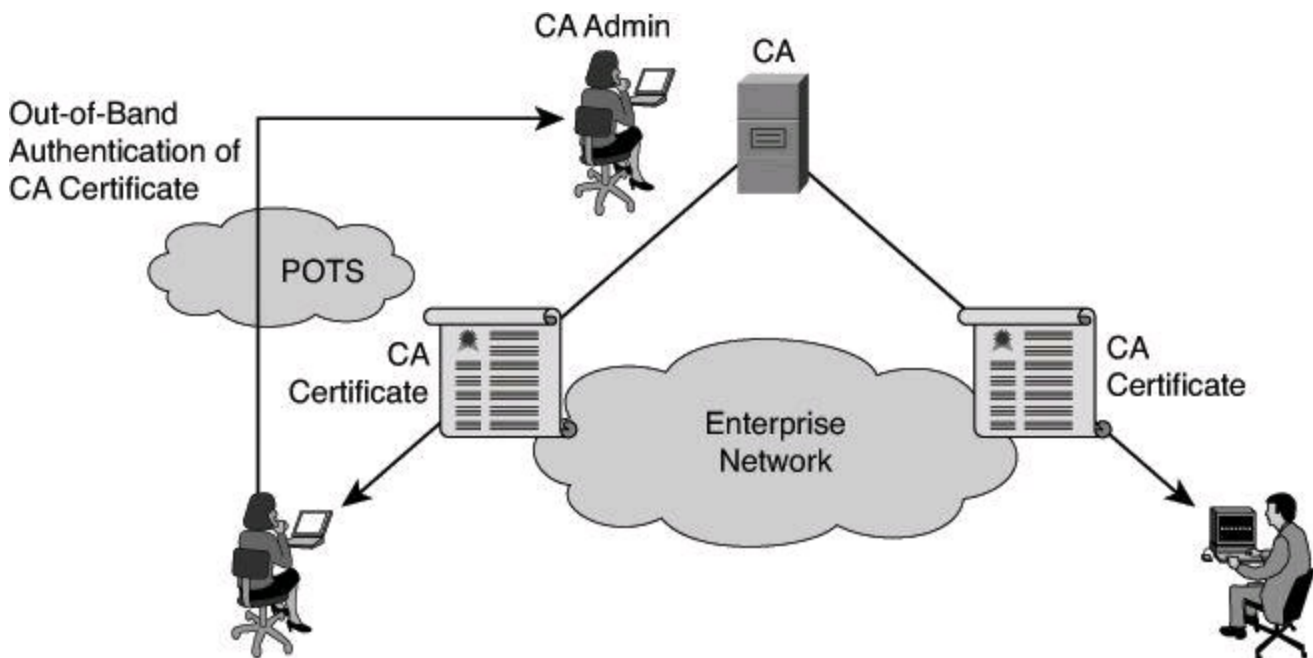


Figure 12-33. Retrieving a CA Certificate

After retrieving the CA certificate, Alice and Bob perform the following steps to submit certificate requests to the CA, as shown in [Figure 12-34](#):

Step 1. Alice's and Bob's systems forward a certificate request that includes their public keys along with some identifying information. All of this information is encrypted using the public key of the CA.

Step 2. Upon receipt of the certificate requests, the CA administrator telephones Alice and Bob to confirm their submittals and the public keys.

Step 3. The CA administrator issues the certificate by adding some additional data to the certificate request, and digitally signing it all.

Step 4. Either the end user manually retrieves the certificate or SCEP automatically retrieves the certificate, and the certificate is installed onto the system.

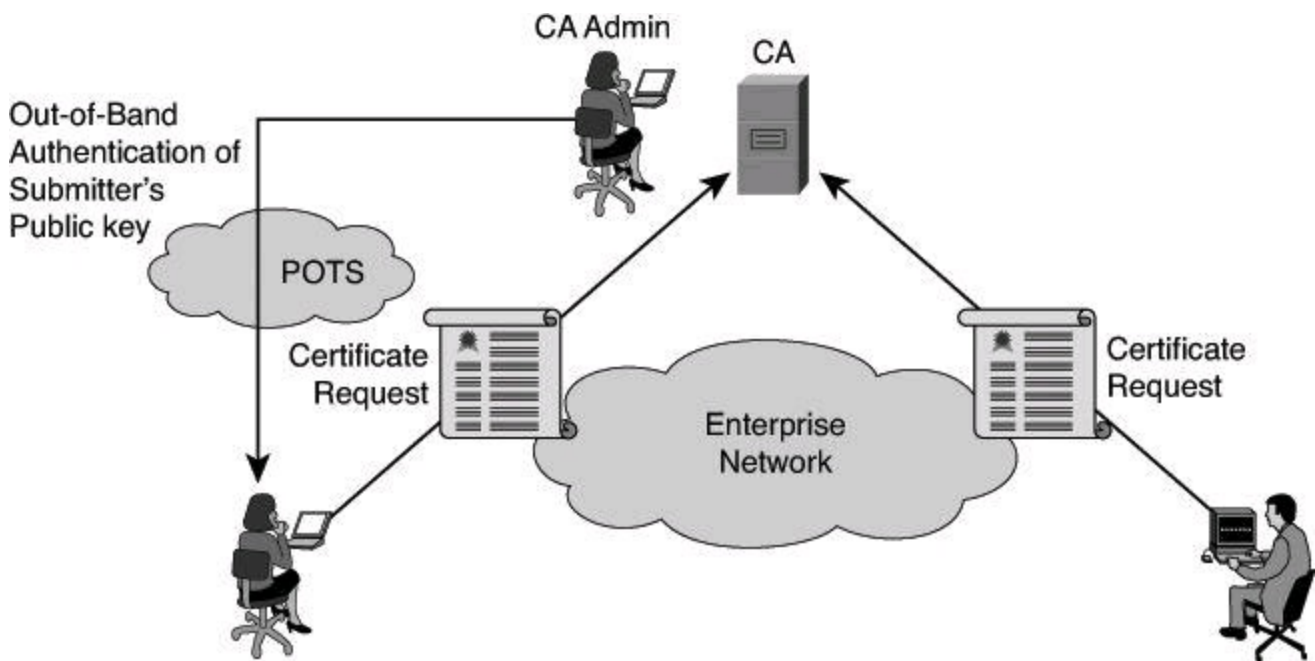


Figure 12-34. Certificate Enrollment

Having installed certificates signed by the same CA, Bob and Alice are now ready to authenticate each other, as shown in [Figure 12-35](#):

Step 1. Bob and Alice exchange certificates. The CA is no longer involved.

Step 2. Each party verifies the digital signature on the certificate by hashing the plaintext portion of the certificate, decrypting the digital signature using the CA public key, and comparing the results. If the results match, the certificate is verified as being signed by a trusted third party, and the verification by the CA that Bob is Bob and Alice is Alice is accepted.

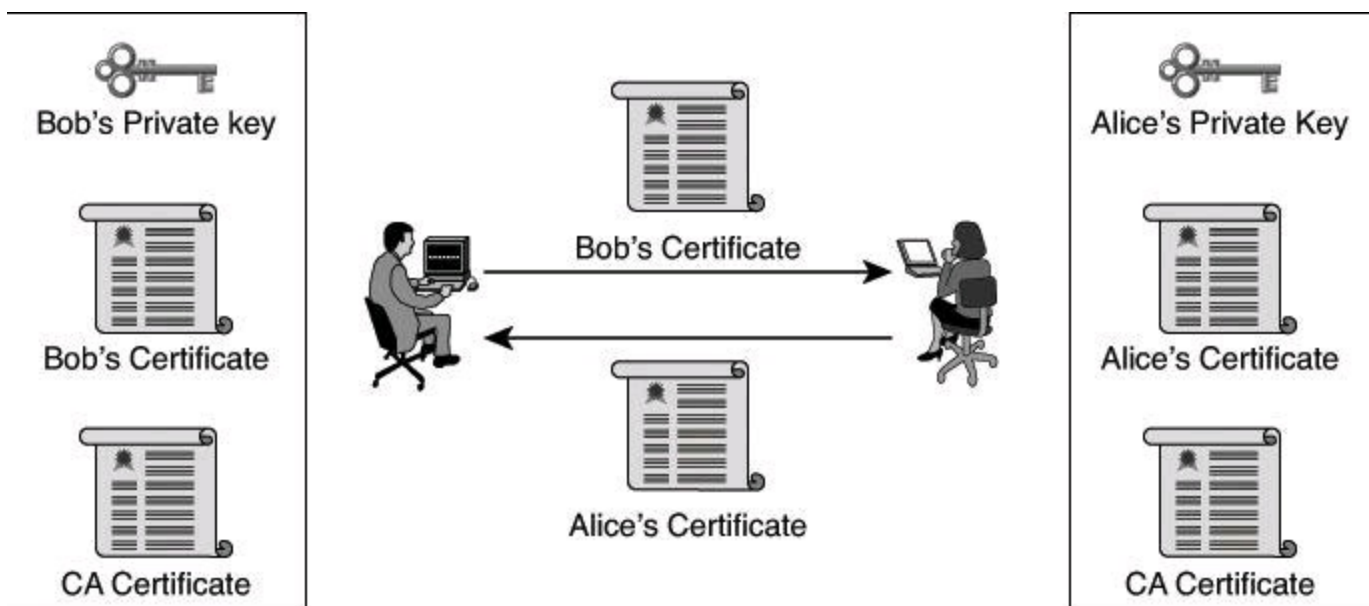


Figure 12-35. Authentication Using Certificates

Certificate Revocation

Digital certificates can be revoked if keys are thought to be compromised, or if the business use of the certificate calls for revocation (for example, VPN access, network logon, and so on). If keys are thought to be compromised, generating new keys will force the creation of a new digital certificate, rendering the old certificate invalid and a candidate for revocation. On the other hand, a consultant could obtain a digital certificate for VPN access into the corporate network only for the duration of the contract.

Certificate revocation is also a centralized function, providing “push” and “pull” methods to obtain a list of revoked certificates frequently or on demand from a centralized entity. In some instances, the CA server often acts as the issuer of certificate revocation information, as shown in [Figure 12-36](#).

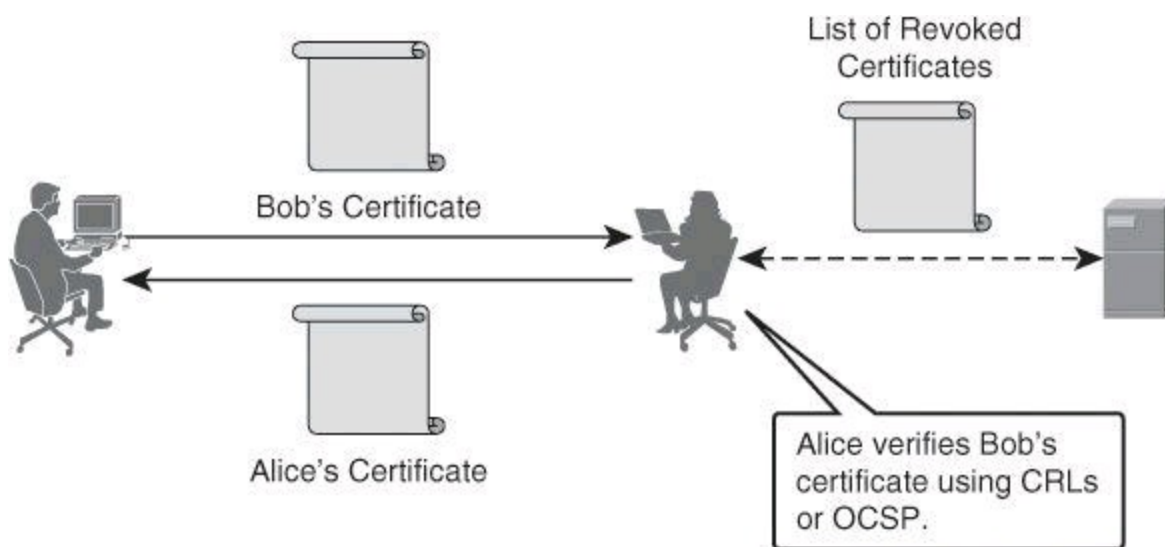


Figure 12-36. Certificate Revocation Process

Several methods can be used for certificate revocation. The most prevalent today is the use of CRLs, but other methods such as Online Certificate Status Protocol (OCSP) overcome some of the limitations of CRLs and are often used. An alternative method is Cisco-proprietary and involves querying the AAA server to see if the user exists. If the user exists, the entity checking the validity of the certificate presumes that the certificate is good because the user is valid in the AAA server. [Table 12-5](#) lists benefits and limitations of each method.

Table 12-5. Certificate Revocation Methods

Certificate Revocation List	Online Certificate Status Protocol	AAA Server Certificate Authorization
A list of revoked certificate serial numbers is distributed as a time-stamped, CA-signed file.	Revocation information is immediately pushed to an online database.	Cisco-proprietary alternative to OCSP.
PKI entities regularly poll the CRL repository to pull the current CRL. Clients cache a local copy of CRLs. The client downloads a newer CRL once the current one expires.	Entities can query the OCSP server at any time to check for validity of the received certificate.	Entities can query the AAA server at any time to check for validity of the received certificate.
There is a window of opportunity for attackers while the new CRL is not yet propagated. This is why it is recommended to have an overlapping period between the two CRLs.	This protocol is not widely deployed.	It is not integrated with the PKI, so it requires a separate authorization database.

Certificates were traditionally used at the application layer to provide strong authentication for applications. Each application may have a different implementation of the actual authentication process, but they all use a similar type of certificate in the X.509 format.

Certificate Use

SSL is probably the most widely used certificate-based authentication. SSL includes the negotiation of keys that are used to encrypt the SSL session. Many applications use SSL to provide authentication and encryption. The most widely used application is HTTPS. Other well-known applications that were using poor authentication and no encryption were modified to use SSL, such as Simple Mail Transfer Protocol (SMTP), LDAP, and Post Office Protocol version 3 (POP3).

Email has experienced many extensions. One of the important extensions was the introduction of Multipurpose Internet Mail Extensions (MIME), which allowed arbitrary data to be included in an email. Another extension was to provide security to entire mail messages or parts of mail messages. S/MIME authenticates and encrypts email messages.

Pretty Good Privacy (PGP) is an application that was originally developed by Phil Zimmerman, a privacy advocate, so that end users could engage in confidential communications using encryption. The most frequent use of PGP has been to secure email.

Certificates are also used at the network layer, or at the application layer, by network devices. Cisco routers, Cisco VPN concentrators, and Cisco ASA firewalls can use certificates to authenticate IPsec peers.

Cisco switches can use certificates to authenticate end devices connecting to LAN ports. Authentication uses IEEE 802.1X between the adjacent devices. The authentication can be proxied to a central access control server (ACS) via EAP-TLS.

Cisco routers can also provide Telnet 3270 support that does not include encryption or strong

authentication. Cisco routers can now use SSL to establish Secure Telnet 3270 sessions.

[Figure 12-37](#) illustrates a network where certificates are used for various purposes. A single CA server can facilitate many applications that require digital certificates for authentication purposes. Using CA servers is therefore a solution that simplifies the management of authentication and provides strong security due to the strength of cryptographic mechanisms that are used in combination with digital certificates.

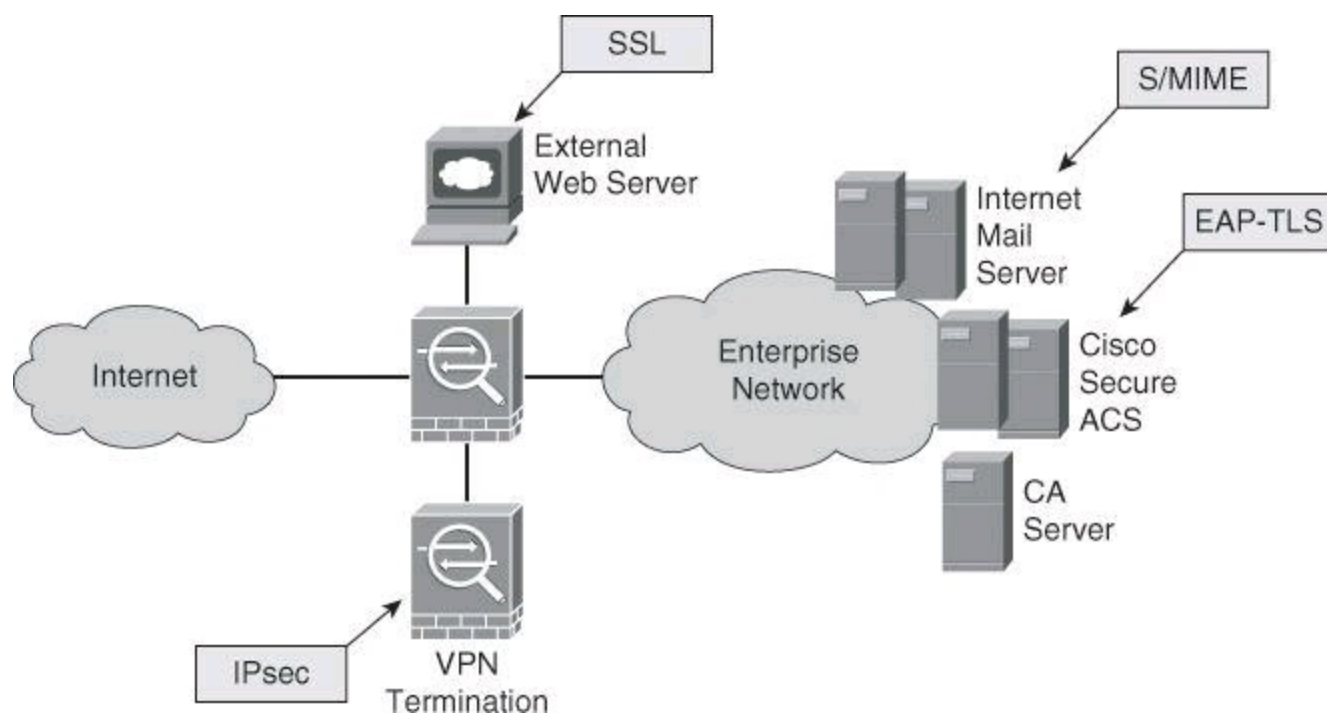


Figure 12-37. Where We Find Certificates Being Used

Digital Certificates and CAs

Compared to other authentication mechanisms, a PKI has the following characteristics:

- To authenticate each other, users have to obtain the certificate of the CA and their own certificate. These steps require the out-of-band verification of the processes. After this verification is complete, the presence of the CA is no longer required until one of the certificates that is involved expires.
- Public-key systems use asymmetric keys where one is public and the other one is private. One of the features of these algorithms is that whatever is encrypted using one key can only be decrypted using the other key. This provides nonrepudiation.
- Key management is simplified because two users can freely exchange the certificates. The validity of the received certificates is verified using the public key of the CA, which the users have in their possession.
- Because of the strength of the algorithms involved, you can set a very long lifetime for the certificates, typically a lifetime measured in years.

The disadvantages of using trusted third parties relate to key management:

- **A user certificate is compromised (stolen private key):** Other users should not accept compromised certificates. The only way to prevent the compromised certificates from being used is to keep a list of all revoked certificates. A server, not necessarily the CA

server, must be accessible to users so that they can periodically download the latest CRL and use the list when authenticating other users. If the CRL lists the received certificate of the user, the authentication fails.

- **The certificate of the CA is compromised (stolen private key):** This invalidates all certificates signed by the CA. A single CA environment requires the creation of a new CA certificate and the creation of new user certificates. A hierarchical CA environment requires the use of an authority revocation list (ARL), where all child certificates of the compromised CA become invalid if the ARL lists the CA.
- **The CA administrator:** The human factor is an additional limitation of the CA-based solution. To lessen the impact, the CA administrator should follow strict rules when issuing certificates to users. A security policy should define the steps required to create certificates (for example, mandatory out-of-band verification of all initial enrollment procedures or verification steps for CA administrators before approving certificate requests).

When you use certificates in IP networks, you might need to combine public-key authentication with another authentication mechanism to increase the level of security and provide more authorization options. For example, IPsec using certificates for authentication and Extended Authentication (XAUTH) with one-time password hardware tokens would be a superior authentication scheme when compared to certificates alone. XAUTH is used for user authentication during IPsec remote access VPN, while the digital certificate is used typically for asset authentication (machine authentication).

Summary

The key points covered in this chapter are as follows:

- A cryptosystem is made up of a combination of hashing, symmetric, and asymmetric algorithms.
- Symmetric algorithms use a single key for encrypting and decrypting. Generally speaking, symmetric algorithms are the strongest and fastest algorithms and therefore are used for most encryption.
- Hashing algorithms use a one-way process designed to provide integrity. Usually, successful decryption of a digest provides proof of integrity and authenticity.
- Asymmetric algorithms use a key pair for the encrypting/decrypting process. One key encrypts, and the other key decrypts.
- RSA is a widely used algorithm for public-key cryptography.
- A PKI uses asymmetric encryption to provide confidentiality, integrity, and authentication services.
- PKI solutions are based on digital certificates and a trusted third party trust model.
- X.509v3, PKCS, and others provide standards for certificate formats and interoperability.
- CRL, OCSP, and AAA server certificate authorization are means to validate a certificate.
- The hierarchical trust model of PKI solutions includes CAs and RAs.

References

For additional information, refer to these resources.

Books and Articles

- Giry, D. "Cryptographic Key Length Recommendation," <http://www.keylength.com/en/3/>
- Kahn, D. *The Codebreakers: The Comprehensive History of Secret Communication from Ancient Times to the Internet* (Scribner, 1996).
- Kaliski, B. "TWIRL and RSA Key Size," <http://www.rsasecurity.com/rsalabs/node.asp?id=2004>.
- Singh, S. *The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography* (Knopf Publishing Group, 1999).

Standards

- IETF. *Public-Key Infrastructure (X.509) (pkix)*, <http://tools.ietf.org/html/rfc5280>
- NIST. "Advanced Encryption Standard (AES)" (FIPS PUB 197), <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>.
- NIST. "Secure Hash Standard" (FIPS PUB 180-1), <http://www.itl.nist.gov/fipspubs/fip180-1.htm>.
- RSA Laboratories. *Public-Key Cryptography Standards (PKCS)*, <http://www.rsa.com/rsalabs/node.asp?id=2124>.
- Wikipedia. "Diffie-Hellman," http://en.wikipedia.org/wiki/Diffie%E2%80%93Hellman_key_exchange.

Encryption Regulations

- U.S. Department of Commerce, <http://www.commerce.gov>.
- "EAR Controls for Items That Use Encryption," <http://www.bis.doc.gov/encryption/default.htm>.

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, "[Answers to Chapter Review Questions](#)."

1. Which three business benefits are directly associated with VPNs?
 - a. Cost savings
 - b. Traffic segregation
 - c. Telemetry
 - d. Scalability
 - e. Productivity
2. Asymmetric encryption utilizes which of the following?
 - a. A complex solution to leverage certificates
 - b. Key pairs to accomplish encryption
 - c. All types of authentication as long as they support PKI procedures
 - d. A certificate authority that signs public keys in a network
3. Select all the desirable features of an encryption algorithm.

- a. Resistance to known cryptanalytic attack
 - b. Short lengths and scalability of the key for easy storage
 - c. No export or import restrictions
 - d. Susceptible to the avalanche effect; a small change in plaintext causes substantial changes in ciphertext
4. A cipher suite is _____.
- a. The place of storage of cipher keys
 - b. A combination of cryptographic methods that work together
 - c. A chain of crypto keys
 - d. The prevalent set of algorithms at a given point in time
5. The selection of asymmetric encryption algorithms is based on a tradeoff between security and _____.
- a. Operational efficiency
 - b. Reliability
 - c. Performance
 - d. Key length
6. RC4 is a(n) _____ encryption algorithm.
- a. Asymmetric
 - b. Polyalphabetic
 - c. Elliptic curve
 - d. Symmetric
7. Which two options determine the strength of protection of modern cryptography algorithms?
- a. Length of key
 - b. Mathematical complexity
 - c. Number of permutations
 - d. Level of trust toward the algorithm itself
 - e. Resistance to brute force attacks
8. Which combination does asymmetric encryption use in confidentiality scenarios?
- a. The private key is used to encrypt and the public key is used to decrypt.
 - b. The public key is used to encrypt and the private key is used to decrypt.
 - c. A digital certificate is used to encrypt and the private key is used to decrypt.
 - d. The public key is used to encrypt and the corresponding digital certificate is used to decrypt.
9. Which of the following is not involved in RSA user authentication?
- a. A hashing function
 - b. A public key

- c. Digital signatures
- d. A shared secret

10. Which of the following cipher categories transforms plaintext into ciphertext by operating bit by bit?

- a. Transposition
- b. Stream
- c. Polyalphabetic
- d. Block

11. DES operates in which two block cipher modes?

- a. ECB
- b. CFB
- c. CBC
- d. OFB

12. Which algorithm is used in AES?

- a. Twofish
- b. RC6
- c. RC4
- d. Rijndael

13. Which statement best describes MD5?

- a. MD5 is a one-way function that makes it difficult to compute a hash from the given input data, but makes it feasible to compute input data given only a hash.
- b. MD5 is a one-way function that makes it easy to compute a hash from the given output data, but makes it infeasible to compute input data given only a hash.
- c. MD5 is a one-way function that makes it difficult to compute a hash from the given output data, but makes it feasible to compute input data given only a hash.
- d. MD5 is a one-way function that makes it easy to compute a hash from the given input data, but makes it infeasible to compute the exact input data given only a hash.

14. Which of the following statements regarding public-key authentication is true?

- a. When the private key is used to encrypt, the corresponding public key is used to decrypt.
- b. Because the public key is present on only one system, authentication is assured when its private key decrypts the message.
- c. Great effort is made to maintain the secrecy of the public keys.
- d. Public-key scenario is used for producing fingerprint.

15. Which set of algorithms provides the most secure communication?

- a. AES, SHA-1
- b. 3DES, SHA-1

c. 3DES, MD5

d. AES, MD5

16. Which of the following statements best describes a digital signature?

a. A digital signature is a message digest encrypted with the sender's public key.

b. A digital signature is a message digest encrypted with the receiver's public key.

c. A digital signature is a message digest encrypted with the sender's private key.

d. A digital signature is a message digest encrypted with the receiver's public key.

17. Complete the sentence with the best statement: The Vigenère cipher is a(n) _____.

a. Polyalphabetic cipher

b. Polymorphic cipher

c. Polybius square cipher

d. Alphabetum cipher

18. Which of the following is not a term that is related to hierarchical CA topologies?

a. Certificate chain

b. Subordinate CAs

c. PKI

d. Single point of failure

19. Which protocol encrypts at the session layer of the OSI model?

a. IPsec

b. Enigma

c. SSL

d. MD5

20. Which statement is most accurate when describing aspects of a birthday attack?

a. An attacker tries every possible key with the decryption algorithm.

b. The attacker has the ciphertext of several messages, all of which have been encrypted using the same encryption algorithm.

c. If some function, when supplied with a random input, returns one of k equally likely values, then by repeating the function with different inputs, the same output would be expected after $1.2k^{1/2}$ number of times.

d. The attacker knows a portion of the plaintext and the corresponding ciphertext.

21. Which of the following terms identifies a standard that describes the structure of digital certificates?

a. PKCS#10

b. SCEP

c. X.509v3

d. RSA

22. Out-of-band authentication, in the context of certificate authorities, typically involves which of the following? (Choose two.)

- a.** Certificate request
- b.** Separate network
- c.** Digital signature
- d.** Public key
- e.** Serial number

Chapter 13. IPsec Fundamentals

This chapter addresses the protocols and algorithms that IPsec uses and the different security services that IPsec provides. More precisely, this chapter

- Analyzes the architecture of the IPsec protocol
- Details the role and operational impact of IPsec's main components
- Describes IPsec modes of operation in various scenarios
- Describes the phases of IPsec connectivity
- Describes the role and component of IKE
- Provides an overview of the operations of IPv6 VPNs

An IP Security (IPsec) virtual private network (VPN) is an essential tool for providing a secure network for business communication. This chapter addresses the different protocols and algorithms that IPsec uses and the different security services that IPsec provides. This chapter also introduces the different VPN technologies and some of the best practices that you should use with them.

IPsec Framework

IPsec is an Internet Engineering Task Force (IETF) standard that defines how a VPN can be set up using the IP addressing protocol; it was originally defined in RFCs 2401 to 2412. IPsec works at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers). IPsec is not bound to any specific encryption, authentication, or security algorithms or keying technology. IPsec is a framework of open standards.

Because IPsec is not bound to specific algorithms, IPsec allows newer and better algorithms to be implemented without patching the existing IPsec standards. IPsec provides data confidentiality, data integrity, origin authentication, and antireplay protection between participating peers at the IP layer. IPsec secures a path between a pair of gateways or between a gateway and a host, as shown in [Figure 12-2](#) in [Chapter 12](#), “[Fundamentals of Cryptography and VPN Technologies](#).” IPsec could also secure a path between two hosts.

IPsec provides the following essential security functions:



- **Confidentiality:** IPsec ensures confidentiality by using encryption. Data encryption prevents third parties from reading the data, especially data transmitted over public networks or over wireless networks.
- **Integrity:** IPsec ensures that data arrives unchanged at the destination; that is, that the data has not been manipulated at any point along the communication path. IPsec ensures data integrity by using checksums, which are a simple redundancy check. The IPsec protocol adds up the basic components of a message, typically the number of bytes, and stores the total value. IPsec performs a checksum operation on received data and compares the result to the authentic checksum. If the sums match, the assumption is that the data has not been

manipulated.

• **Authentication:** Authentication ensures that the connection is made with the desired communication partner. IPsec uses Internet Key Exchange (IKE) to authenticate users and devices that can carry out communication independently. IKE uses several types of authentication:

- Username and password
- One-time password
- Biometric
- Preshared keys (PSK)
- Digital certificates

Note

IKE is discussed in more detail in the section, “[IKE Protocol](#),” later in the chapter.

• **Antireplay protection:** Antireplay protection verifies that each packet is unique and is not duplicated. IPsec packets are protected by comparing the sequence number of the received packets with a sliding window on the destination host or security gateway. A packet that has a sequence number that comes before the sliding window is considered either late or a duplicate packet. Late and duplicate packets are dropped.

Essential services provided by IPsec are as follows:

- Confidentiality
- Integrity
- Authentication
- Antireplay

The main reason we use IPsec is typically to provide confidentiality to our transmission. Confidentiality is reached by encrypting sensitive payload. The following explains the recent improvement to encryption.

Suite B Cryptographic Standard

In most cryptographic functions, the key length is an important security parameter. Both academic and private organizations provide recommendations and mathematical formulas to approximate the minimum key size requirement for security. Despite the availability of these publications, choosing an appropriate key size to protect systems and networks from attacks remains a challenge.

As computing power increases, the feasibility of deploying complex encryption algorithms commercially also increases. The main principle is that the degree of security depends on the length of the key of the encryption algorithm. The time that it takes to process all of the possibilities is a function of the computing power of the computer. Therefore, the shorter the key, the easier it is to break the key.

In 2005, the U.S. National Security Agency (NSA) identified a set of cryptographic algorithms that, when used together, provides the preferred method for ensuring the security and integrity of

information that is passed over public networks such as the Internet. The NSA called the set of algorithms “Suite B.” Today, Suite B is globally recognized as an advanced, publicly available standard for cryptography. Suite B provides a security level of 128 bits or higher, significantly higher than many commonly used standards.

Integrated into IETF standards, Suite B algorithms make it easier to collaborate in environments where costs or logistics traditionally hindered information sharing. Secure sharing of information over the Internet and other nontrusted networks supports various missions at all levels of government. Examples include intelligence agencies and the military in a government context, or private sector companies that are required to increase the security of transmitting sensitive content such as intellectual property or private customer and employee information.

Another advantage of Suite B is that it helps public and private sector organizations meet compliance requirements. These requirements include compliance with Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), Federal Information Processing Standards (FIPS), and others.

As described in RFC 4869, “Suite B Cryptographic Suites for IPsec,” Suite B cryptography secures information traveling over networks using four well-established, public-domain cryptographic algorithms:

- Encryption that is based on AES using 128- or 256-bit keys
- Digital signatures with the Elliptic Curve Digital Signature Algorithm (ECDSA) using curves with 256- and 384-bit prime moduli
- Key exchange, either preshared or dynamic, using the Elliptic Curve Diffie-Hellman (ECDH) method.
- Hashing (digital fingerprinting) based on Secure Hash Algorithm 2 (SHA-2)

The NSA has stated that these four algorithms in combination provide adequate information assurance for classified information. RFC 4869 has gained acceptance in the industry.

Note

In December 2006, NSA submitted an Internet Draft on implementing Suite B for IPsec, which was accepted for publication by the IETF as RFC 4869. However, RFC 4869 was later obsoleted by RFC 6379 in October 2011. The curriculum of IINSv2 exam 640-554 was developed with RFC 4869 in mind.

Also note that with the use of 384-bit elliptic curve, SHA-384 and AES with 256-bit keys are necessary for the protection of top secret information.

Commercial Suite B devices do not require the special handling requirements that are traditionally associated with government-specific cryptographic devices. The less stringent requirements simplify adoption, strengthen the overall architecture security, and minimize operational costs.

Essential services provided by the Suite B cryptographic standard include the following:

- Confidentiality = AES
- Integrity = SHA-2
- Authentication = ECDSA
- Key management = ECDH

Encryption Algorithms

[Chapter 12](#) provided an in-depth analysis of the different encryption technologies. You saw that the degree of security depends on the length of the key and on the encryption algorithm. The time that it takes to process all the possibilities is a function of the computing power of the computer. Therefore, the shorter the key, the easier it is to break the key.

The following are some of the encryption algorithms and key lengths that VPNs use:

- **Date Encryption Standard (DES):** DES was developed by IBM. DES uses a 56-bit key, ensuring high-performance encryption. DES is a symmetric key cryptosystem.
- **3DES:** The 3DES algorithm is a variant of the 56-bit DES. 3DES operates in a way that is similar to how DES operates, in that data is broken into 64-bit blocks. 3DES then processes each block three times, each time with an independent 56-bit key. 3DES provides significant encryption strength over 56-bit DES. 3DES is a symmetric key cryptosystem.
- **Advanced Encryption Standard (AES):** The National Institute of Standards and Technology (NIST) has recently adopted AES to replace the existing DES encryption in cryptographic devices. AES provides stronger security than DES and is computationally more efficient than 3DES. AES offers three different key lengths: 128-, 192-, and 256-bit keys. AES is the standard prescribed in RFC 4869, “Suite B Cryptographic Suites for IPsec.”
- **Rivest, Shamir, and Adleman (RSA):** RSA is an asymmetrical key cryptosystem. It uses a key length of 512, 768, 1024, or larger. IPsec does not use RSA for data encryption. IKE uses RSA encryption only during the peer-authentication phase.
- **Software-Optimized Encryption Algorithm (SEAL) algorithm:** SEAL is a stream cipher, developed in 1993 by Phillip Rogaway and Don Coppersmith, and uses a 160-bit key for encryption.

Key Exchange: Diffie-Hellman

Encryption algorithms, such as DES and 3DES, explained in [Chapter 12](#), require a symmetric shared-secret key to perform encryption and decryption. You can use email, courier, or overnight express to send the shared-secret keys to the administrators of the devices. But the easiest key-exchange method is a public-key exchange method between the encrypting and decrypting devices. The method has two variants:

- The Diffie-Hellman (DH) key agreement, explained in great detail in [Chapter 12](#), is a cryptographic protocol that provides a way for two peers to establish a shared-secret key,

which only they know, even though they are communicating over an unsecure channel.

- ECDH is a variant of the DH protocol using elliptic curve cryptography (ECC). It is part of the Suite B standards.

That shared-secret key, created by the Diffie-Hellman key exchange method, is then used as the encryption key to exchange data.

These algorithms are used within IKE to establish session keys. They support different key sizes that are identified by different DH or ECDH groups. DH groups determine the strength of the key that is used in the key exchange process. Higher group numbers are more secure, but they require additional time to compute the key. The DH groups are referred to as DHx; for example, Diffie-Hellman Group 1 is referred to as DH1. The following is a quick listing of the different DH groups (which are covered in more detail later in this chapter, in the “[IKE Protocol](#)” section):

- DH1: 768-bit key
- DH2: 1024-bit key
- DH5: 1536-bit key
- DH7: 163-bit ECDH key
- DH14: 2048-bit key
- DH15: 3072-bit key
- DH16: 4096-bit key
- DH19: 256-bit ECDH key
- DH20: 384-bit ECDH key
- DH24: 2048-bit ECDH key

Data Integrity

VPN data is typically transported over the public Internet. Potentially, this data could be intercepted and modified. To guard against this problem, you can use a data-integrity algorithm, explained in [Chapter 12](#). A data-integrity algorithm adds a hash to the message, which guarantees the integrity of the original message. If the transmitted hash matches the received hash, the message has not been tampered with. However, if there is no match, the message was altered.

A Hashed Message Authentication Code (HMAC) is a data-integrity algorithm that guarantees the integrity of the message. At the local end, the message and a shared-secret key are sent through a hash algorithm, which produces a hash value. The message and hash are sent over the network. Refer to [Figure 12-18](#) in [Chapter 12](#) for a detailed presentation on HMAC.

General HMAC concepts were explained in [Chapter 12](#). Now, let’s look at three common HMAC algorithms:

- **HMAC-Message Digest 5 (HMAC-MD5):** HMAC-MD5 uses a 128-bit shared-secret key of any size. The variable-length message and shared-secret key are combined and run through the HMAC-MD5 hash algorithm. The output is a 128-bit hash. The hash is appended to the original message and is forwarded to the remote end.
- **HMAC-Secure Hash Algorithm 1 (HMAC-SHA-1):** HMAC-SHA-1 uses a secret key of any size. The variable-length message and the shared-secret key are combined and run

through the HMAC-SHA-1 hash algorithm. The output is a 160-bit hash. The hash is appended to the original message and is forwarded to the remote end.

- **HMAC-Secure Hash Algorithm 2 (HMAC-SHA-2):** The SHA-2 family of HMACs is based on the same base algorithm as SHA-1. The SHA-2 family (the second generation of SHA algorithms) includes the algorithms of 256 and 384 bit, referred to as SHA-256 bit hash algorithm and SHA-384 bit hash algorithm. This functionality is part of the Suite B standard.

Authentication

When you are conducting business long distance, it is necessary to know who is at the other end of the phone, email, or fax. The same is true of VPN networks. The device on the other end of the VPN tunnel must be authenticated before the communication path is considered secure. Four peer-authentication methods exist:

- **Preshared keys:** A secret key value is entered into each peer manually and is used to authenticate the peer. This is a shared secret that both parties must exchange ahead of time. Think of it as a secret password that they offer to each other to confirm the identity of the other party. At each end, the PSK is combined with other information to form the authentication key.
- **RSA signatures:** In addition to sending its digital certificate to the remote end, the sender also includes the hash value of a message encrypted with its private key. This acts like a signature. At the remote end, once the receiver has validated the sender's digital certificate, the encrypted hash is decrypted using the public key of the sender, found in the sender's digital certificate. If the decrypted hash matches the recomputed hash, the signature is genuine.
- **RSA encrypted nonces:** A *nonce* is a random number generated by the peer. It is similar to a nonce word, used only once. RSA encrypted nonces use RSA to encrypt the nonce value and other values. This method requires that the public key of the two peers be present on the other peer before the third and fourth messages of an IKE exchange can be accomplished. For this reason, public keys must be manually copied to each peer as part of the configuration process, and therefore this method is limited by the available memory of the receiver. This method is the least used of the four authentication methods.
- **Elliptic Curve Digital Signature Algorithm (ECDSA):** The ECDSA is the elliptic curve analog of the DSA, which is part of the Digital Signature Standard (DSS) method. ECDSA signatures are smaller than RSA signatures of similar cryptographic strength. ECDSA public keys (and certificates) are smaller than similar-strength DSA keys, resulting in improved communications efficiency. Furthermore, on many platforms, ECDSA operations can be computed more quickly than similar-strength RSA operations. These advantages of signature size, bandwidth, and computational efficiency may make ECDSA an attractive choice for many IKE and IKE version 2 (IKEv2) implementations.

IPsec is a framework of open standards. IPsec spells out the messaging to secure the communications but relies on existing algorithms to implement the encryption, authentication, and key exchange. [Figure 13-1](#) illustrates some of the standard algorithms that IPsec uses.

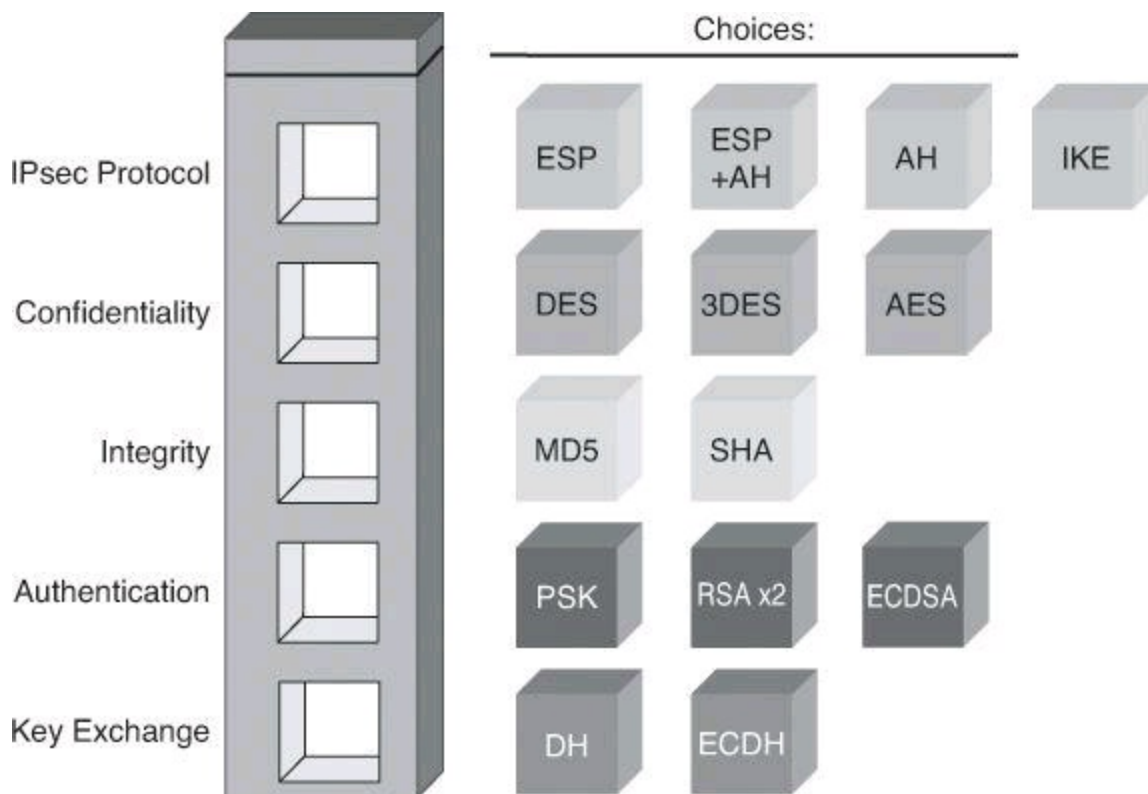


Figure 13-1. IPsec Framework Components

When you configure an IPsec gateway to provide security services, you will call from the components that make the IPsec framework. This framework includes IPsec protocols and many other categories of components, as shown in [Figure 13-1](#). You must first choose an IPsec protocol. The choices are Encapsulating Security Payload (ESP) or ESP with Authentication Header (AH). AH can also be used alone. The second element is an encryption algorithm. Choose the encryption algorithm that is appropriate for the desired level of security: DES, 3DES, or AES. The third element is authentication. Choose an authentication algorithm to provide data integrity: MD5 or SHA. The last element is the DH algorithm group. Choose which group to use, DH Group 1 (DH1), DH Group 2 (DH2), or DH Group 5 (DH5). IPsec provides the framework, and the administrator chooses the algorithms that are used to implement the security services within that framework.

IPsec Protocol

IPsec is a framework of open standards. There are two distinct stages to building IPsec tunnels: IKE and the IPsec protocol itself. IKE is responsible for establishing the rules of engagement and confirming the identity of the peers with whom communication is about to start. IPsec tunnels are responsible for passing data in a secure communication. To make an analogy, consider the FTP protocol, which uses two ports. Port 21 is used as the control port and port 20 is used as the data port. Port 21 is responsible for validating the identity of the parties and negotiating the requests. The actual data (file) exchange is done through the data port, port 20. The process is similar with the IPsec framework. IKE has a similar role as port 21, with responsibility for negotiating the type of connection, the identity of the entities, and so forth. IPsec tunnels are similar to port 20, because they are the secured carrier of the data exchanged between the two peers (the procedure for which will be covered in detail later in this chapter). Think of IKE as the control plane and IPsec as the data plane. For now, let's focus on the actual tunnels that provide secure data exchanges between peers, the actual IPsec tunnels.

IPsec spells out the messaging to secure the communications but relies on existing algorithms, listed in [Figure 13-1](#). There are two main IPsec framework protocols, as depicted in [Figure 13-2](#):

- **Authentication Header (AH):** AH, which is IP protocol 51, is the appropriate protocol to use when confidentiality is not required or permitted. It provides data authentication and integrity for IP packets passed between two systems. It is a means of verifying that any message that is passed from Router A to Router B has not been modified during transit. It verifies that the origin of the data was either Router A or Router B. AH does not provide data confidentiality (encryption) of packets. All text is transported in the clear. If the AH protocol is used alone, it provides weak protection. Consequently, Encapsulating Security Payload uses the AH protocol to provide data encryption and tamper-aware security features.
- **Encapsulating Security Payload (ESP):** ESP is a security protocol that can provide confidentiality and authentication. ESP, which is IP protocol 50, provides confidentiality by performing encryption on the IP packet. IP packet encryption conceals the data payload and the identities of the ultimate source and destination. ESP provides authentication for the inner IP packet and ESP header. Authentication provides data-origin authentication and data integrity. Although both encryption and authentication are optional in ESP, at a minimum, one of them must be selected.

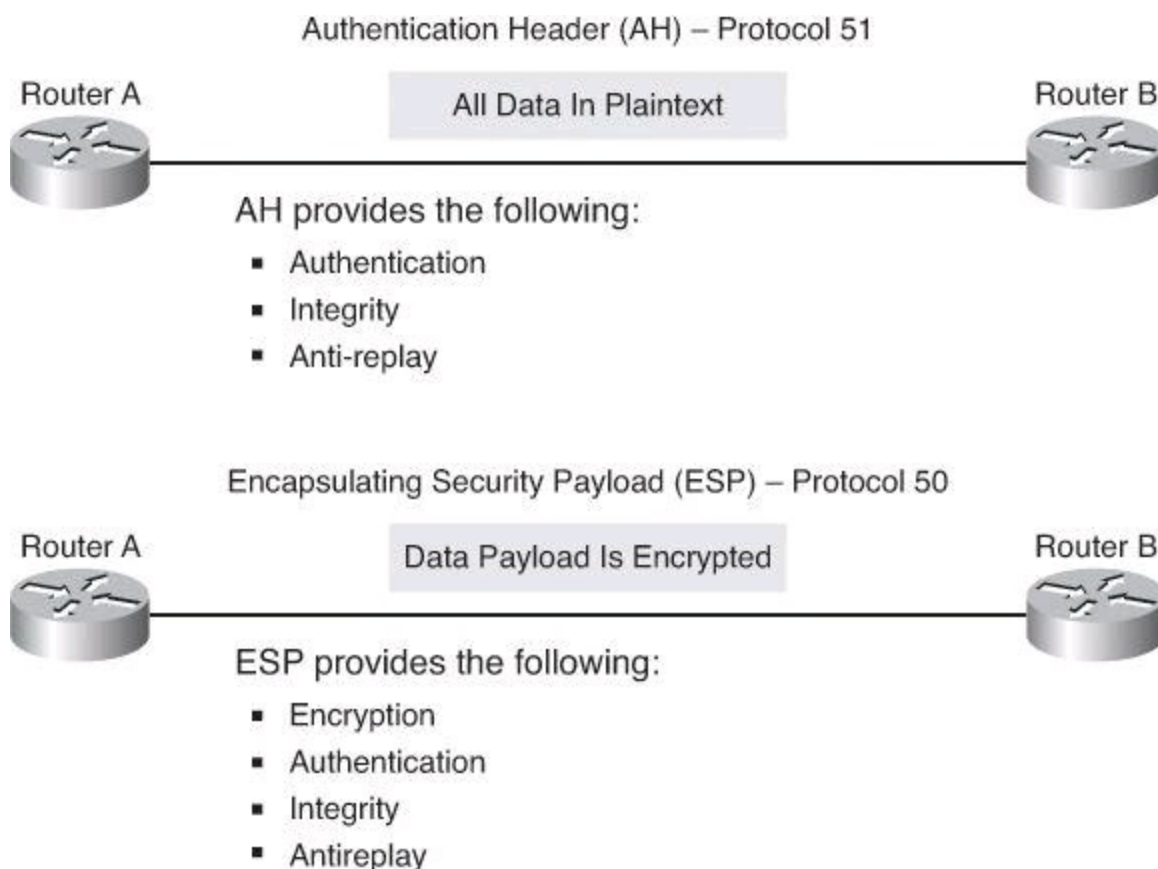


Figure 13-2. IPsec Security Protocols

Authentication Header

AH achieves authenticity by applying a keyed one-way hash function to the packet to create a hash, or message digest. The hash is combined with the text and is transmitted. The receiver detects changes in any part of the packet that occur during transit by performing the same one-way hash

function on the received packet and comparing the result to the value of the message digest that the sender has supplied. The fact that the one-way hash also involves the use of a shared-secret key between the two systems means that authenticity is assured.

The AH function is applied to the entire datagram except for any mutable IP header fields that change in transit; for example, Time to Live (TTL) fields that are modified by the routers along the transmission path. [Figure 13-3](#) illustrates the AH process, which is as follows.

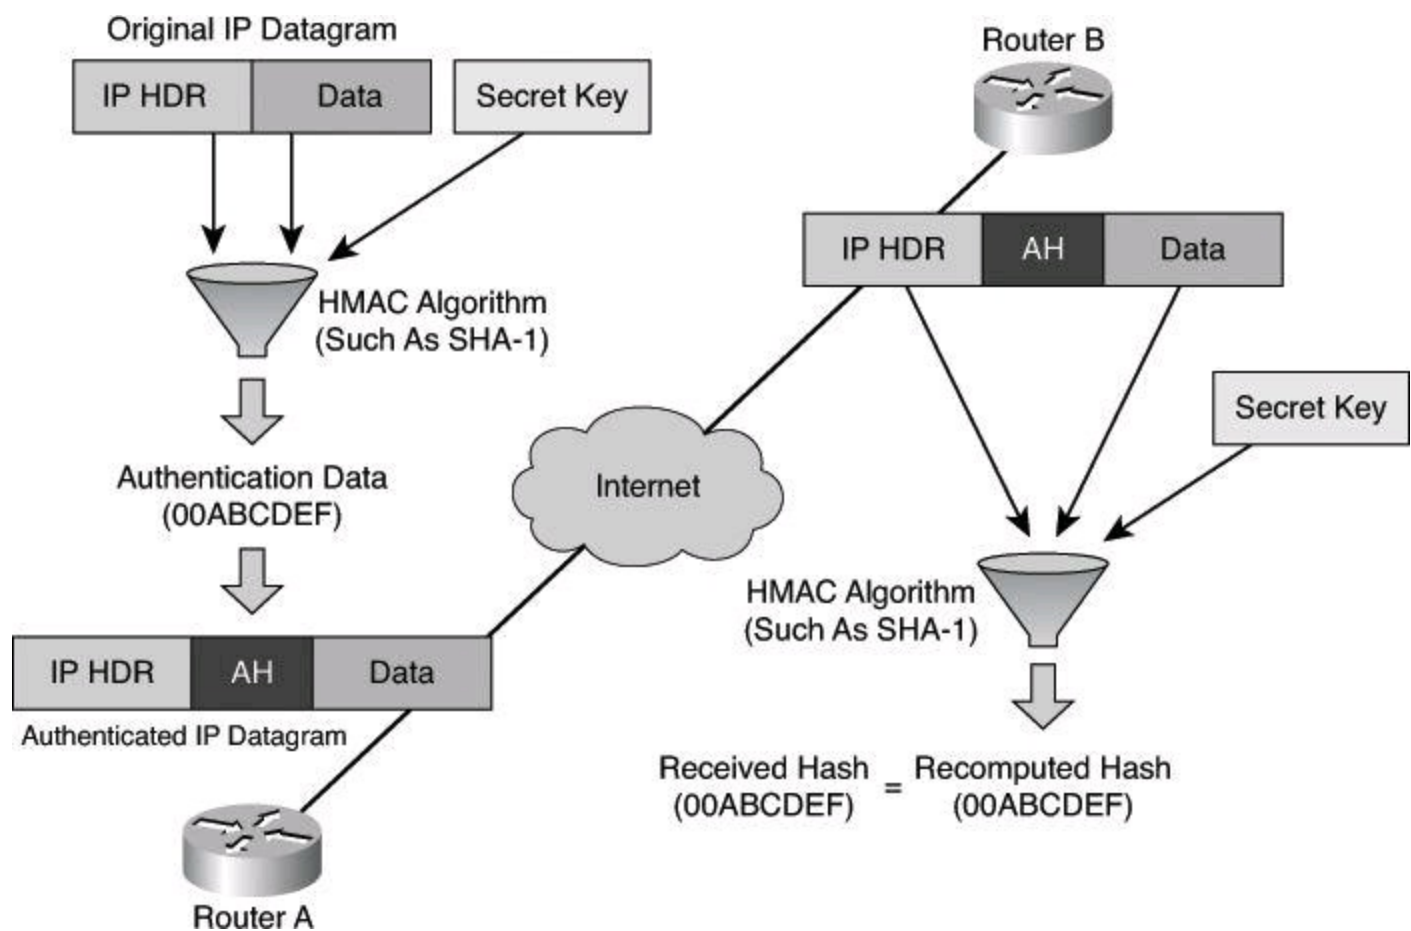


Figure 13-3. AH Authentication and Integrity

Step 1. The IP header and data payload are hashed.

Step 2. The hash builds a new AH header, which is prepended to the original packet.

Step 3. The new packet is transmitted to the IPsec peer router.

Step 4. The peer router hashes the IP header and data payload, extracts the transmitted hash from the AH header, and compares the two hashes. The hashes must match exactly. Even if 1 bit is changed in the transmitted packet, the hash output on the received packet changes and the AH header will not match.

AH supports the HMAC-MD5 and HMAC-SHA-1 algorithms.

Note

AH protects the entire IP packet, which includes the IP header, against tampering during transmission. Because NAT modifies the IP header, it is incompatible with AH. Fixes, such as tunnel mode, discussed later, remedy this problem.

Also, note that the Cisco firewalls, starting at code 7, do not offer the AH option. Only ESP

can be performed.

Essential services provided by AH are

Key
Topic

- Data integrity through hashing
 - Data origin authentication through hashing
 - Antireplay protection
-

AH in Production

AH has its place for secure communication. When most of us hear the term “secure communication,” we think encrypted payload. However, there are situations where confidentiality is not necessary but authenticity and integrity are. Take the example of a car manufacturer that places an order to its brake pads supplier every morning at 8 a.m. for just-in-time delivery the next day. There is nothing secret about today’s order of 8000 brake pads. So, payload encryption is not necessary. However, the brake pads supplier, upon receiving the order, wants to confirm that the order really came from the car manufacturer and not from Johnny the Joker, and that the order of 8000 brake pads wasn’t tampered with. The order for 8000 brake pads would be released with the payload in cleartext; however, an encrypted hash of the packet is included in the transmission, thus providing a proof of origin and integrity.

Encapsulating Security Payload

ESP provides confidentiality by encrypting the payload. It supports a variety of symmetric encryption algorithms. The lowest common algorithm for IPsec is 56-bit DES. Cisco products also support the use of 3DES and especially AES for stronger encryption.

ESP can also provide integrity and authentication of the datagrams. First, the payload is encrypted. Next, the encrypted payload is sent through a hash algorithm, HMAC-MD5 or HMAC-SHA-1. The hash provides authentication and data integrity for the data payload.

Optionally, ESP can also enforce antireplay protection by requiring that a receiving host set the replay bit in the header to indicate that the packet has been seen.

The original data is well protected by ESP because the entire original IP datagram is encrypted, as shown in [Figure 13-4](#). An ESP header and trailer are added to the encrypted payload. With ESP authentication, the encrypted IP datagram and the ESP header and trailer are included in the hashing process. Lastly, a new IP header is prepended to the authenticated payload. The new IP address is used to route the packet through the Internet.




Figure 13-4. ESP Protocol

When both authentication (AH) and encryption (ESP) are selected, encryption occurs first. One reason for this order of processing is that it facilitates rapid detection and rejection of replayed or bogus packets by the receiving device. Before decrypting the packet, the receiver can authenticate inbound packets. By doing this, it can quickly detect problems and potentially reduce the impact of denial-of-service (DoS) attacks.

Essential services provided by ESP are



- Data confidentiality through encryption
- Data integrity through hashing
- Data origin authentication through hashing
- Antireplay protection

IPsec Modes of Operations

ESP and AH can be used in two different ways, or *modes*. The encapsulation can be done in tunnel mode or in transport mode. The encapsulation performed on an ESP packet with each mode is illustrated in [Figure 13-5](#).

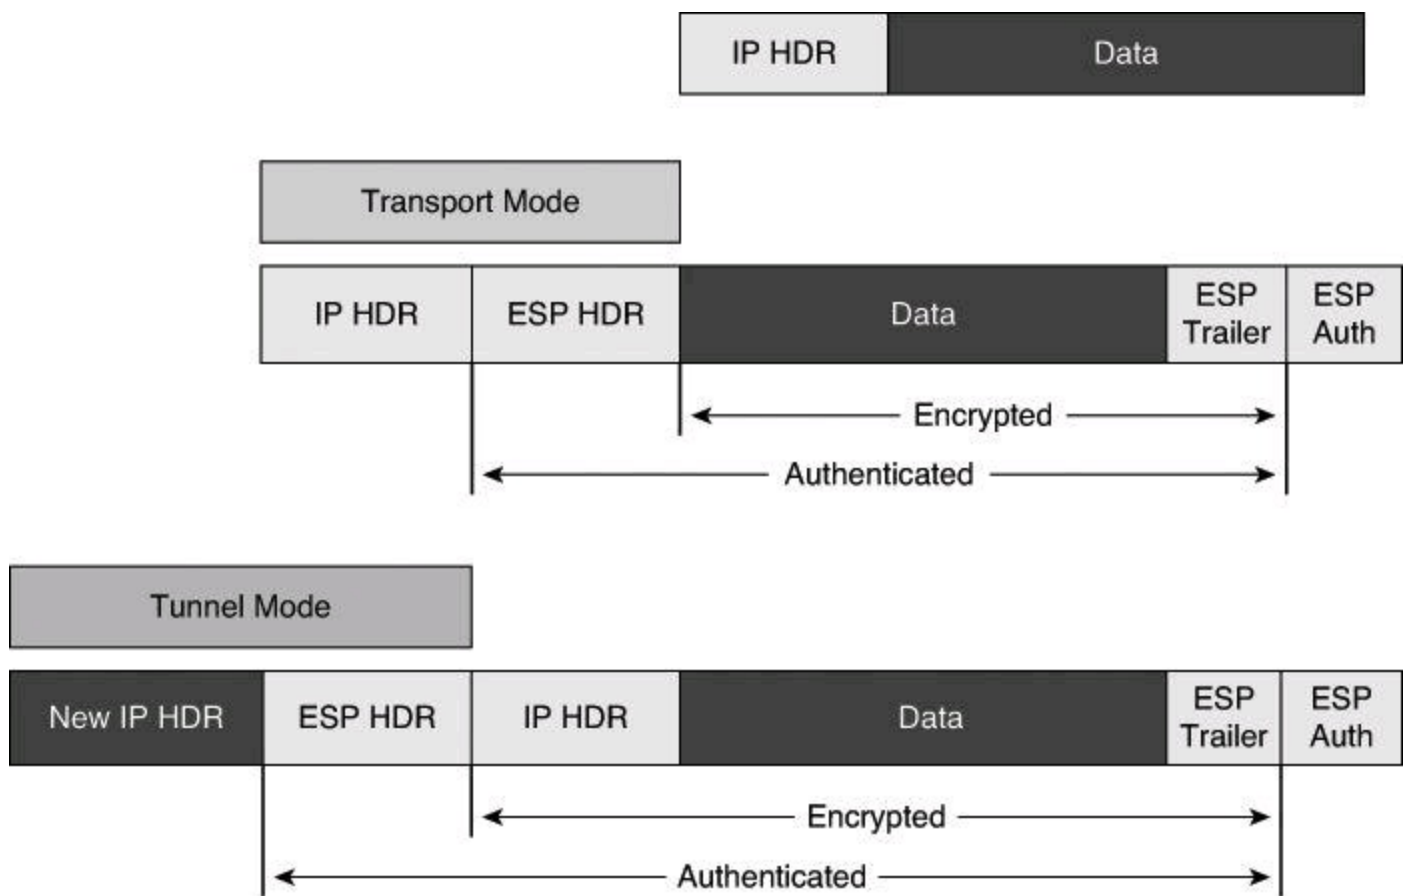


Figure 13-5. Encapsulation with Tunnel Mode and Transport Mode

Transport Mode

In transport mode, security is provided only for the transport layer and above. Transport mode protects the payload of the packet but leaves the original IP address in the clear. The original IP address is used to route the packet through the Internet. ESP transport mode is used between hosts and is therefore not compatible with Network Address Translation (NAT).

Note

Transport mode works well with generic routing encapsulation (GRE) because GRE hides the addresses of the end stations by adding its own IP header.

Tunnel Mode

ESP tunnel mode provides security for the complete original IP packet. The original IP packet is encrypted, and then it is encapsulated in another IP packet. The outside IP address is used to route the packet through the Internet.

ESP tunnel mode is used between a host and a security gateway or between two security gateways, as shown in [Figure 13-6](#). For gateway-to-gateway applications, rather than load IPsec on all the computers at the remote and corporate offices, it is easier to have the security gateways perform the IP-in-IP encryption and encapsulation.

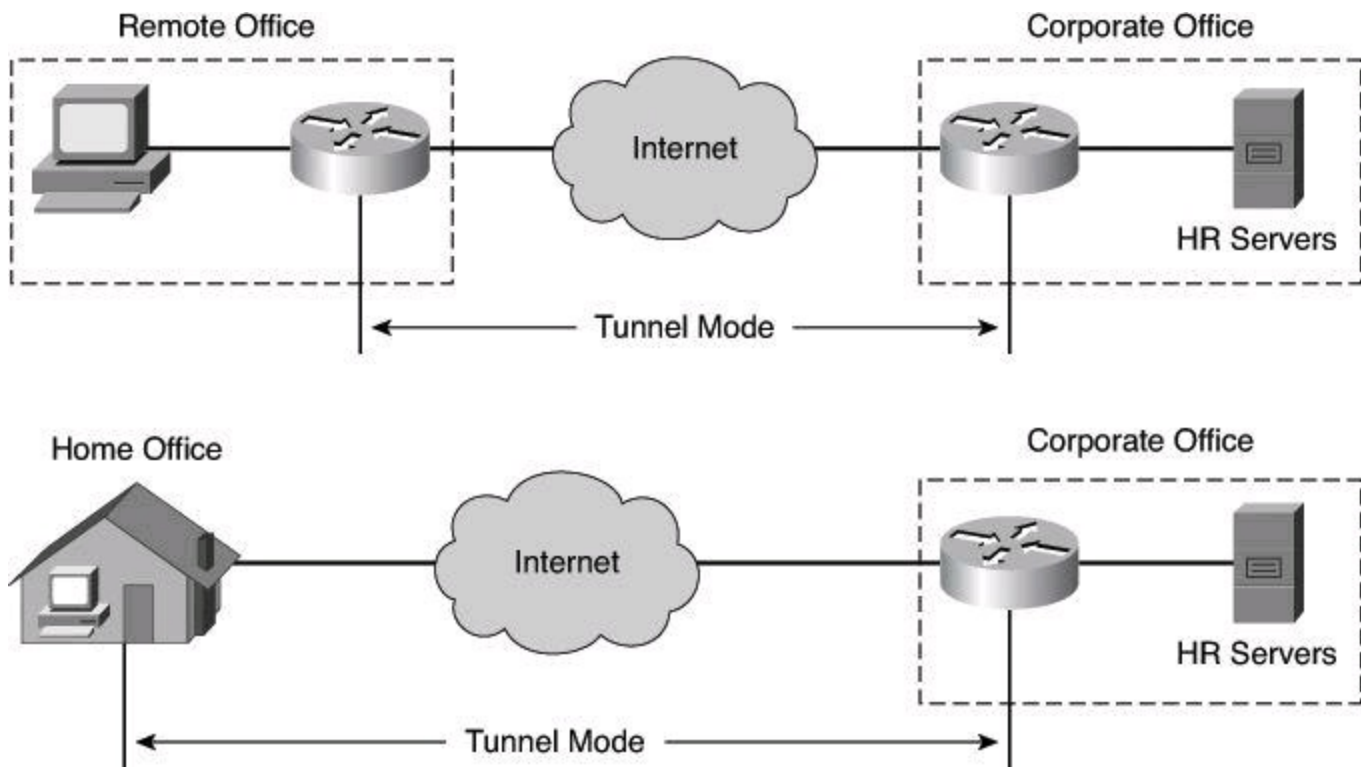


Figure 13-6. IPsec Tunnel Mode

ESP tunnel mode is used in the IPsec remote-access application. At a home office, there might be no router to perform the IPsec encapsulation and encryption. In this case, an IPsec client running on the PC performs the IPsec IP-in-IP encapsulation and encryption. At the corporate office, the router de-encapsulates and decrypts the packet.

IKE Protocol

IPsec implements a VPN solution using an encryption process that involves the periodic changing of encryption keys. IPsec uses the IKE protocol to authenticate a peer computer and to generate encryption keys. IKE negotiates a security association (SA), which is an agreement between two peers engaging in an IPsec exchange and consists of all the required parameters necessary to establish successful communication. An IPsec peer accepting incoming IKE requests listens on UDP port 500.

IPsec uses the IKE protocol to provide these functions:

- Negotiation of SA characteristics
- Automatic key generation
- Automatic key refresh
- Manageable manual configuration

There are two versions of the IKE protocol: IKEv1 and IKEv2. IKEv2 was created to overcome some of the limitations of IKEv1. IKEv2 provides the following enhancements:

- Simplicity, by requiring fewer transactions to establish security associations. A simplified initial exchange of messages reduces latency and increases connection establishment speed.
- Stronger security, through DoS protection and other functions.
- Reliability, by using sequence numbers, acknowledgements, and error correction.

- Flexibility, through support for Extensible Authentication Protocol (EAP) as a method for authenticating VPN endpoints.
- Mobility, by using the IKEv2 Mobility and Multihoming Protocol (MOBIKE) extension. This enhancement allows mobile users to roam and change IP addresses without disconnecting their IPsec session.

Because the IKEv2 base specification includes all the functionality of IKEv1, we'll look at the IKEv1 modes and phases first, followed by a brief overview of the enhancements in IKEv2, and then summarize the differences between these two versions.

Note

Because a picture is worth a thousand words, [Figure 13-7](#) is the visual representation of IPsec sequence of events when tunnels are being built with IKEv1. The process starts at the bottom of the figure with IKE Phase 1, followed by IKE Phase 2, and finally with the two unidirectional IPsec tunnels coming up to exchange the encrypted data.

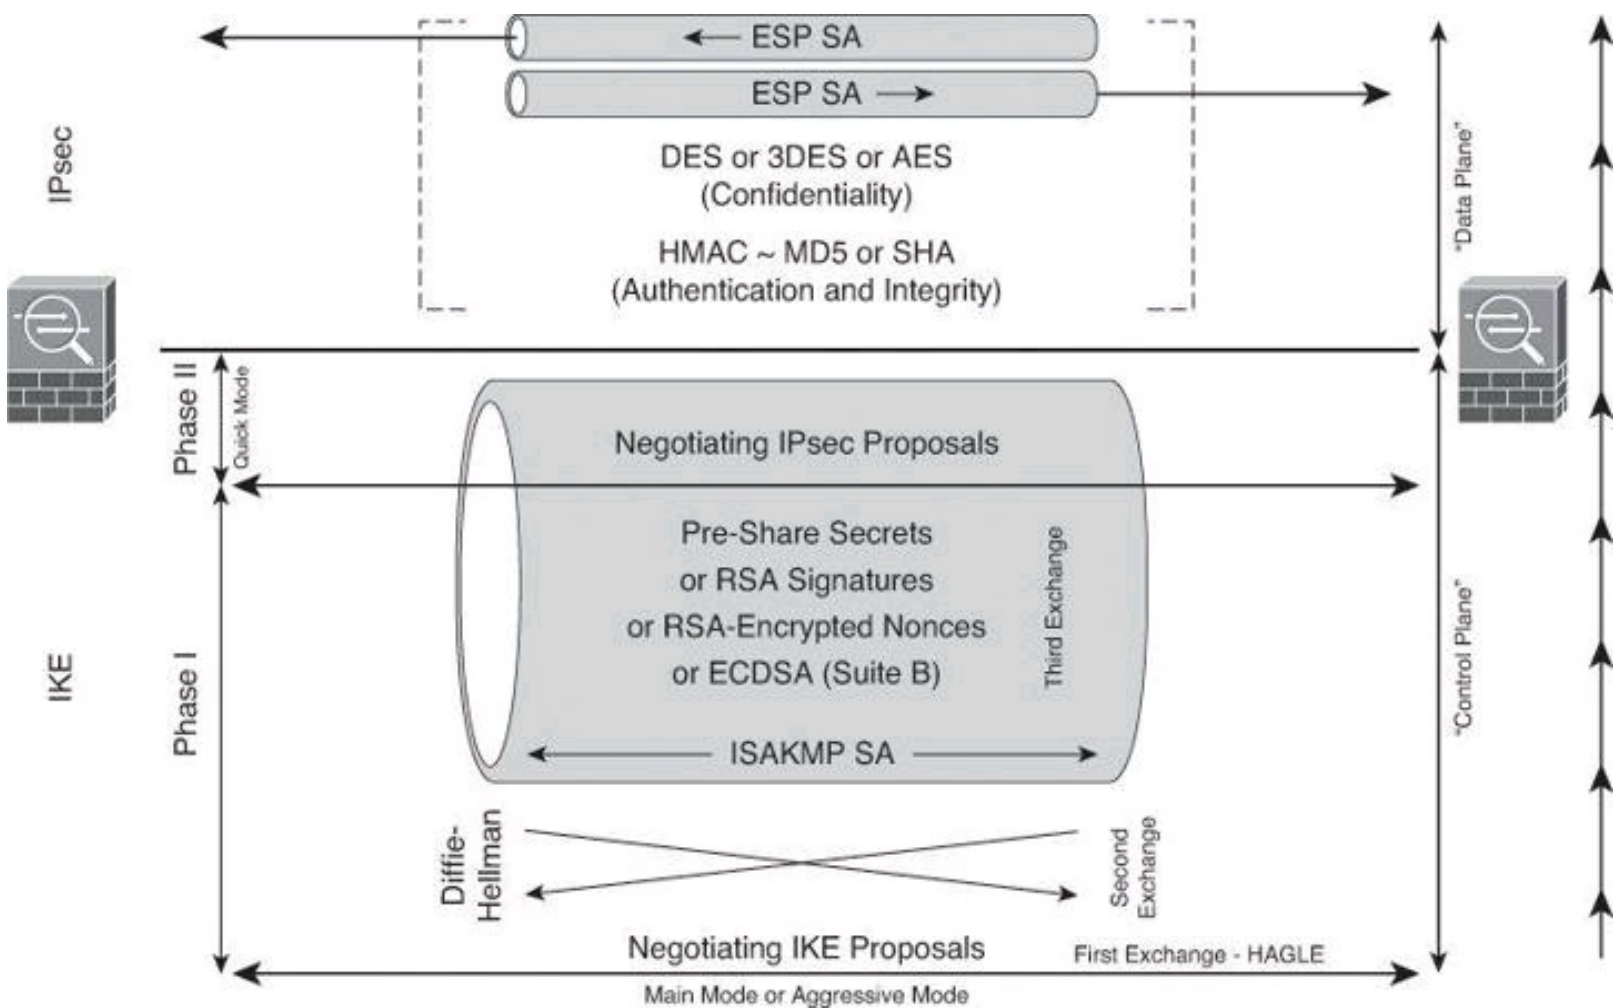


Figure 13-7. Visual Representation of IKEv1 and IPsec Tunnels Being Built from the Ground Up

IKEv1 Modes

To establish a secure communication channel between two peers, the IKE protocol uses the following three modes of operation:

- **Main mode:** In main mode, which takes place in IKEv1 Phase 1, an IKE session begins

with one computer (the initiator) sending a proposal or proposals to another computer (the responder). The proposal sent by the initiator defines what encryption and authentication protocols are acceptable, how long keys should remain active, and whether Perfect Forward Secrecy (PFS) should be enforced. (PFS is covered later in this chapter.) There are three exchanges typical of main mode, for a total of six packets:

- The first exchange between the initiator and the responder establishes the basic security policy. The initiator sends all its Internet Security Association and Key Management Protocol (ISAKMP) policies. The responder chooses a proposal best suited to the security situation and then sends that proposal to the initiator. This exchange totals two packets.
 - The next exchange passes DH public keys between the two users. DH key exchange is a cryptographic protocol that allows two parties that have no prior knowledge of each other to establish a shared-secret key over an unsecure communications channel. All further negotiation is encrypted within the IKE SA.
 - The third exchange authenticates an ISAKMP session. The fifth and sixth packets are the first packets exchanged encrypted, as depicted earlier in [Figure 13-7](#), which shows authentication being conducted inside the ISAKMP tunnel. Once the IKE SA is established, IPsec quick mode negotiation begins.
- **Aggressive mode:** Aggressive mode compresses into three packets the IKE SA negotiation phases just described. In aggressive mode, which takes place in IKEv1 Phase 1, the initiator passes all data that is required for the SA. The responder sends the proposal, key material, and ID and authenticates the session in the next packet. The initiator replies by authenticating the session. Negotiation is quicker, and the initiator and responder IDs pass in plaintext.
 - **Quick mode:** Quick mode IPsec negotiation, which happens in IKEv1 Phase 2, takes place after successful IKE SA negotiation. Quick mode is similar to aggressive mode IKE negotiation, except that negotiation is protected within an IKE SA. Quick mode negotiates the SA for the data encryption and manages the key exchange for that IPsec SA, as shown in [Figure 13-7](#). This means that it negotiates a shared IPsec transform, derives shared-secret keying material that the IPsec security algorithms will use, and establishes IPsec SAs. During quick mode, the peers exchange the list of protected networks: the subnets at each end that must use the encrypted tunnel when communicating amongst themselves.

IKEv1 Phases

To establish a secure communication channel between two peers, the IKE protocol executes the following phases:

- **IKE Phase 1:** Two IPsec peers perform the initial negotiation of SAs. In this phase, the SA negotiations are bidirectional; data may be sent and received using the same encryption key. In IKE Phase 1, the transform sets, hash methods, and other parameters are determined. Optionally, IKE Phase 1 can include authentication, in which each peer in the SA negotiation is able to verify the identity of the other. Even if the SA negotiation data stream between the two IPsec peers is compromised, there is little chance that the encryption keys could be guessed and thus the traffic decrypted.

- **IKE Phase 2:** SAs are negotiated by the IKE process ISAKMP on behalf of other services, such as IPsec, that need encryption key material for operation. Quick mode negotiates the IKE Phase 2 SAs. In this phase, the SAs that IPsec uses are unidirectional; therefore, a separate key exchange is required for each data flow.

Note

Other, optional parameters are sometimes negotiated between IKE Phase 1 and IKE Phase 2, such as level of DH, encryption algorithms, and methods of authentication.

IKEv1 Phase 1

The basic purpose of IKE Phase 1, shown in [Figure 13-8](#), is to negotiate IKE policy sets, authenticate the peers, and set up a secure channel between the peers. IKE Phase 1 occurs in two modes: main mode and aggressive mode.

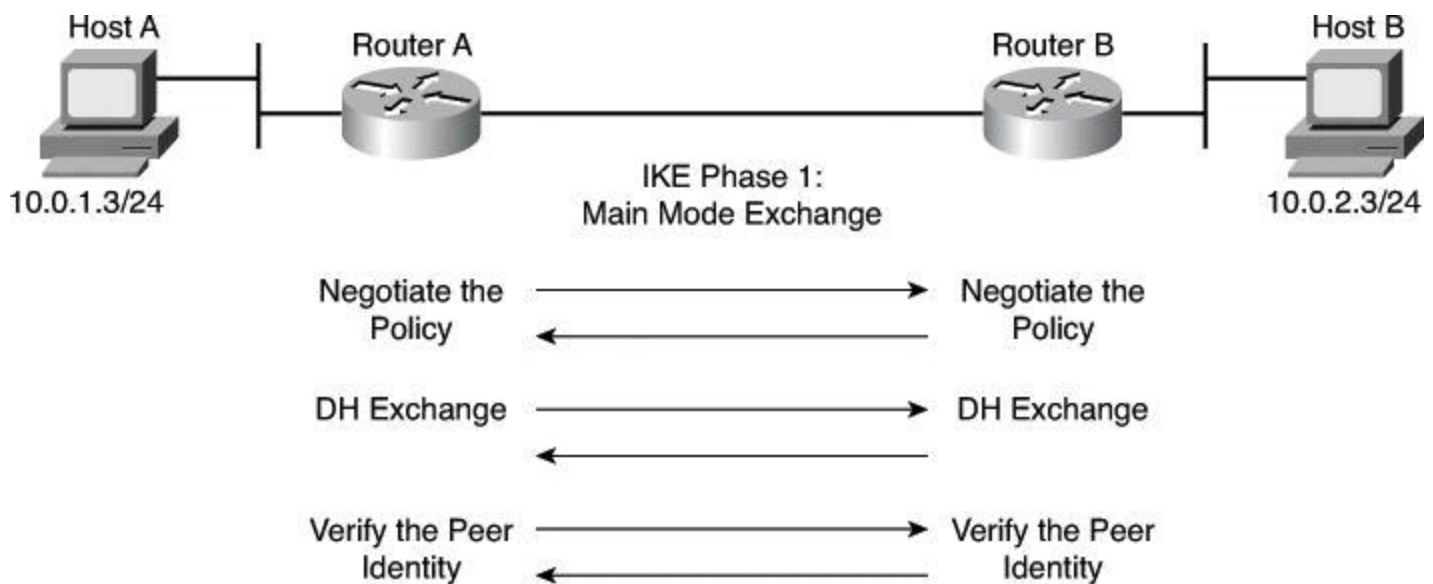


Figure 13-8. IKE Phase 1

Main Mode

Main mode has three two-way exchanges between the initiator and receiver:

- **First exchange:** Peers negotiate and agree on the algorithms and hashes that will be used to secure the IKE communications.
- **Second exchange:** DH generates public and private values. The peers exchange their public values, and the result is a shared secret. The shared-secret key is used to generate all the other encryption and authentication keys.
- **Third exchange:** The identity of the other side is verified. The main outcome of main mode is a secure communications path for subsequent exchanges between the peers.

Note

In aggressive mode, fewer exchanges are done, and the exchanges have fewer packets. As an example, with aggressive mode, the peers will exchange their identity and a hash of the PSK even though the DH calculations have not been performed yet, and therefore there is

no shared secret to send data in an encrypted format.

Aggressive Mode

Aggressive mode, as explained earlier, compresses the IKE SA negotiation phases into a total of three messages. Negotiation is quicker, and the initiator and responder IDs pass in plaintext.

Main Mode Versus Aggressive Mode

The use of main mode or aggressive mode is a tradeoff between performance and security. Main mode requires more messages but does not expose the identity of the peers. This identity is protected by the policies and keys that are negotiated in the first two exchanges. Aggressive mode requires fewer messages and therefore is more efficient. However, peer identities are exposed before negotiating a secure channel, so it assumes a trusted transport and a more protected environment outside of IPsec. The choice of mode is a configurable option in most IPsec implementations.

IKEv1 Phase 1 Example

When you are trying to make a secure connection between Host A and Host B through the Internet, a secure path (a tunnel) is established between Router A and Router B. Through the tunnel, the IPsec encryption, authentication, and other protocols are negotiated.

IKEv1 First Exchange: IKE Policy Is Negotiated

The ultimate goal of IPsec is to have a secure tunnel through which we can transmit encrypted data. Before we start exchanging encrypted data with a peer, though, we need to confirm who we are dealing with and how we will be sending encrypted data to it. You would never send your credentials in cleartext to your peer for authentication, so IPsec needs to start with transmission of other information in cleartext. This is the beginning of IKEv1 Phase 1.

In the cleartext, the two peers agree on the following:

- **Encryption algorithm they will use for their IKE tunnel:** DES, 3DES, or EAS
- **Hashing algorithm they will use for integrity check:** MD5 or SHA
- **Authentication method they will use:** PSK, RSA Signatures, RSA encrypted nonces, or ECDSA
- **DH group:** The DH group used to generate the symmetrical key used to encrypt IKE traffic between the two peers, such as authentication credentials, which shouldn't be sent in cleartext. An IKE tunnel is made of different protocols, one of which is called Internet Security Associate Key Management Protocol (ISAKMP). ISAKMP is an encrypted tunnel through which peers will be exchanging, among other things, their credentials. To have an encrypted tunnel, both peers must have the same symmetrical key. The DH group designates the length of the symmetrical key that will be generated during the upcoming Diffie-Hellman operation.
- **Duration of ISAKMP tunnel:** This is referred to as the *lifetime*.

IKE HAGLE

Use the acronym HAGLE to assist you in remembering the five elements that peers agree

upon during IKEv1 Phase 1:

H = Hash

A = Authentication

G = Group (DH group)

L = Lifetime

E = Encryption

Rather than negotiate each protocol individually, the protocols are grouped into sets, called *IKE policy sets*. IKE policy sets are exchanged during the IKE main mode, first exchange phase. If a policy match is found between peers, main mode continues. If no match is found, the tunnel is torn down.

Of the five values that both peers must agree on, only lifetime can be adjusted dynamically. It will be adjusted to the smallest values exchanged between the two peers. All the other elements of IKEv1 Phase 1 negotiations (hash, authentication, DH group, encryption) must have a perfect match between the two peers.

In [Figure 13-9](#), Router A sends IKE policy sets 10, 20, and 30 to Router B. Router B compares its sets with those received from Router A. In this instance, there is a policy match. Policy set 10 of Router A matches policy set 20 of Router B and is selected over policy set 30 on both sides, because the lowest policy priority number wins.

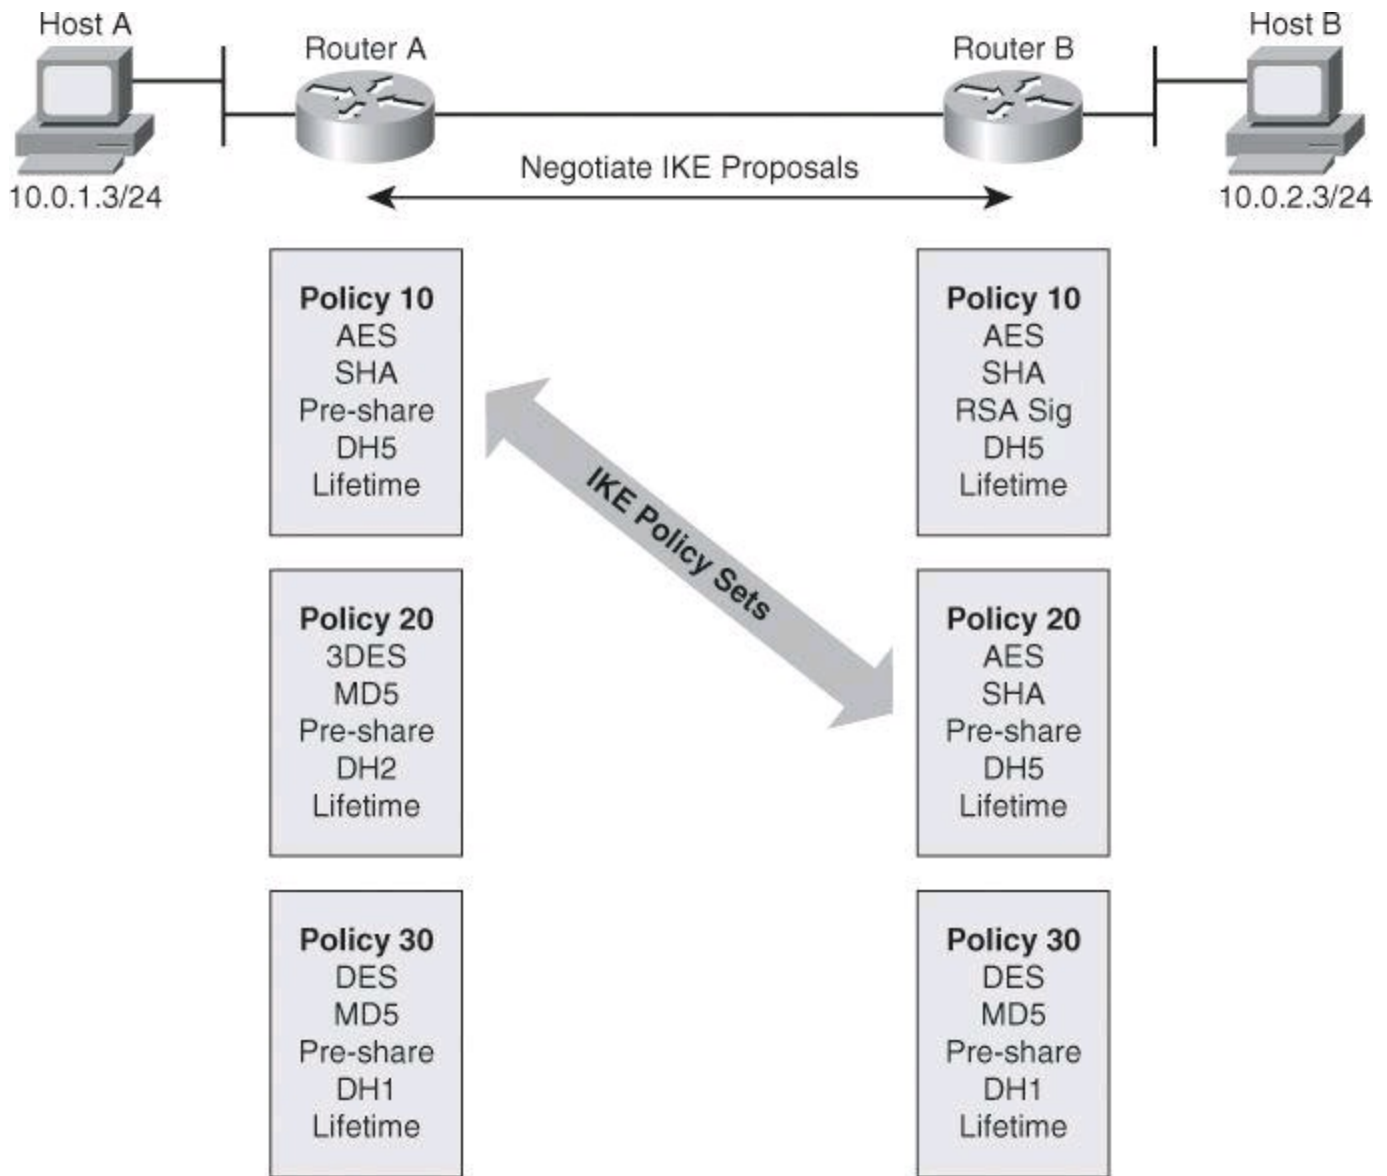


Figure 13-9. IKEv1 Phase 1, First Exchange: Policy Is Negotiated

Note

Policy set numbers are only locally significant to a VPN device. The policy set numbers do not have to match between two VPN peers.

In a point-to-point application, each end may need only a single IKE policy set to be defined. However, in a hub-and-spoke environment, the central site may require multiple IKE policy sets to satisfy all the remote peers.

IKEv1 Second Exchange: DH Key Exchange

When the IKE policy is agreed on, including the size of the prime number to be used for DH, the two peers run the DH key exchange protocol. You will remember from [Chapter 12](#) that, as shown in [Figure 13-10](#), the result of DH will be the creation of a shared secret, where $K = K'$, that is needed by the various symmetric encryption and hashing algorithms upon which IKE and IPsec will ultimately agree.

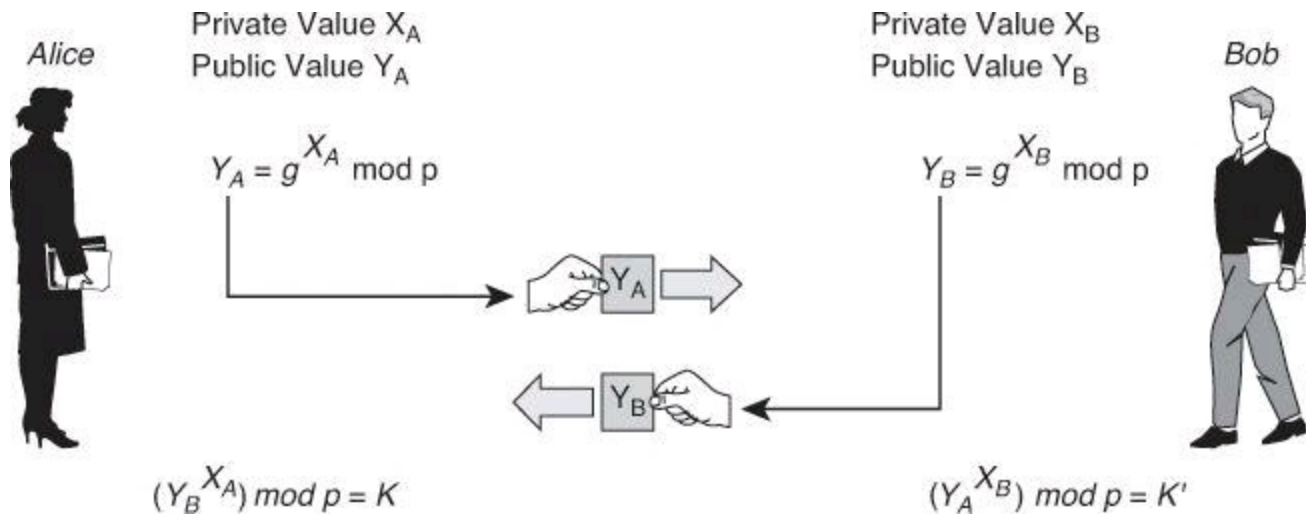


Figure 13-10. IKEv1 Phase 1, Second Exchange: DH Key Exchange

Several levels of DH key exchanges are available in Cisco IOS Software:

- **DH Group 1:** A key exchange that uses a 768-bit prime number. This group is the usual choice when the encryption algorithm is DES.
- **DH Group 2:** A key exchange that uses a 1024-bit prime number. This group is the usual choice when using 3DES for encryption.
- **DH Group 5:** A key exchange that uses a 1536-bit prime number. DH 5 should be used with AES.
- **DH Group 7 (ECC):** A key exchange that generates IPsec keys when the elliptic curve field is 163 bits. This group was designed to be used with low-powered hosts such as PDAs.
- **DH Group 14:** A key exchange that results in 2048 bits of keying material.
- **DH Group 15:** A key exchange that results in 3072 bits of keying material.
- **DH Group 16:** A key exchange that results in 4096 bits of keying material.
- **DH Group 19:** A key exchange that results in 256 bits of keying material.
- **DH Group 20:** A key exchange that results in 384 bits of keying material.
- **DH Group 24:** A key exchange that results in 2048 bits of keying material.

Note

Groups 5, 14, 15, and 16 should be used with AES. Groups 19, 20, and 24 should be used with ECDH.

The group that is chosen must be strong enough (have enough bits) to protect the IPsec keys during negotiation. A generally accepted guideline recommends the use of a 2048-bit group after the year 2013 (until the year 2030). Either DH14 or DH24 can be selected to meet this guideline. Even if a longer-lived security method is needed, the use of ECC is recommended, but G15 and DH16 can also be considered.

IKEv1 Third Exchange: Authenticate Peer Identity

When you are conducting business over the Internet, it is necessary to know who is at the other end of the tunnel. The device on the other end of the VPN tunnel must be authenticated before the communications path is considered secure. The last exchange of IKE Phase 1 authenticates the remote peer, as illustrated in [Figure 13-11](#).

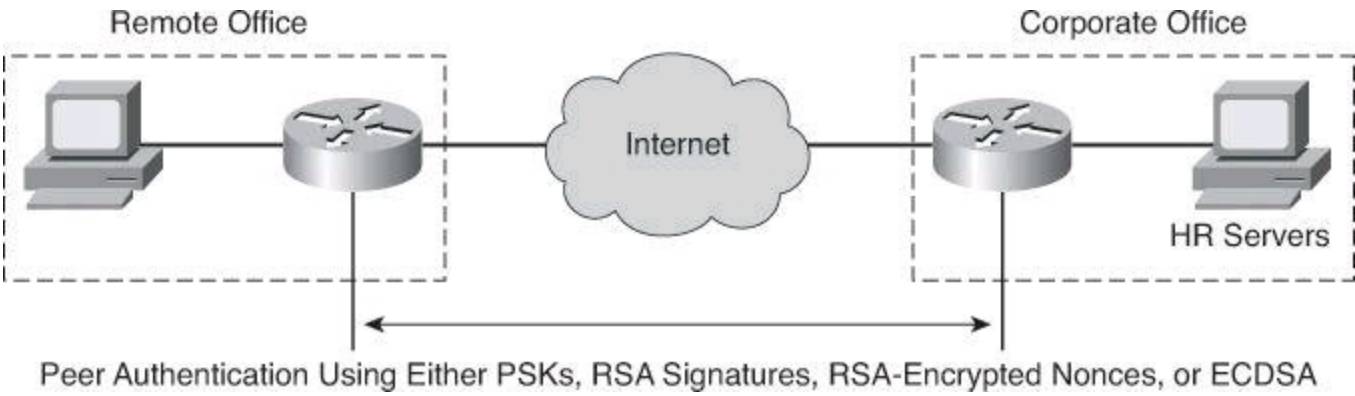


Figure 13-11. IKEv1 Phase 1, Third Exchange: Authenticate Peer Identity

As mentioned earlier, there are four data origin authentication methods with IKEv1:

- **PSKs:** Pre-shared keys are a secret key value that is entered into each peer manually and is used to authenticate the peer.
- **RSA signatures:** RSA signatures are the exchange of digital certificates that is used to authenticate the peers in addition to sending a hash value of a message encrypted with its private key as proof of its identity.
- **RSA encrypted nonces:** Nonces are random numbers that are generated by each peer and then encrypted and exchanged between peers. The two nonces are used during the peer-authentication process.
- **ECDSA signatures:** Exchange of certificates. ECDSA certificates are smaller than RSA signatures of similar cryptographic strength, resulting in improved communications efficiency. ECDSA is available with Suite B.

IKEv1 Phase 2

The purpose of IKEv1 Phase 2 is to negotiate the IPsec security parameters, as shown in [Figure 13-12](#), that will be used to secure the IPsec tunnel. IKEv2 Phase 2 performs the following functions.

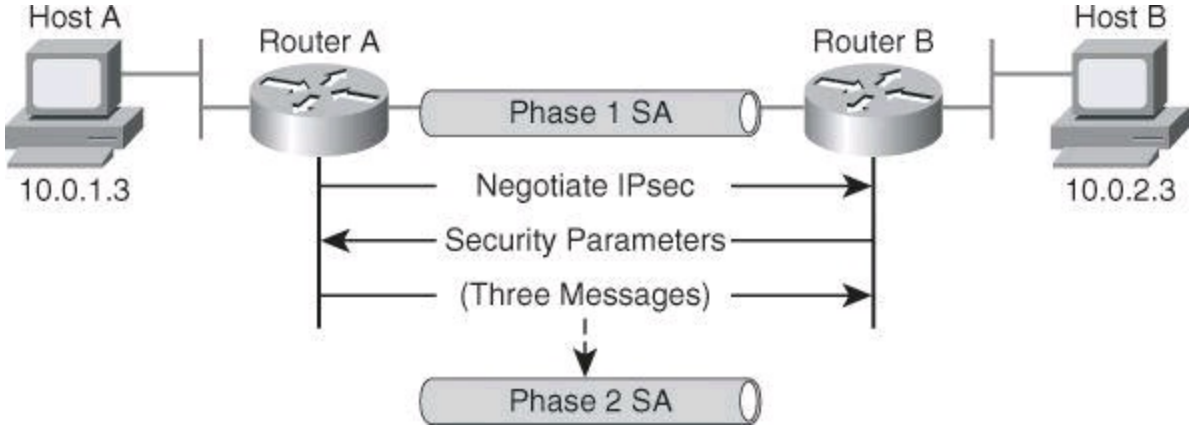


Figure 13-12. IKE Phase 2

- Negotiates IPsec security parameters, known as IPsec transform sets.

- Establishes IPsec SAs.
- Periodically renegotiates IPsec SAs to ensure security.
- Optionally, performs an additional DH exchange to generate IPsec SA keys that have no relation to the IKE keys. Generating IPsec keys from scratch for the purpose of IPsec SAs is referred to as Perfect Forward Secrecy (PFS), which is described after IKEv2 quick mode.

IKEv2 Quick Mode

IKE Phase 2 has one mode, called quick mode. Quick mode occurs after IKE has established the secure tunnel in Phase 1. It negotiates a shared IPsec transform, derives shared-secret keying material that the IPsec security algorithms will use, and establishes IPsec SAs. Quick mode also exchanges nonces that are used to generate new shared-secret key material and to prevent replay attacks from generating false SAs.

Quick mode also renegotiates a new IPsec SA when the IPsec SA lifetime expires. Basically, quick mode refreshes the keying material that creates the shared-secret key, which is based on the keying material derived from the DH exchange in Phase 1.

Perfect Forward Secrecy

The basic principle of PFS is that new keys (shared secrets) cannot be derived from older keys. The idea is that if keys are compromised, previous and subsequent keys are secured because they were generated from scratch and not derived. Therefore, PFS exists when a DH exchange is done at each rekeying interval, which is preferable, but might not be necessary, instead of deriving the new keys from the previous keys.

IKE Version 2

In IKEv2, there is a simplified initial exchange of messages that, compared to IKEv1, reduces latency and increases the connection establishment speed. The IKEv2 base specification includes all the functionality of IKEv1 as well as additional functionality. It preserves most of the features of version 1, including the two negotiation phases. However, the specific operations of the two phases differ:

- **Phase 1:** In IKEv2, there is only one mode or set of network flows defined for Phase 1 negotiation. The set requires four messages, two from the initiator and two from the responder. The first two of these four messages flow in the clear on the network. The other two messages are encrypted. Still, this mode is equivalent to main mode in IKEv1, and not to aggressive mode.

The negotiation of the first Phase 2 SA is accomplished within these four flows, still in Phase 2. When the fourth message flows, both Phase 1 and Phase 2 SAs are activated. Either the initiator or the responder can later activate additional Phase 2 SAs using this Phase 1 SA.

- **Phase 2:** The purpose of Phase 2 negotiation is to establish a set of parameters, known as an SA, which is used to protect specific types of IP traffic. The Phase 2 SA contains the keys that are used to encrypt and decrypt IPsec packets on the host, authenticate IPsec packets on the host, or both. Because the first Phase 2 SA is created during Phase 1, all subsequent Phase 2 SAs are called child SAs. The identity of the remote peer is implicitly

authenticated when the Phase 2 SA is used.

IKEv2, illustrated in [Figure 13-13](#), uses fewer messages to accomplish the same (and more) objectives as IKEv1. In addition, the parent-child relationship of Phase 2 SAs streamlines the process. The first Phase 2 SA is created during phase 1, and negotiated parameters are added here for child Phase 2 SAs.

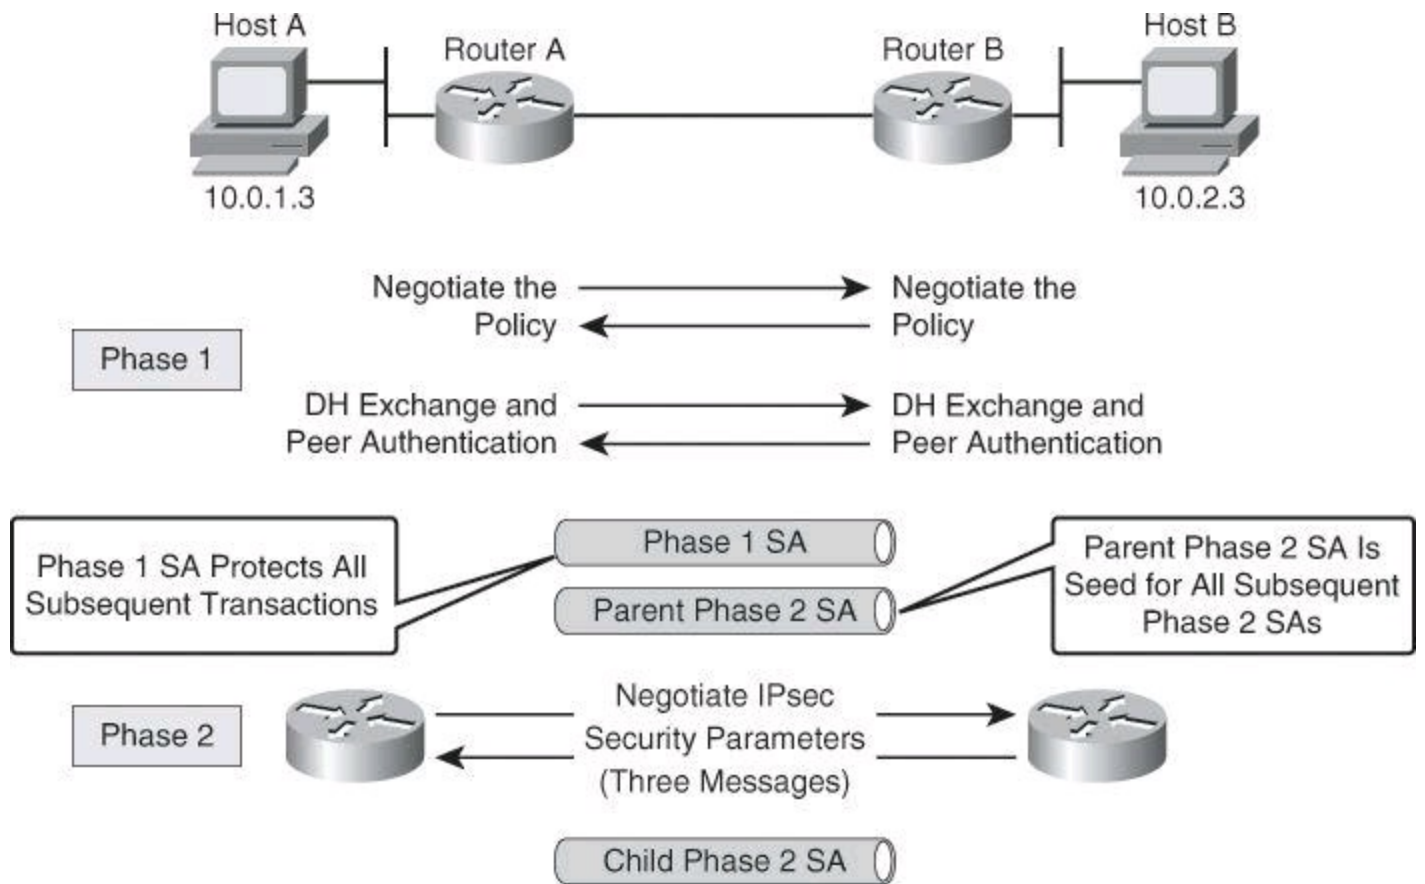


Figure 13-13. IKEv2: A Simplified Approach

For instance, rekeying in IKEv2 negotiates new keys, but it does not reauthenticate the identity of the peer. When an IKEv1 Phase 1 SA is refreshed, new keys are negotiated, and identity of the IKE peer is reauthenticated.

IKEv1 Versus IKEv2

A full lecture on IKEv2 is beyond the scope of this book; however, [Table 13-1](#) provides you with a summary of the differences between IKEv1 and IKEv2.

Table 13-1. Differences Between IKE Version 1 and Version 2

Benefit	Enhancement	IKEv1	IKEv2
Simplicity	RFC documents	RFC 2407, 2408, 2409, 4109, 4995	RFC 5996
	Negotiation		
	Phase 1	6 messages (main mode) or 3 messages (aggressive mode)	4 messages
	Phase 2	3 messages	2 messages
Security	Identity hiding	Optional, only in main mode	Always
	DoS protection	No	Yes, optional
Reliability	Message ACK	None	Reliable; all messages are acknowledged and sequenced
Flexibility	Backward compatibility	No	Yes
	Authentication	None	Extensible Authentication Protocol
	Rekeying	Requires reauthentication	No reauthentication required
	Mobility	None	MOBIKE for Layer 3 roaming

The notable differences with IKE version 2 are

- **Fewer RFCs:** The specifications for IKEv1 were covered in at least three RFCs and even more. IKEv2 combines the specifications into one RFC.
- **Simplified message exchange:** IKEv2 has one four-message initial exchange mechanism while IKEv1 provides eight distinctly different initial exchange mechanisms, each one of which has slight advantages and disadvantages.
- **Security:** Identities are always hidden during Phase 1 in IKEv2. DoS attacks are mitigated because IKEv2 does not process a request until it determines the requester. IKEv1 can be tricked to perform much cryptographic (and expensive) processing from false locations (spoofing).
- **NAT traversal (NAT-T):** ESP as is couldn't be used with PAT because there are no ports that could be translated by the Internet-facing router configured for PATting outbound connections. To mitigate this problem, NAT-T was created. VPN clients, sitting behind a PAT device, encapsulate IKE and ESP in UDP port 4500, which enables these protocols to pass through a device performing NAT. NAT-T is used by both IKEv1 and IKEv2.
- **Reliability:** IKEv2 uses sequence numbers and acknowledgments to provide reliability

and mandates some error processing logistics and shared state management.

- **Support for EAP:** By using EAP, IKEv2 can leverage existing authentication infrastructure and credential databases. EAP allows users to choose a method suitable for existing credentials, and also makes separation of the IKEv2 responder (VPN gateway) from the EAP authentication endpoint (backend authentication, authorization, and accounting [AAA] server) easier.
- **Mobility:** Through MOBIKE, IKEv2 allows for Layer 3 roaming that changes the peer's IP address. The main scenario for MOBIKE is enabling a remote-access VPN user to move from one address to another without reestablishing all SAs with the VPN gateway. For instance, a user could start from fixed Ethernet in the office and then disconnect the laptop and move to the office wireless LAN. When the user leaves the office, the laptop could start using General Packet Radio Service (GPRS). When the user arrives home, the laptop could switch to the home wireless LAN.

IPv6 VPNs

IPsec is the baseline protocol that was established for enabling packet security during transmission of packets in IPv6. Even though the IPv6 standards mandate that IPsec be implemented, they do not mandate that IPsec be used for all IPv6 communications. Such a mandate might not be realistic, given the limited computing resources of a personal digital assistant (PDA), telephone, printer, sensor, toaster, or refrigerator. These small embedded devices can lack the computing resources to perform encryption functions.

IPsec is also native to IPv6. It defines new protocol headers that add authentication and confidentiality to IPv6 packets. The original IPsec RFCs defined the use of an AH to secure the header information and the content of the packet and the use of an ESP to secure the contents of the packet. The IPsec architecture for IPv4 and the IPsec architecture for IPv6 are similar as far as the standards are concerned. In IPv4, AH and ESP were IP protocol headers. The difference is that IPv6 uses the extension header approach, as learned in the CCNA and as shown in [Figure 13-14](#). In IPv6, ESP uses the next-header value of 50 and AH uses the next-header value of 51.




Figure 13-14. IPv6 Extension Header Approach

Because of the ubiquitous nature of the IPv6 Internet (at least in the future), some of the more recent features and functionalities that are found in IPsec become crucial when the protocol is used in IPv6 networks. The latest cryptography suites (Suite B and beyond) and the mobility enhancements that are found in IKEv2 are key.

Here is a summary of IPsec for IPv6 characteristics:

- IPsec is mandatory for IPv6.

- IPsec is native to IPv6.
- Includes built-in confidentiality, integrity, authentication, and antireplay.
- Offers flexibility and low overhead through extension headers.
- The IPsec framework and behavior are the same as IPsec for IPv4.
- Strong encryption (Suite B) and mobility enhancements (IKEv2) are key in IPv6.
- Only site-to-site tunnel mode VPNs are supported in Cisco IOS as of version 15.1.

IPsec Services for Transitioning to IPv6

In the early stages of the transition to IPv6, much of the world's networks and the Internet will continue to use IPv4. Several options are available for transitioning to IPv6. Some of them involve tunneling IPv6 in IPv4 (such as manual IPv6 tunnels, automatic IPv4-compatible tunnels, GRE, automatic 6to4 tunnels, and Intra-Site Automatic Tunnel Addressing Protocol [ISATAP] tunnels).

Using these methods, you can transport IPsec-protected IPv6 traffic over an IPv4 WAN. If you only have IPv4 connectivity between your remote locations and your central site, you can use a tunnel. The tunnel interface can carry the IPv6 packets within an IPv4 tunnel.

[Figure 13-15](#) shows the resulting packet structure. The IPv4 packet is encapsulated by IPsec with a header and a trailer. The result is encapsulated with an IPv4 header that allows forwarding across the IPv4 Internet. The recommended IPsec mode (per RFC 4891) is transport mode.

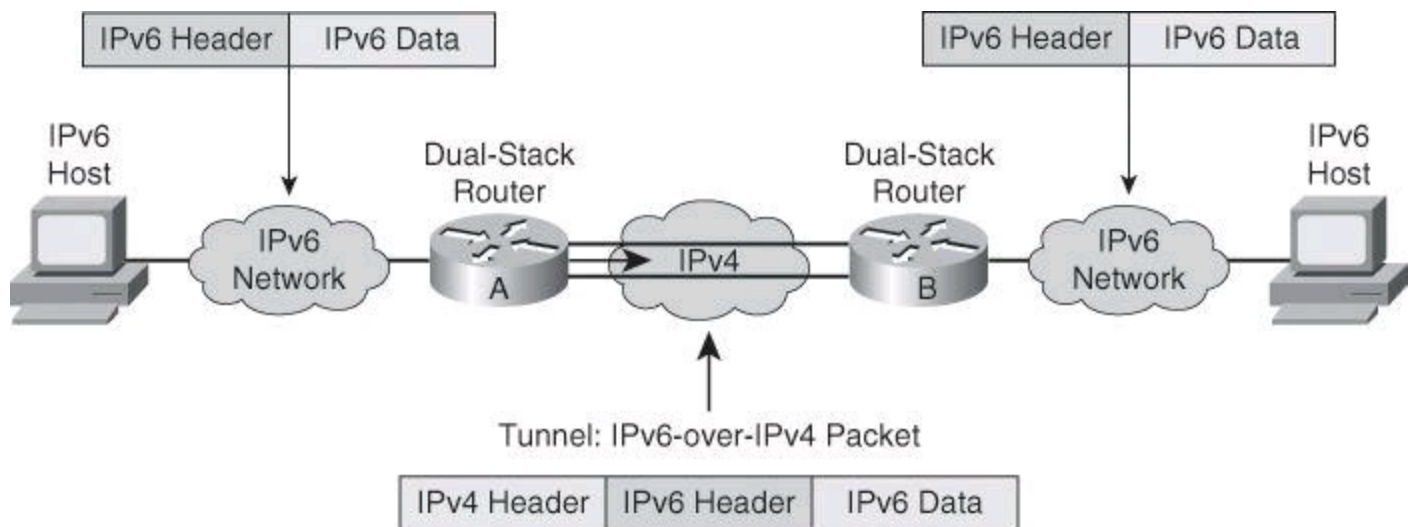


Figure 13-15. IPv6 Transition Using IPv6-over-IPv4 Tunnel

ESP uses the next-header value of 50 and AH uses the next-header value of 51.

Solutions for transitioning IPsec services to IPv6 are

- **Tunneling IPv6 over IPv4:** A common transition strategy
- **IPv4 IPsec:** Can be used to provide VPN services to IPv6
- **IPsec transport mode:** Typically used

Summary

Key
Topic

The key points covered in this chapter are as follows:

- IPsec is a ubiquitous VPN technology that provides confidentiality, data-integrity, authentication, and antireplay services.
- ESP, AH, and IKE perform major functions within IPsec.
- IKEv1 phases 1 and 2 are enhanced by IKEv2.
- IPsec is mandatory and native to IPv6.

References

For additional information, refer to these resources.

Books

Carmouche, J. H. *IPSec Virtual Private Network Fundamentals* (Cisco Press, 2007).

Naganand, D., and Harkins, D. *IPSec: The New Security Standard for the Internet, Intranets, and Virtual Private Networks, Second Edition* (Prentice Hall, 2003).

Cisco.com Resources

“Cisco IOS IPsec,”
http://www.cisco.com/en/US/products/ps6635/products_ios_protocol_group_home.html
“Export Compliance & Regulatory Affairs: Encryption Control Guidance,”
<http://www.cisco.com/wvl/export/faq.html>

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

- 1.** Which of the following is not a security function provided by IPsec?
 - a. Confidentiality
 - b. Key management
 - c. Antireplay
 - d. Authentication
 - e. Integrity
- 2.** Which of the following are peer-authentication methods with IPsec? (Choose all that apply.)
 - a. PSK
 - b. RSA signatures
 - c. RSA encrypted nonces
 - d. HMAC-SHA1
 - e. HMAC-MD5
- 3.** Which three modes are used by IKE to secure communications?
 - a. Main mode

- b.** Basic mode
- c.** Advanced mode
- d.** Quick mode
- e.** Aggressive mode
- f.** Simple mode

4. Match the Diffie-Hellman group with the size of its keying material.

Groups

- a.** Group 5
- b.** Group 2
- c.** Group 7
- d.** Group 1
- e.** Group 16
- f.** Group 19

Bits

- g.** 4096
- h.** 163
- i.** 1024
- j.** 1536
- k.** 256
- l.** 768

5. Which statement is true regarding how ESP segments are encapsulated in IP packets?

- a.** In transport mode, security is provided only for the transport layer and below.
- b.** In tunnel mode, security is provided for the complete original IP packet.
- c.** In tunnel mode, security is provided only for the transport layer and above.
- d.** In transport mode, security is provided for the entire IP packet.

6. Match the IPsec component with its category.

- a.** ESP
- b.** IKE
- c.** EDCH
- d.** EDCSA
- 1.** Key Exchange
- 2.** Authentication
- 3.** Negotiation
- 4.** Confidentiality

7. In IPsec tunnel mode, the original IP address is used to route the packet through the untrusted network. True or false?

- a. True
- b. False

8. You require DoS protection and mobility for your IPsec phase 1 negotiation. Which protocol would you select?

- a. Diffie-Hellman.
- b. IKEv2
- c. IKEv1
- d. AH

9. In IPv6 environments, IPsec is mandatory for all IPv6 communications. True or false?

- a. True
- b. False

10. You want to prevent filtering IPsec traffic across the path of a VPN. Which three protocols should be allowed in the firewall policy?

- a. IP protocol 50
- b. TCP 50
- c. UDP 51
- d. TCP 51
- e. UDP 500
- f. UDP 4500

Chapter 14. Site-to-Site IPsec VPNs with Cisco IOS Routers

Building a site-to-site IPsec VPN is an essential part of many plans to meet the security requirements of customers. In this chapter, you will learn how to use the command-line interface (CLI) to configure a site-to-site IPsec VPN to securely connect two or more subnets over the Internet or an intranet.

This chapter teaches you how to configure a site-to-site IPsec VPN with preshared keys, using Cisco Configuration Professional. This ability includes being able to meet these objectives:

- Evaluate the requirements and configuration of site-to-site IPsec VPNs
- Use Cisco Configuration Professional to configure site-to-site IPsec VPNs
- Use CLI commands and Cisco Configuration Professional monitoring options to validate the VPN configuration
- Use CLI commands and Cisco Configuration Professional monitoring options to monitor and troubleshoot the VPN configuration

Site-to-Site IPsec: Planning and Preparation

Site-to-site VPNs are the option of choice for organizations of all kinds in implementing a corporate network across public and private networks. Internet-based VPN environments and Multiprotocol Label Switching (MPLS) VPN environments benefit from the flexibility of deployment and standards-based implementation of cryptographic mechanisms.

The choice of device-terminating VPNs, such as a firewall or a router, becomes a key factor in implementing site-to-site VPNs. Organizations benefit from leveraging their existing network elements and using an integrated approach to VPN deployments. This chapter highlights the use of Cisco IOS routers as site-to-site VPN termination points in IP Security (IPsec) environments.

Site-to-Site IPsec VPN Operations

IPsec VPN negotiation can be broken down into five steps, as shown in [Figure 14-1](#), including Phase 1 and Phase 2 of Internet Key Exchange (IKE):


Step 1. An IPsec tunnel is initiated when Host A sends “interesting” traffic to Host B. Traffic is considered interesting when it travels between the IPsec peers and meets the criteria that is defined in the crypto access control list (ACL).

Step 2. In IKE Phase 1, the IPsec peers (routers A and B) negotiate the established IKE SA policy. Once the peers are authenticated, a secure tunnel is created using ISAKMP.


Step 3. In IKE Phase 2, the IPsec peers use the authenticated and secure tunnel to negotiate IPsec SA transforms. The negotiation of the shared policy determines how the IPsec tunnel is established.

Step 4. The IPsec tunnel is created and data is transferred between the IPsec peers based on the IPsec parameters configured in the IPsec transform sets.


Step 5. The IPsec tunnel terminates when the IPsec SAs are deleted or when their lifetime expires.




1. Host A sends interesting traffic to Host B.
2. Router IOS-FW and Router Site-1 negotiate an IKE Phase 1 session.



3. Router IOS-FW and Router Site-1 negotiate an IKE Phase 2 session.



4. Hosts A and B communicate via IPsec tunnel.



5. IPsec tunnel is terminated.

Figure 14-1. Site-to-Site IPsec VPN

Planning and Preparation Checklist

The following items should be considered in planning and preparing for a site-to-site VPN deployment:

- **Verify connectivity between peers:** IPsec uses distinct protocols and ports for its operations. Ensure that protocol 50 (Encapsulating Security Payload [ESP]), protocol 51 (Authentication Header [AH]), and UDP port 500 (ISAKMP) traffic is not blocked at interfaces that are used by IPsec.
- **Define interesting traffic:** It is common in site-to-site VPNs to find deployments where some traffic will use the VPN and some traffic will bypass it. This concept is known as *split tunneling*. It is important to define which traffic will trigger the establishment of the VPN (interesting traffic) and which traffic will travel outside the tunnel (noninteresting traffic).
- **Determine the cipher suite requirements:** The choice of encryption algorithms for Phase 1 and Phase 2 of IPsec, as well as for the different services (confidentiality, integrity, and authentication), define the strength of protection and capacity planning considerations. The cipher suite is defined by the assets the organization is trying to protect and the level of risk it is willing to take. However, other considerations apply, such as compliance regulations. It is important to know that the cipher suite can be different for different peers. Multiple policies can be built on a given device, and they are negotiated differently with different peers. The resulting tunnels will provide varying levels of risk.

- **Manage monitoring, troubleshooting, and change:** It is important to build all three areas in the planning process. This chapter illustrates some of the tools that are available to accomplish these tasks.

Building Blocks of Site-to-Site IPsec

IPsec has specific procedures that require configuration on the security appliance for the successful establishment of a VPN tunnel.

Interesting Traffic and Crypto ACLs

Interesting traffic is defined by crypto ACLs in site-to-site IPsec VPN configurations. Crypto ACLs perform these functions, illustrated in [Figure 14-2](#):

- **Outbound:** For outbound traffic, the crypto ACL defines the flows that IPsec should protect. Traffic that is not selected is sent in plaintext.
- **Inbound:** The same ACL is processed for inbound traffic. The ACL defines traffic that should have been protected by IPsec, and discards packets if they are selected but arrive unprotected (unencrypted).

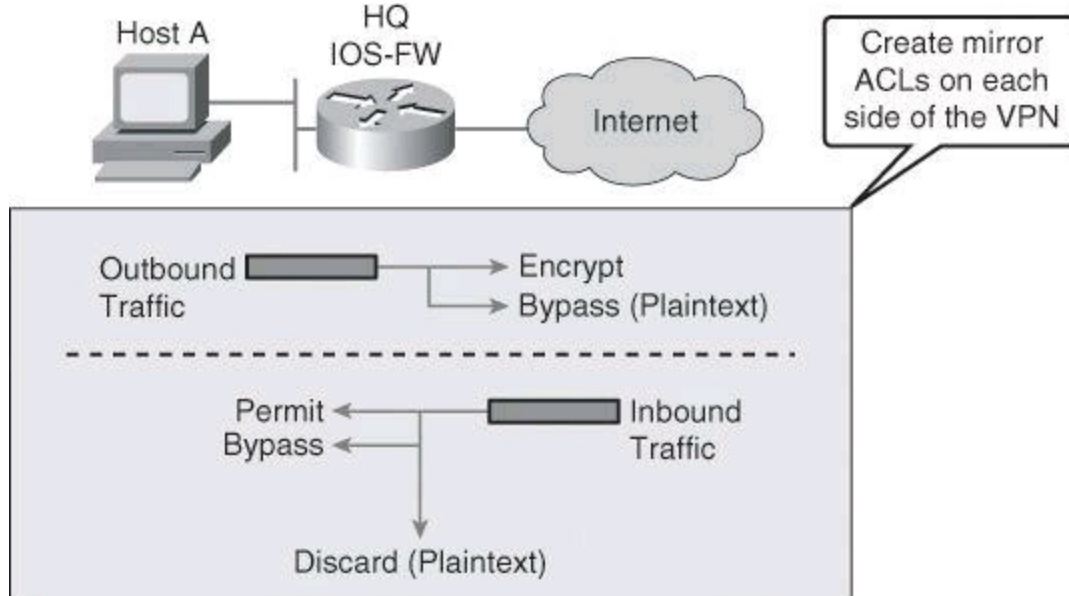


Figure 14-2. Outbound and Inbound Access Control Lists

Ultimately, traffic will be processed against ACL statements. Traffic matching a **permit** action in the ACL will be protected and sent through the VPN. Traffic matching a **deny** action in the ACL will be sent outside the VPN in cleartext.

These recommendations apply when defining interesting traffic:

- Try to be as restrictive as possible when defining which packets to protect in a crypto ACL. If you must use the **any** keyword in a **permit** statement, you must preface that statement with a series of **deny** statements to filter out any traffic (that would otherwise fall within that **permit** statement) that you do not want to be protected.
- Because crypto ACLs are processed for both outbound and inbound traffic, it is important to create mirrored crypto ACLs in each VPN endpoint.

Mirrored Crypto ACLs

When crypto ACLs are incorrectly configured or missing, traffic might only flow in one direction across the VPN tunnel, or it might not be sent across the tunnel at all. When a router receives encrypted packets from an IPsec peer, it uses the same ACL to determine which inbound packets to decrypt by viewing the source and destination addresses in the ACL in reverse order. Any unprotected inbound traffic that matches a **permit** entry in the crypto ACL for a **crypto map** entry that is highlighted as IPsec is dropped. This drop occurs because this traffic was expected to be protected by IPsec.

In [Figure 14-3](#), the VPN protects traffic from both LANs behind each router acting as VPN peer. The crypto ACL created on router Site-1 is the exact mirror of the crypto ACL created on router IOS-FW. The definition of source traffic on router IOS-FW becomes the definition of destination traffic on router Site-1. Similarly, the definition of destination traffic on router IOS-FW becomes the definition of source traffic on router Site-1.

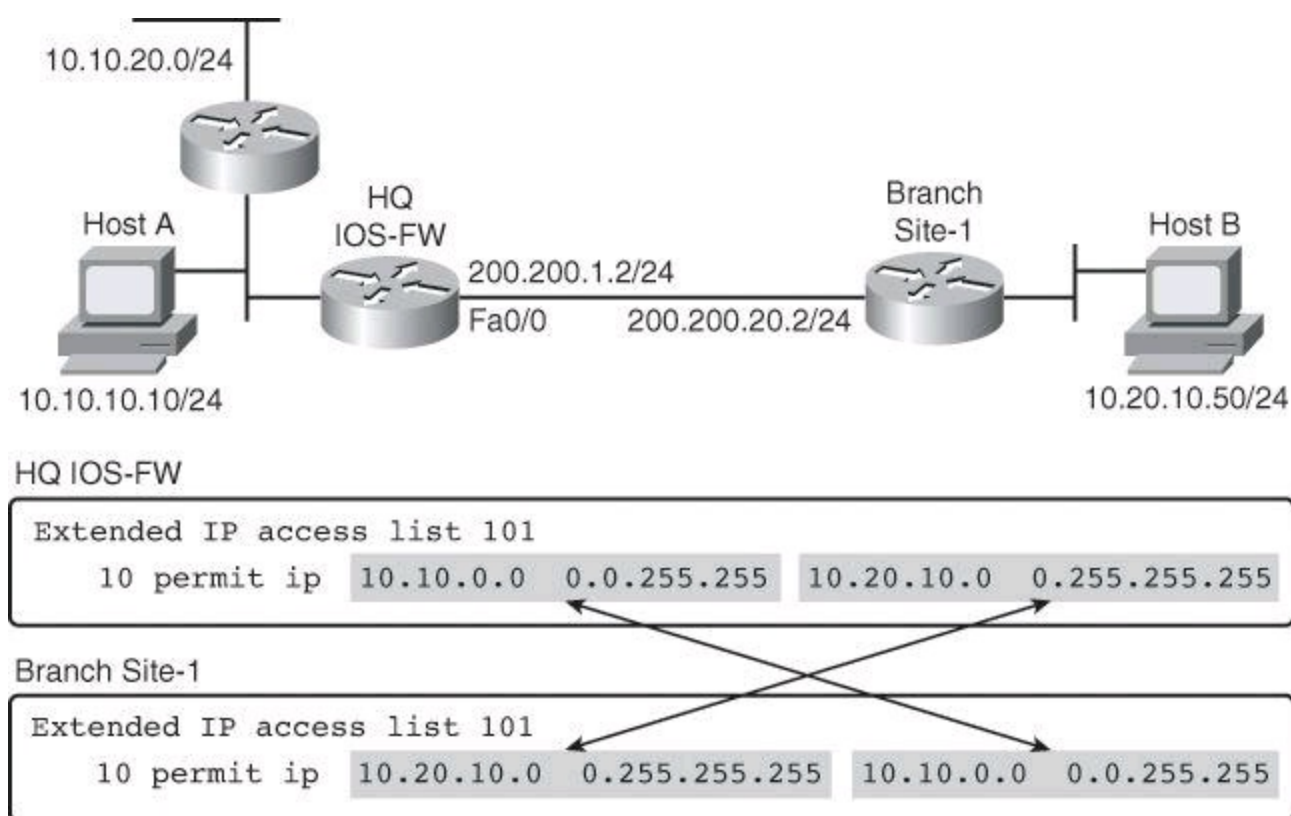


Figure 14-3. Mirrored Crypto ACLs

You will notice in [Figure 14-3](#) that we have included all the HQ subnets in our crypto ACLs, thus the wildcard mask of 16 bits for 10.10.0.0. The branch office has only one subnet, 10.20.10.0, thus the wildcard mask of 24 bits.

Cipher Suite

The cipher suite selection should follow guidelines in terms of encryption algorithms, key sizes, and key lifetimes. [Table 14-1](#) illustrates an example. Advanced Encryption Standard (AES), for instance, is considered a stronger encryption algorithm than Triple Data Encryption Standard (3DES). Similarly, Secure Hash Algorithm 2 (SHA-2) is considered a stronger hashing function than SHA-1 and Message Digest 5 (MD5). In fact, some of these algorithms might have been broken already, in real networks or in a lab.

Table 14-1. Example of Cipher Suite Selection Decision

Parameter	Weaker	Stronger
Encryption algorithm	3DES	AES
Hash algorithm	MD5	SHA-1/SHA-2
Authentication method	Preshare	RSA signature
Diffie-Hellman key exchange	DH Group 1	DH Group 5
IKE SA lifetime	86,400 seconds	< 86,400 seconds

Similar considerations apply to authentication methods, where digital certificates are considered a stronger method to uniquely identify VPN peers, and they are a more scalable solution in the presence of a high volume of peers. Diffie-Hellman (DH) groups define the size of encryption keys, and will have an effect on device performance.

Key lifetimes also define the level of risk and the strength of the cryptography solution. The shorter the lifetimes, the more frequently peers will perform rekeying and refresh the cryptography materials, mitigating the effect of man-in-the-middle attacks and other exploits that expose keys and other components.

Crypto Map

Crypto map entries that you create for IPsec combine the needed configuration parameters of IPsec SAs, including the following parameters:

- Which traffic should be protected by IPsec using a crypto ACL
- The granularity of the flow to be protected by a set of SAs
- Who the remote IPsec peer is, which determines where the IPsec-protected traffic is sent
- The local address that is to be used for the IPsec traffic (optional)
- Which IPsec security should be applied to this traffic, choosing from a list of one or more transform sets

Crypto maps, illustrated in [Figure 14-4](#), are applied to router interfaces in order for the IPsec VPN to start working. You can apply only one crypto map set to a single interface. Crypto maps can negotiate different policies for different peers on the same interface. In that case, the same crypto map will have multiple entries. These multiple entries are known as a *crypto map set*.

Crypto maps define the following:

- Interesting traffic (crypto ACLs)
- Remote VPN peers
- Cipher policy (transform set)
- Key management method
- SA lifetimes




Figure 14-4. Crypto Map and Its Role

You must create multiple crypto map entries for a given interface if any of these conditions exist:

- Different data flows are to be managed by separate IPsec peers.
- You want to apply different IPsec security to different types of traffic (to the same or separate IPsec peers). An example would be if you want traffic between one set of subnets to be authenticated and traffic between another set of subnets to be both authenticated and encrypted. In this case, you should define the different types of traffic in two separate ACLs, and you must create a separate crypto map entry for each crypto ACL.
- If you are not using IKE to establish a particular set of SAs, and you want to specify multiple ACL entries, you must create separate ACLs (one per permit entry) and specify a separate crypto map entry for each ACL.

Configuring a Site-to-Site IPsec VPN Using CCP

The following pages show you how to configure a site-to-site IPsec VPN with preshared keys by using Cisco Configuration Professional wizards. All screen shots and configuration examples that follow use the same scenario that is depicted in [Figure 14-5](#). The goal is to establish a site-to-site VPN between router IOS-FW and router Site-1, protecting traffic that travels between the two LAN segments behind each router.




Figure 14-5. Scenario for Configuring a Site-to-Site IPsec VPN with Preshared Keys Using CCP VPN Wizard

To ensure secure communications, a strong cipher suite is required, which involves using preshared keys, AES, SHA-1, and DH5, as an example.

Initiating the VPN Wizard

The Cisco Configuration Professional VPN wizard can be used to simplify the process of configuring the site-to-site VPN. The wizard is located by following this option path: **Configure > Security > VPN > Site-to-Site VPN**.

The Create Site to Site VPN tab includes two types of wizard:

- **Create a Site to Site VPN:** This option allows you to create a VPN network connecting two routers.
- **Create a Secure GRE Tunnel (GRE over IPsec):** This option allows you to configure a Generic Routing Encapsulation (GRE) protocol tunnel between your router and a peer system.

As shown in [Figure 14-6](#), click **Create a Site-to-Site VPN** and click **Launch the Selected Task**.

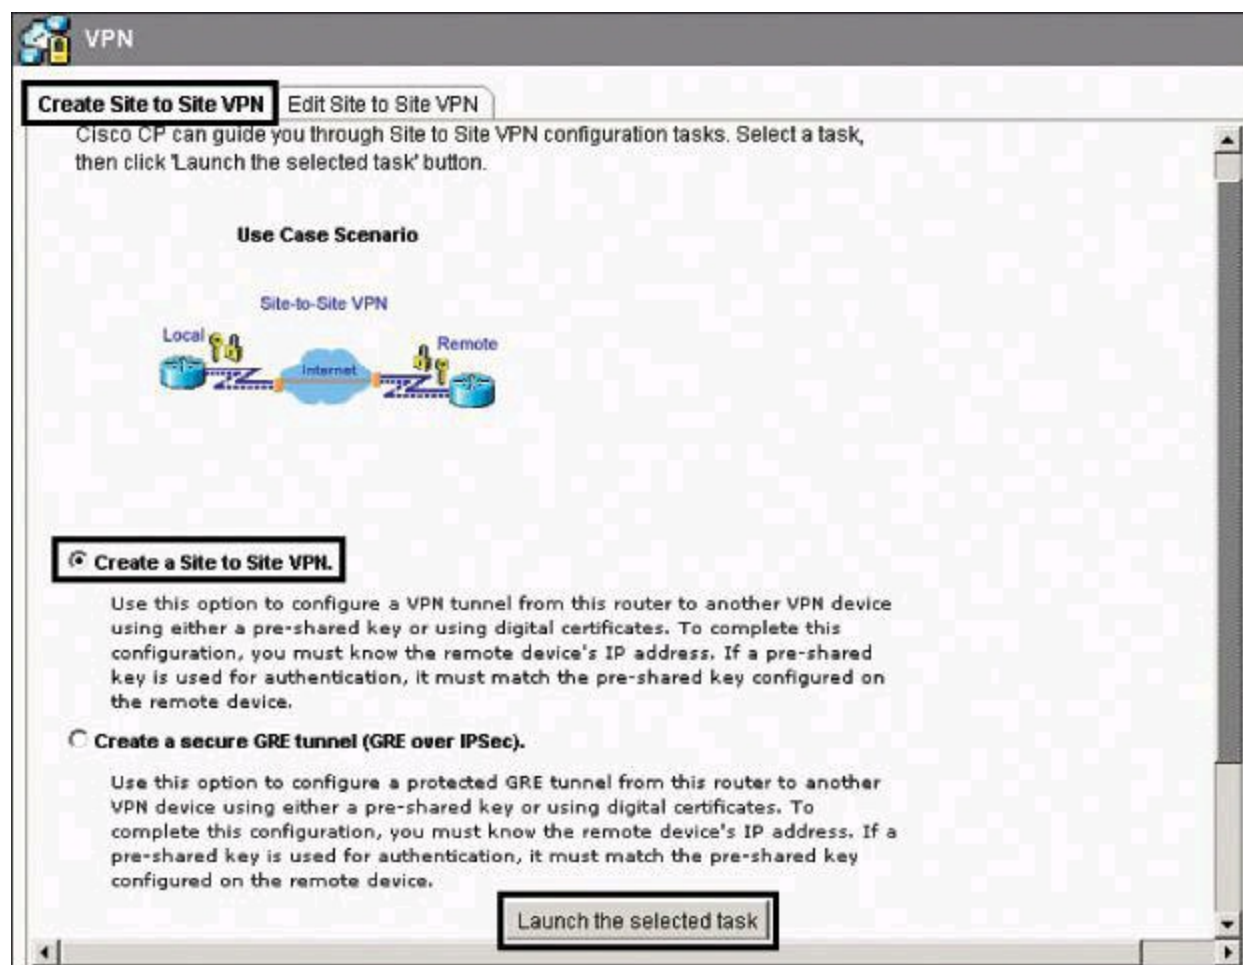
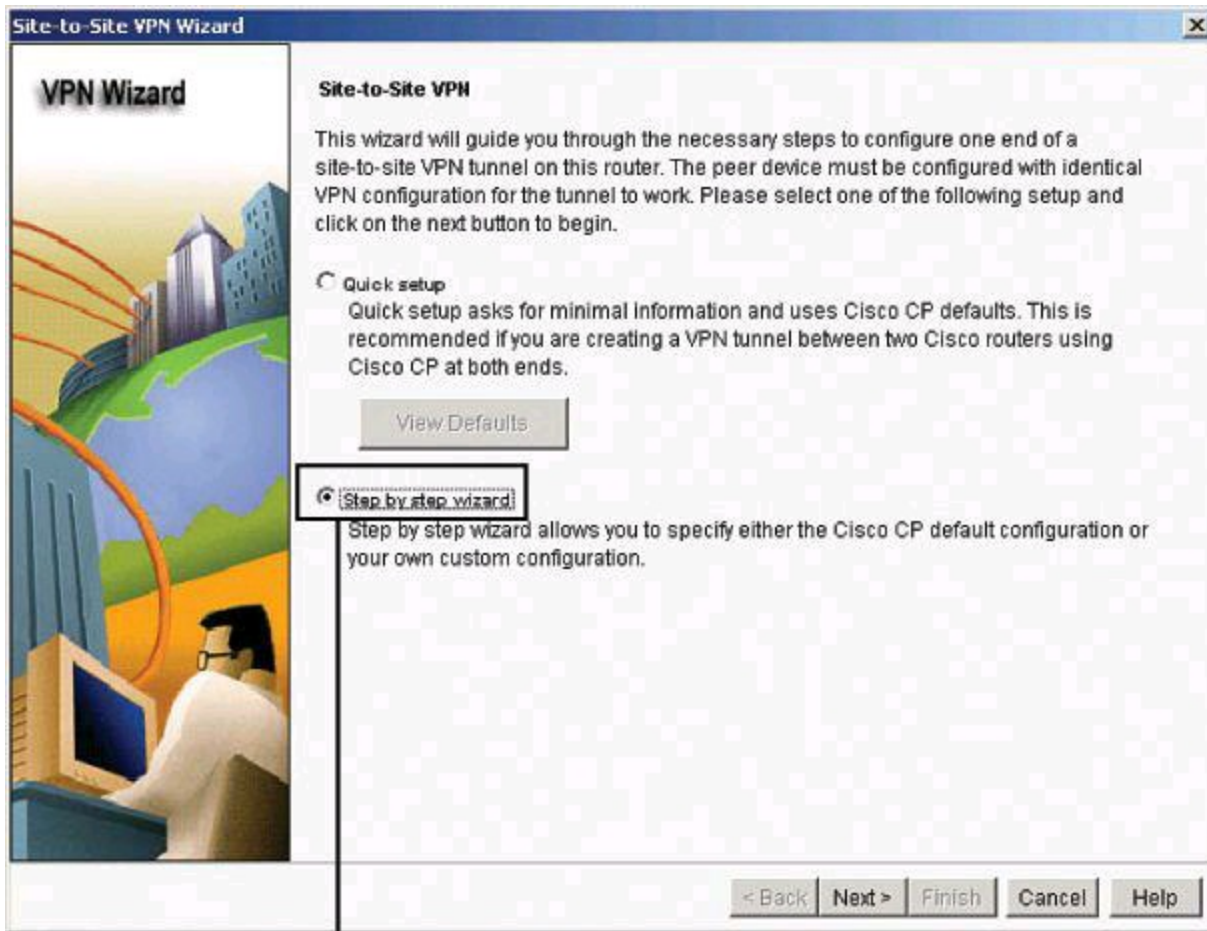


Figure 14-6. Launching CCP Site-to-Site VPN Wizard

As shown in [Figure 14-7](#), the Site-to-Site VPN Wizard starts with a choice of quick setup or a step-by-step wizard:

- Using quick setup, CCP will automatically provide a default IKE policy to govern authentication, a default transform set to control the encryption of data, and a default IPsec rule that will encrypt all traffic between the router and the remote device. Quick setup is best used when both the local router and the remote system are Cisco routers using CCP. Quick setup will configure 3DES encryption if it is supported by the Cisco IOS image. Otherwise, it will configure DES encryption. If you need AES or Software-Optimized

Encryption Algorithm (SEAL) encryption, you must use the step-by-step wizard. Using quick setup, you can view the default IKE policy, transform set, and IPsec rule that will be used to configure the one-step VPN.



Step-by-Step Wizard Enables
You to Change the Defaults

Figure 14-7. Wizard Gives a Choice Between Quick Setup or Step-by-Step Approach

- Using the step-by-step wizard, you can configure a site-to-site VPN using parameters that you specify. The result is a custom configuration for the VPN, allowing the use of any of the CCP defaults that you need. The step-by-step wizard, selected in [Figure 14-7](#), allows you to specify stronger encryption than the quick setup wizard allows.

VPN Connection Information

You can use the VPN Connection Information page of the Site-to-Site VPN Wizard, shown in [Figure 14-8](#), to identify the IP address or hostname of the remote site that will terminate the VPN tunnel that you are configuring, to specify the router interface to use, and to enter the preshared key that both routers will use to authenticate each other.

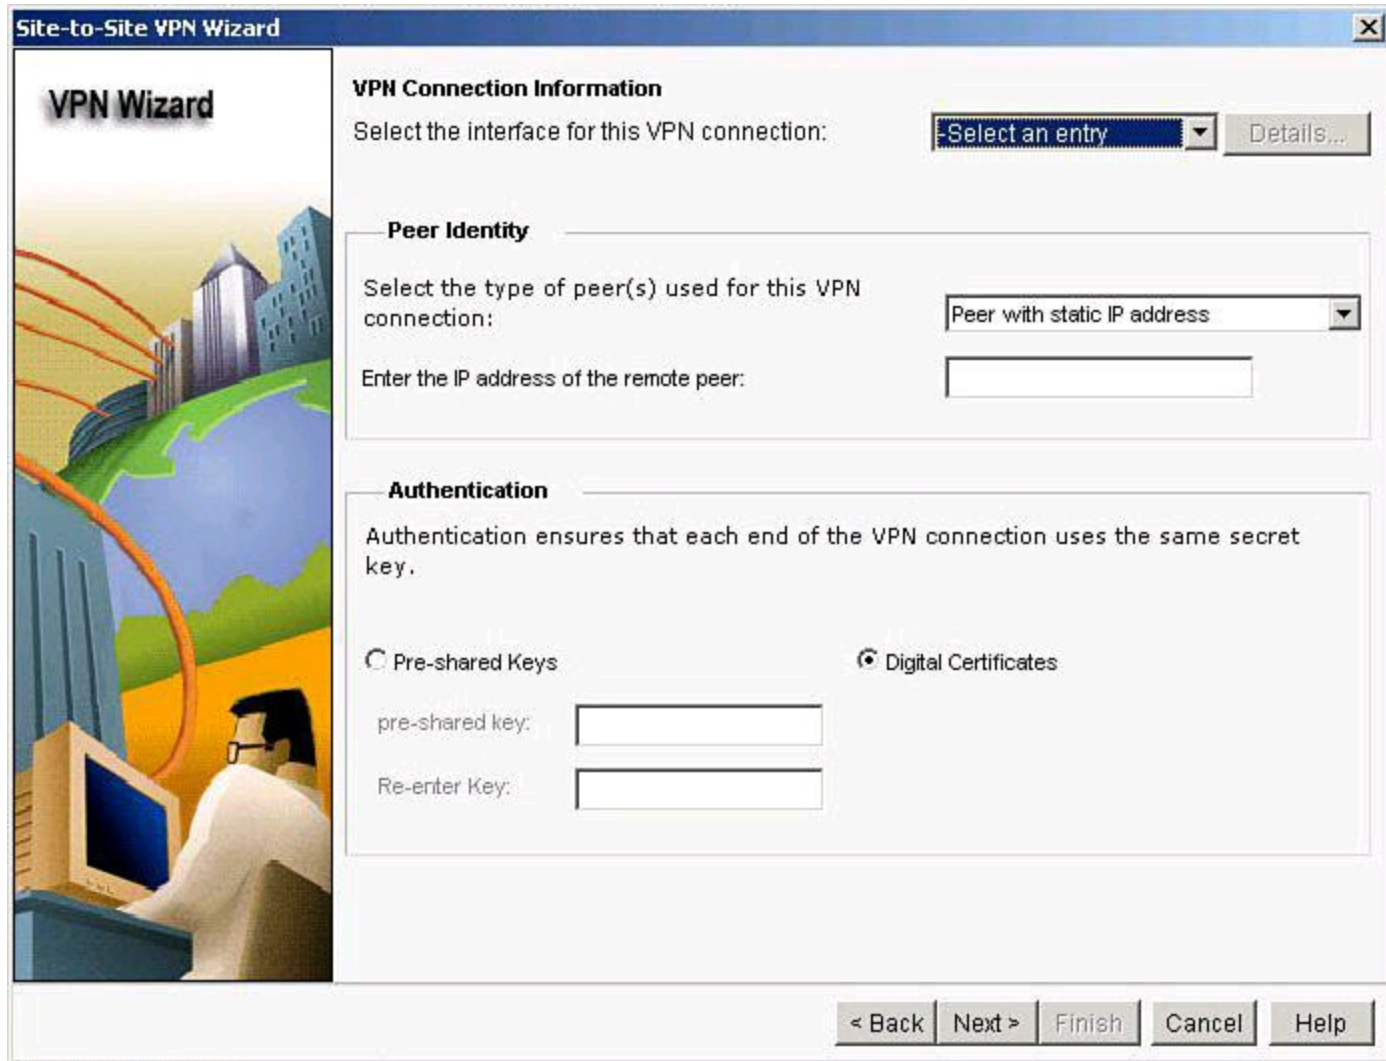


Figure 14-8. VPN Connection Information Page

The VPN Connection Information page of the wizard has three distinct sections, as shown in [Figure 14-8](#).

The first component of the VPN Connection Information page is the local router interface that will initiate and terminate the tunnel, shown in [Figure 14-9](#). From the interface drop-down menu, select the interface on this router that connects to the remote site. The router presented in [Figure 14-9](#) is the HQ IOS-FW. As also shown in [Figure 14-9](#), you can click Details to view the details of the interface, in terms of other configured features, such as Network Address Translation (NAT), existing ACLs, quality of service (QoS) rules, or Cisco IOS Intrusion Prevention System (IPS) rules. This is important to determine the compatibility of those features with IPsec configurations on the same interface.

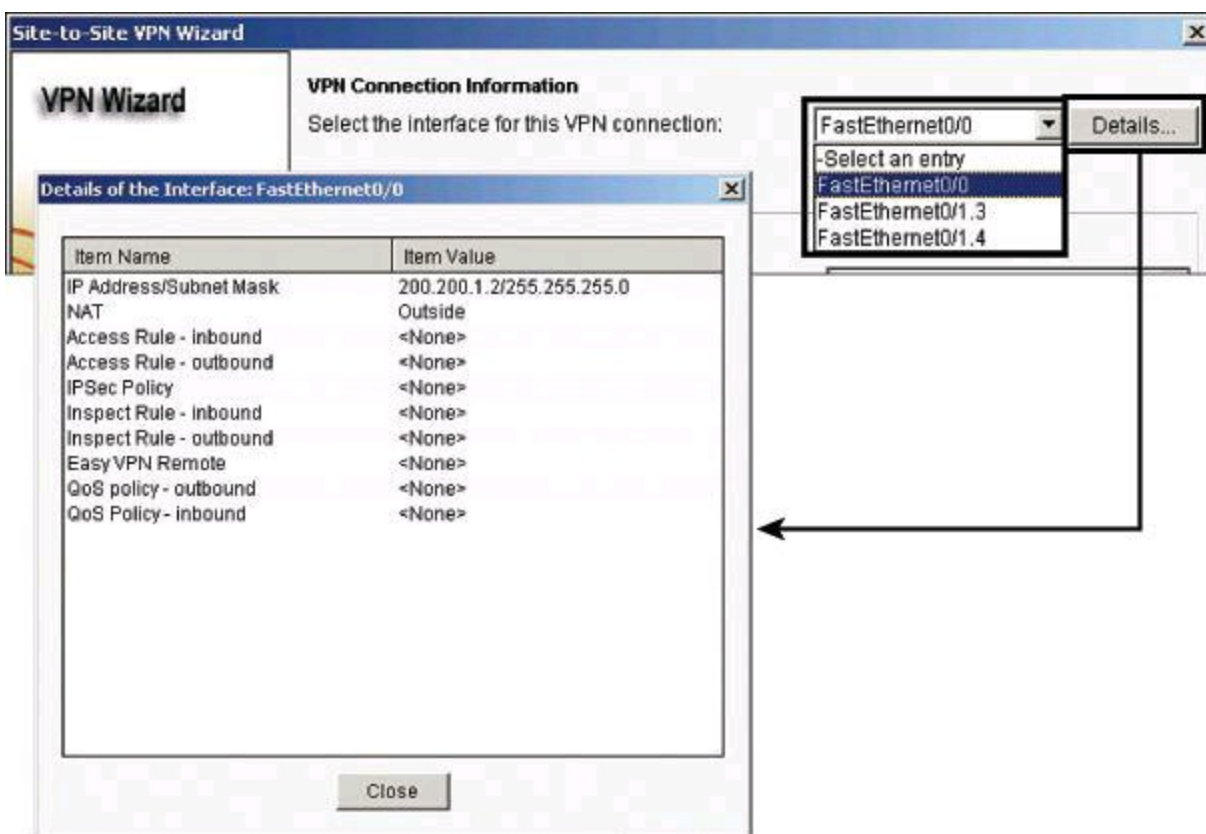


Figure 14-9. First Component of VPN Connection Information Page: Interface Selection

The second component of the VPN Connection Information page is the Peer Identity section, shown in [Figure 14-10](#). Enter the IP address of the remote IPsec peer that will terminate the VPN tunnel that you are configuring. The remote IPsec peer might be another router, a VPN concentrator, or any other gateway device that supports IPsec. Three options are available in the Peer Identity section, the first two of which are located in the drop-down menu:

- **Peer(s) with dynamic IP addresses:** Choose this option if the peers that the router connects to use a dynamically assigned IP address.
- **Peer with static IP address:** Choose this option if the peer that the router connects to uses a fixed IP address.
- **Enter the IP address of the remote peer:** This is enabled when Peer With Static IP Address is chosen. Enter the IP address of the remote peer.

Select Peer with Static IP Address from the drop-down.

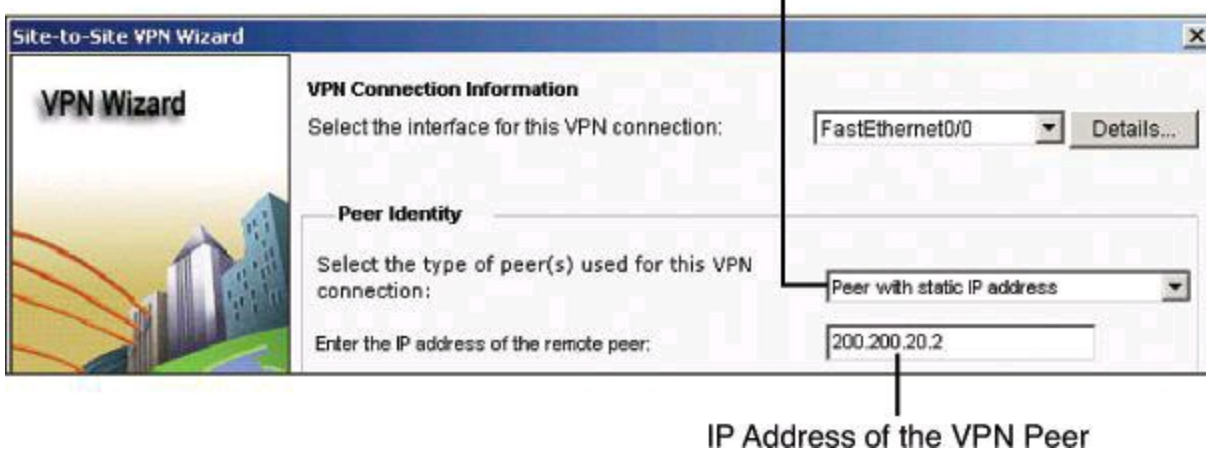


Figure 14-10. Second Component of VPN Connection Information Page: Peer Identity

The third component of the VPN Connection Information page, the Authentication section, is where you specify the preshared key used for IPsec authentication, as shown in [Figure 14-11](#).

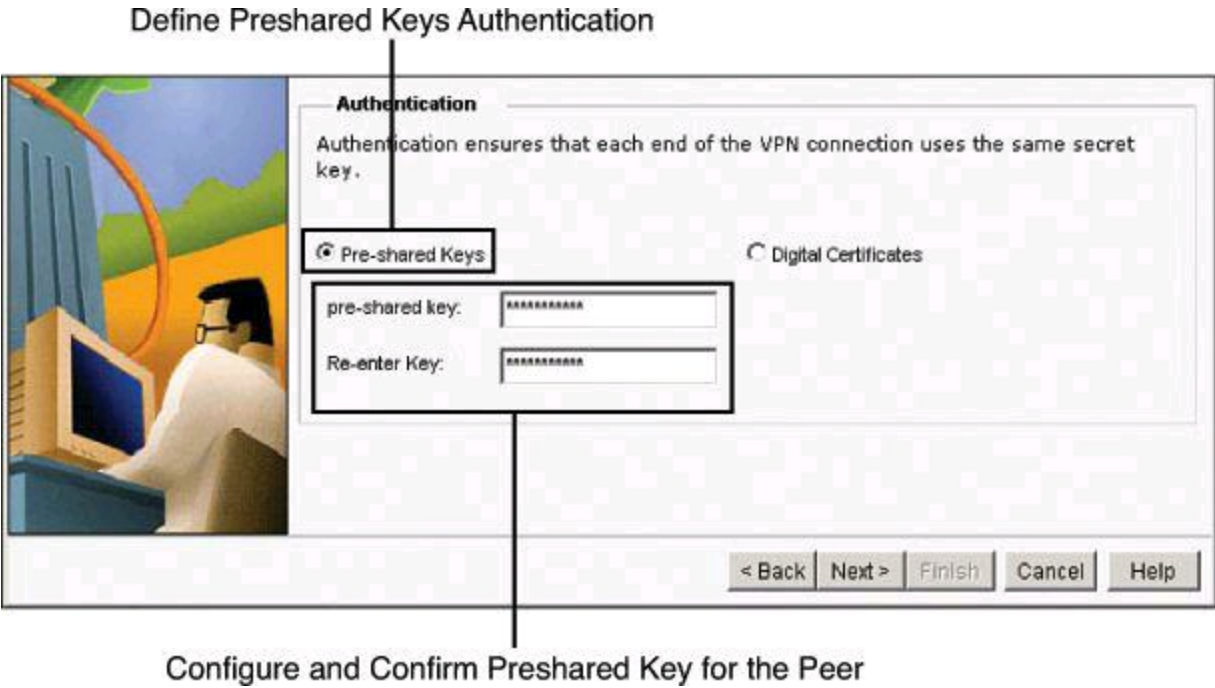


Figure 14-11. Third Component of VPN Connection Information Page: Authentication

In the Authentication section of this page, you can click the Pre-shared Keys button if the VPN peers use a preshared key to authenticate connections from each other. This key must be the same on each side of the VPN. Question marks (?) and spaces must not be used in the preshared key. The preshared key can contain a maximum of 128 characters. Optionally, you can configure digital certificates as the authentication option. Click Next to move to the next page.

IKE Proposals

The next wizard page, IKE Proposals, shown in [Figure 14-12](#), lists all of the IKE policies that have been configured on the router. If no user-defined policies have been configured, the page lists the Cisco Configuration Professional default IKE policy. IKE policies govern the way in which devices in a VPN authenticate themselves during Phase 1. The local router uses the IKE policies listed on this page to negotiate authentication with the remote router.

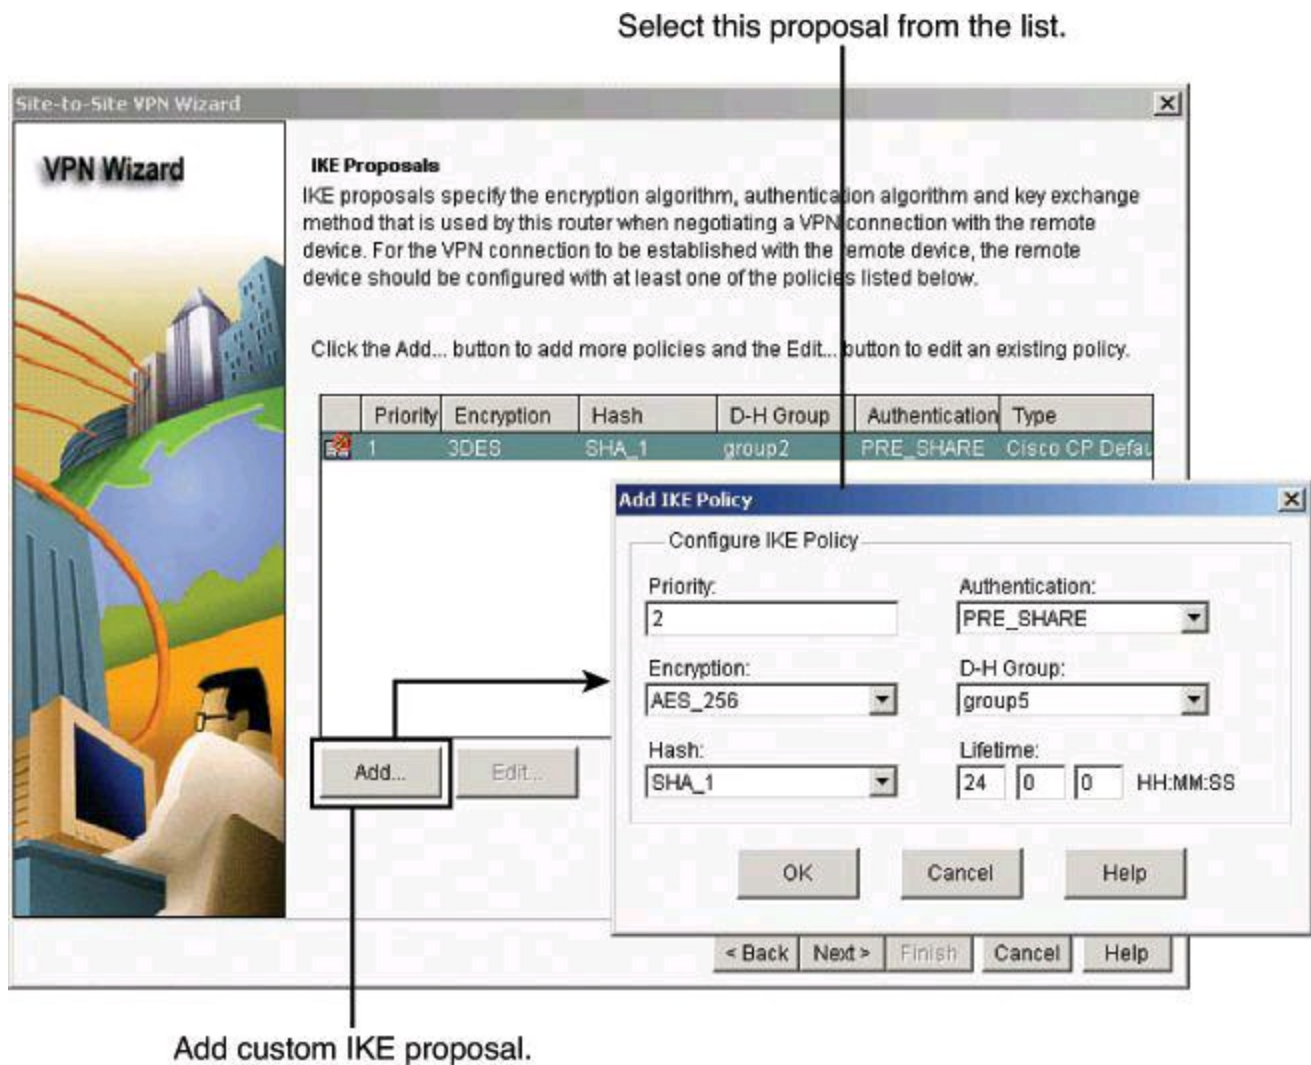


Figure 14-12. IKE Proposals Configured Through the VPN Wizard

The local router and the peer device must both use the same policy. The router that initiates the VPN connection offers all its policies to the receiver, which analyzes them starting with the lowest priority number first. If the remote system rejects that policy, it considers the policy with the next lowest number, and continues in this fashion until it finds a match, as shown previously in Figure 13-8. You must coordinate closely with the administrator of the peer system so that you can configure identical policies on both routers.

If you want to add an IKE policy that is not included in this list, you can click Add and create the policy in the Add IKE Policy dialog box that is displayed. You can edit an existing policy by selecting it and clicking Edit. CCP default policies are read-only and cannot be edited. In [Figure 14-12](#), the Add IKE Policy dialog box displays the desired policy for the requirements, which are encryption using AES, authentication using preshared keys, hashing using SHA-1, and key exchange using DH 5.

Note

Adding a new policy creates a new line with a higher priority in the IKE policy list. In the Add IKE Policy dialog box, you can control the priority number and define the cipher suite of your choice. Be sure to select the new policy from the list of IKE policies before you click **Next** to continue.

Transform Set

The next step of the VPN Wizard is to define the transform sets, which define the cipher suite for IPsec Phase 2. The Transform Set page, shown in [Figure 14-13](#), lists the CCP default transform sets and the additional transform sets that have been configured on this router. These transform sets will be available for use by the VPN. A *transform set* represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow. A *transform* describes a particular security protocol with its corresponding algorithms.

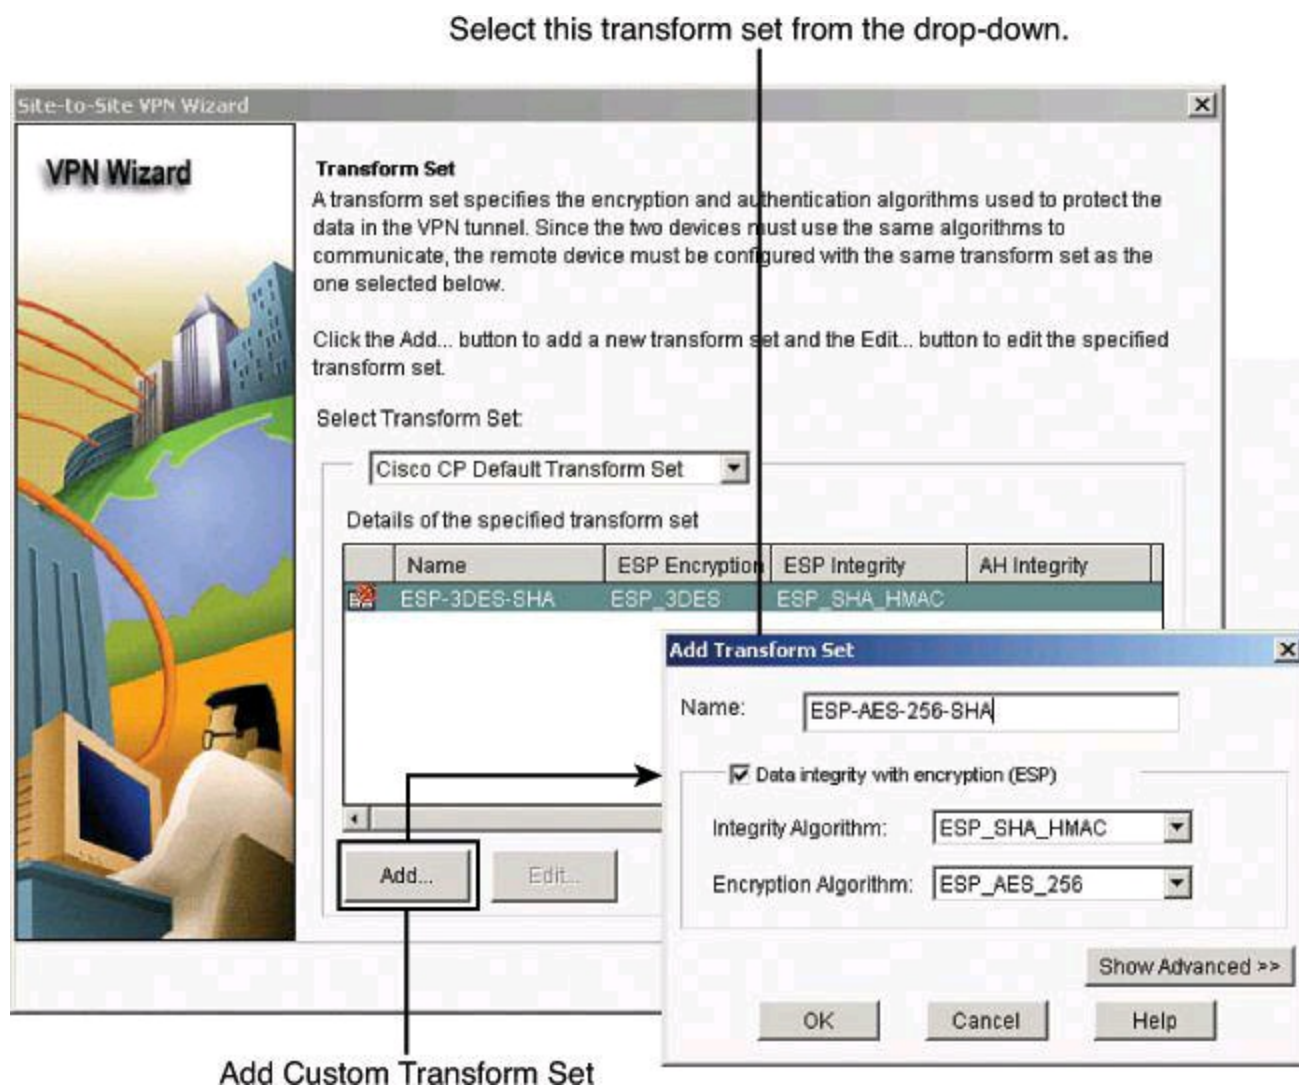


Figure 14-13. Transform Set Configured Through the VPN Wizard

A list of all configured transform sets is available in the Select Transform Set drop-down menu. Selecting one allows you to view the settings and add entries to the transform set.

Click **Add** to create the transform set in the Add Transform Set dialog box. You may choose to perform this action if the router is to negotiate different Phase 2 policies for different peers. This is common in hub-and-spoke VPN topologies. For nondefault transform sets, you can also click Edit to modify the settings.

In [Figure 14-13](#), the Add Transform Set dialog box is configured with the desired policy. ESP is the IPsec protocol that will define authentication, integrity, and payload encryption to create Phase 2 tunnels, using SHA-1 as the integrity algorithm, and AES as the encryption algorithm. The default is tunnel mode.

Click **Next** to continue the wizard and the VPN configuration.

Note

Adding a new policy will add options to the Select Transform Set drop-down menu. Be sure to select the newly created policy from this drop-down menu before you click **Next** to continue with the wizard.

Traffic to Protect

The next step of the wizard is to define interesting traffic. The Traffic to Protect page, shown in [Figure 14-14](#), lets you define the traffic that this VPN protects. The VPN can protect traffic between specified subnets, or protect the traffic that is specified in an IPsec rule that you select. There are two options available:

- **Protect all traffic between the following subnets:** Use this option to specify a single source subnet (a subnet on the LAN) whose outgoing traffic you want to encrypt, and one destination subnet that is supported by the peer that you specified on the VPN Connection Information page. All traffic flowing between other source and destination pairs will be sent unencrypted.
- **Create/Select an access-list for IPsec traffic:** Use this option if you need to specify multiple sources and destinations, or if you need to encrypt specific types of traffic. An IPsec rule can consist of multiple entries, each specifying different traffic types and different sources and destinations. Click the ... button next to the field and specify an existing IPsec rule that defines the traffic you want to encrypt, or create an IPsec rule to use for this VPN.

Define Destination (Remote Network) for Protected Traffic

VPN Wizard

Traffic to protect
IPSec rules define the traffic, such as file transfers (FTP) and e-mail (SMTP) that will be protected by this VPN connection. Other data traffic will be sent unprotected to the remote device. You can protect all traffic between a particular source and destination subnet, or specify an IPSec rule that defines the traffic types to be protected.

Protect all traffic between the following subnets

Local Network
Enter the IP address and subnet mask of the network where IPSec traffic originates.

IP Address: 10.10.0.0
Subnet Mask: 255.255.0.0 or 16

Remote Network
Enter the IP Address and Subnet Mask of the destination Network.

IP Address: 10.20.10.0
Subnet Mask: 255.255.255.0 or 24

Create/Select an access-list for IPsec traffic

< Back Next > Finish Cancel Help

Define Source (Local Network) for Protected Traffic

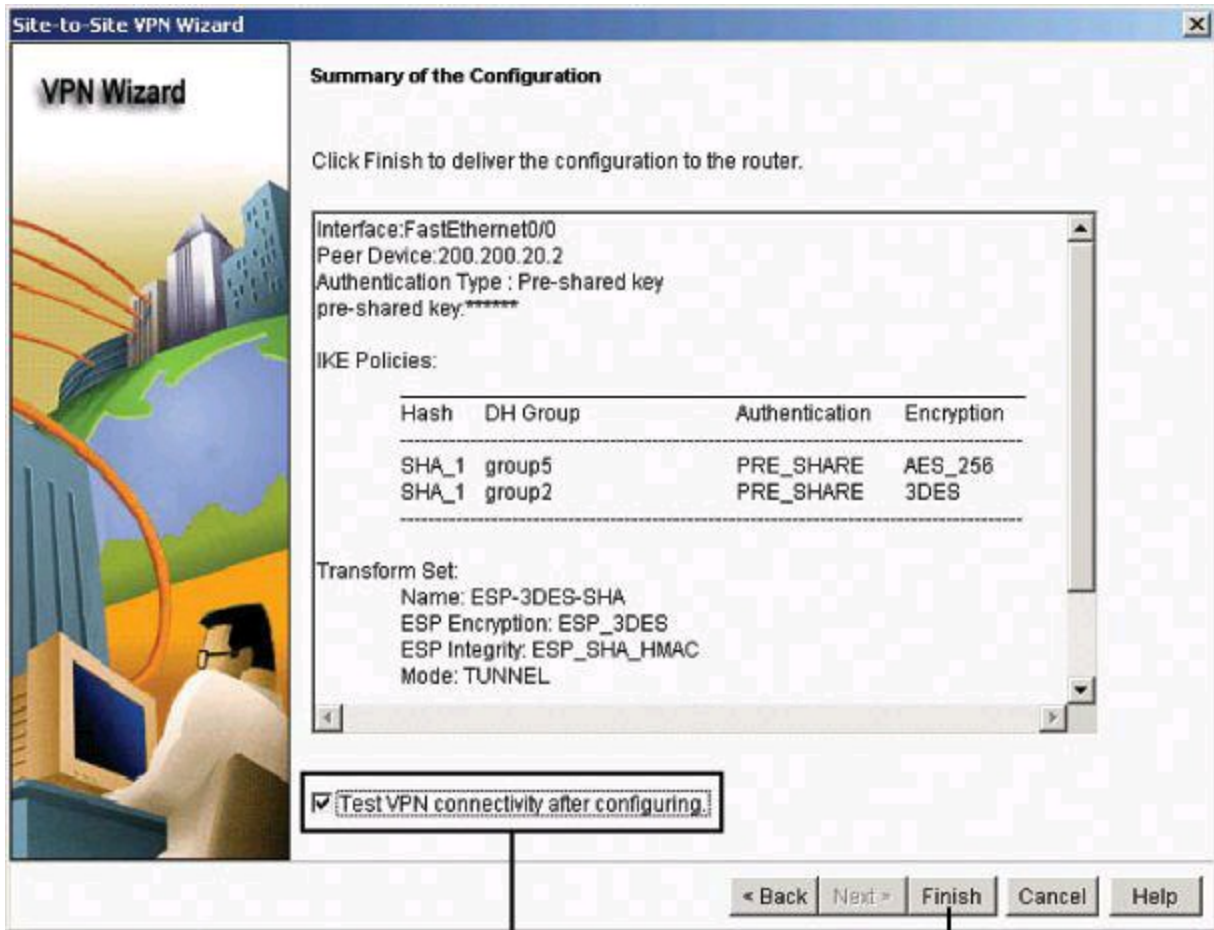
Figure 14-14. Protecting Traffic Through the VPN Wizard

In [Figure 14-14](#), the protected networks from the perspective of router IOS-FW are 10.10.0.0/16 as the source (the LAN behind IOS-FW) and 10.20.10.0/24 as the destination (the LAN behind router Site-1).

Click **Next** to continue the wizard and the VPN configuration.

Configuration Summary

The last step of the wizard is to review and accept the configuration, as shown in [Figure 14-15](#). This page shows you the VPN configuration that you created. After you review the configuration on this page, you can either click the Back button to make changes if you want, or click Finish to complete the wizard. Note that a testing option is available. You can check the Test VPN Connectivity After Configuring check box to test the VPN connection you have just configured. The results of the test will be shown in another window. This option is demonstrated later in this chapter.



Check to test VPN connectivity.

Review configuration and click Finish.

Figure 14-15. Summary of the Site-to-Site VPN Wizard Configuration

Note

If NAT is configured on this router, a CCP warning box, shown in [Figure 14-16](#), will appear after you click **Finish**; this warning will offer to convert your NAT rules with route maps in order for the VPN to work properly. CCP assumes, rightly so, that you would like to skip the NAT process for traffic going through the IPsec tunnel. Therefore, traffic from 10.10.0.0/16 will appear with its original (real) IP address when transiting on the 10.20.10.0/24 network and vice versa. In other words, sessions between 10.10.0.0/16 and 10.20.10.0/24 will not have their source address translated.

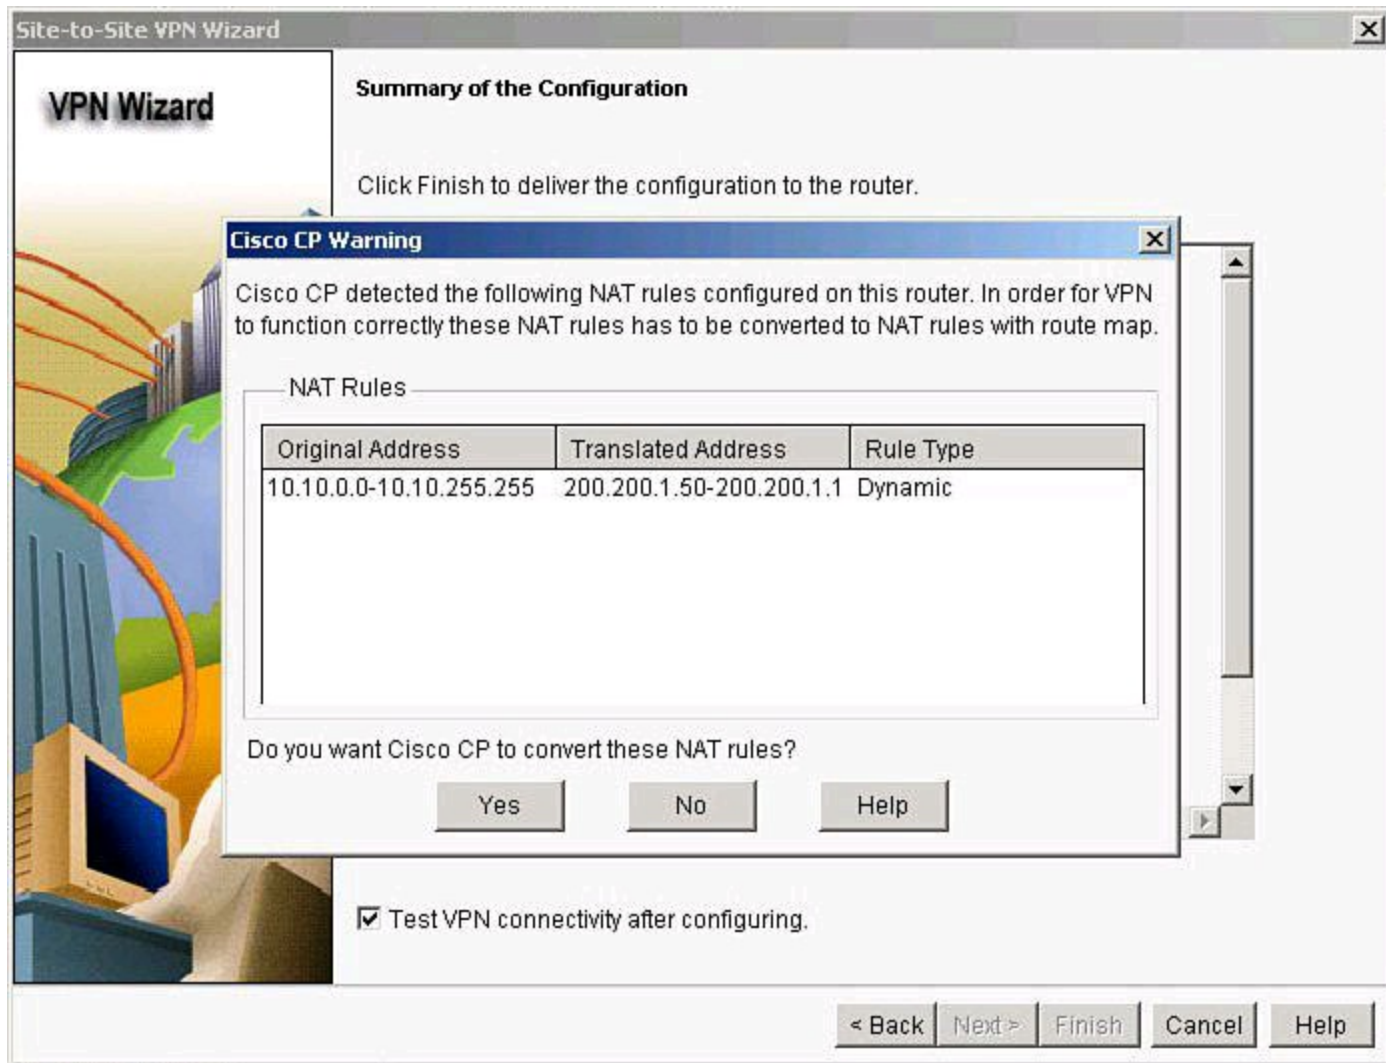


Figure 14-16. VPN Wizard Warning About NAT Rules

Creating a Mirror Configuration for the Peer Site

Site-to-site VPNs require both sides to be configured with matching policies. You also need to create mirrored crypto ACLs in each VPN endpoint and ensure that the peers know each other's IP addresses. Typically, the same process that is configured by the wizard has to be replicated on the other peer. In Cisco Configuration Professional, you can navigate to **Configure > Security > VPN > Site-to-Site VPN**, click the **Edit Site-to-Site VPN** tab, and then click **Generate Mirror** to automatically create a mirror configuration for the peer.

Clicking this button, as shown in [Figure 14-17](#), creates a text file that captures the VPN configuration of the local router so that a remote router can be given a VPN configuration that enables it to establish a VPN connection to the local router. This button is disabled if you have selected a dynamic site-to-site VPN tunnel.

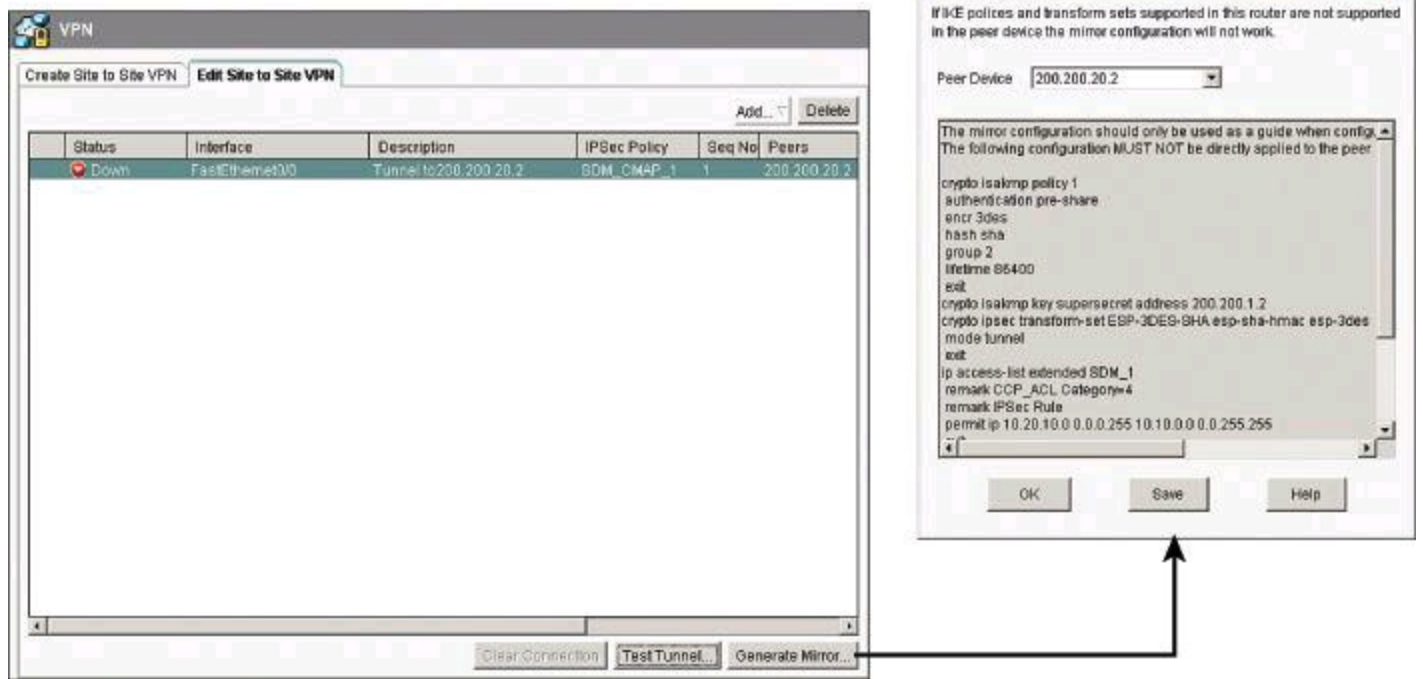


Figure 14-17. Creating a Mirror Configuration for the Peer Router

Note

The text file that you generate must not be copied into the configuration file of the remote system, but it must be edited and used only to show what has been configured on the local router so that the remote device can be configured in a way that is compatible. Identical names for IPsec policies, IKE policies, and transform sets may be used on the remote router, but the policies and transform sets may be different. If the text file is simply copied into the remote configuration file, configuration errors are likely to result.

Note

Any previously configured VPN connections that are detected by Cisco Configuration Professional that do not use ISAKMP crypto maps will appear as read-only entries in the VPN connection table and cannot be edited.

Verifying the IPsec Configuration Using CCP and CLI

You can verify the VPN configuration on the Edit Site to Site VPN tab, as shown in [Figure 14-18](#). Use this tab to manage the VPN connections to remote systems.

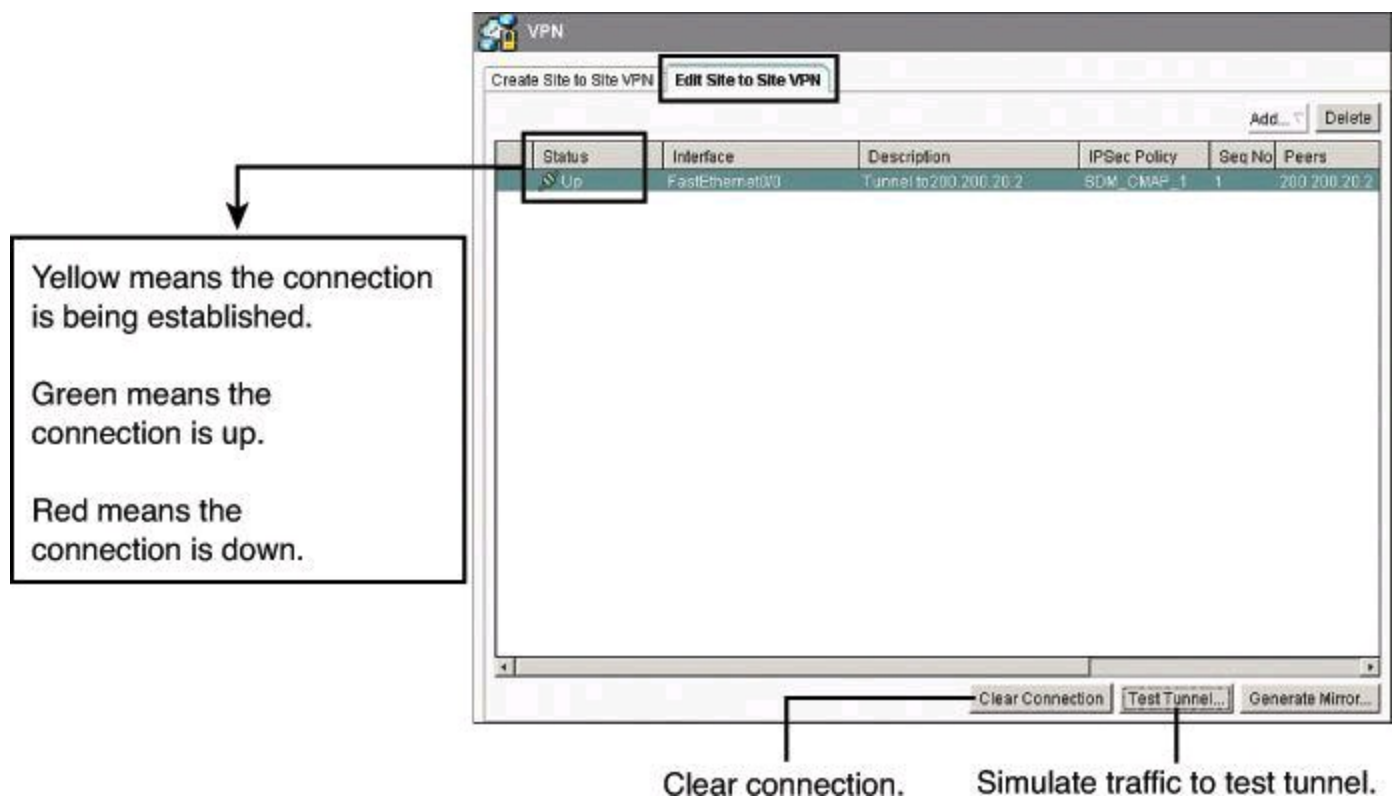


Figure 14-18. Editing VPN Configuration Using CCP

The Status column shows the status of the connection, which is indicated by the intuitive icons. You can also create, edit, and delete VPN connections, and reset existing connections.

The Clear Connection, Generate Mirror, and Test Tunnel buttons are also available from this tab.

Verifying IPsec Configuration Using CLI

The Cisco IOS CLI can also be used to verify the site-to-site VPN configuration. [Table 14-2](#) shows some of the available commands.

Table 14-2. Commands to Verify IPsec Configuration

Parameters	Description
<code>show crypto isakmp policy</code>	Displays configured IKE policies
<code>show crypto ipsec transform-set</code>	Displays configured IPsec transform sets
<code>show crypto map</code>	Displays configured crypto maps
<code>show crypto ipsec sa</code>	Displays established IPsec tunnels

Verifying IKE Policy Using the CLI

Use the `show crypto isakmp policy` command to display configured IKE policies and the default IKE policy settings. This command is useful because it reveals your ISAKMP (IKE) configuration with one command.

The command output, shown in [Example 14-1](#), displays all sections of the IKE policy for router Site-1 according to the topology originally presented in [Figure 14-5](#). The numbers that identify each

section (in this example, section numbers 1 and 2) are used to process different sections for peers with different IKE policies. Sections are processed in order, from lower to higher numbers.

If you do not configure IKE policies, or if traffic does not match any configured policy, a default policy is processed. The default policy includes several sections that start with high section numbers (65507) in order to guarantee their processing as an option of last resort. This default policy can be displayed with the **show crypto isakmp default policy** command.

Example 14-1. Output of *show crypto isakmp policy* Command

[Click here to view code image](#)

```
IOS-FW# show crypto isakmp policy
Global IKE policy
Protection suite of priority 1
  encryption algorithm: Three key triple DES
  hash algorithm: Secure Hash Standard
  authentication method: Pre-Shared Key
  Diffie-Hellman group: #2 (1024 bit)
lifetime: 86400 seconds, no volume limit
```

Verifying IKE Phase 2 Policy Using the CLI

You can use the **show crypto ipsec transform-set** command to show all of the configured transform sets. Because transform sets determine the level of protection that your data will have as it is tunneled, it is important to verify the strength of your IPsec protection policy.

In [Example 14-2](#), the first transform set was created by Cisco Configuration Professional for our example. The remaining transform sets are preconfigured as part of the default policy.

Example 14-2. Output of *show crypto ipsec transform-set* Command

[Click here to view code image](#)

```
IOS-FW# show crypto ipsec transform-set
Transform set ESP-3DES-SHA: { esp-3des esp-sha-hmac }
  will negotiate = { Tunnel, },
Transform set #${default_transform_set_1}: { esp-aes esp-sha-hmac }
  will negotiate = { Transport, },
Transform set #${default_transform_set_0}: { esp-3des esp-sha-hmac }
  will negotiate = { Transport, },
```

Verifying Crypto Maps Using the CLI

To see all of the configured crypto maps, use the **show crypto map** command, as shown in [Example 14-3](#). This command verifies configurations and shows the configured and current peers, as well as the crypto ACL that defines traffic flows to be protected. The output also shows the interface where the crypto map is assigned.

The information provided by the crypto ACL is most important, as it allows you to monitor and

troubleshoot the accuracy and validity of interesting traffic.

The **show running-config** command also reveals many of these same settings.

Note

In production, you could reduce the length of the **show run** command by using the **show run brief** command or the **show run | section crypto** command. The equivalent command on a Cisco ASA firewall would be **show run crypto**.

Example 14-3. Output of *show crypto map* Command

[Click here to view code image](#)

```
IOS-FW# show crypto map
Crypto Map "SDM_CMAP_1" 1 ipsec-isakmp
  Description: Tunnel to200.200.20.2
  Peer = 200.200.20.2
  Extended IP access list 101
  access-list 101 permit ip 10.10.0.0 0.0.255.255 10.20.10.0
  0.0.0.255
  Current peer: 200.200.20.2
  Security association lifetime: 4608000 kilobytes/3600 seconds
  Responder-Only (Y/N): N
  PFS (Y/N): N
  Transform sets={
    ESP-3DES-SHA: { esp-3des esp-sha-hmac },
  }
  Interfaces using crypto map SDM_CMAP_1:
  FastEthernet0/0
```

Monitoring Established IPsec VPN Connections

Cisco Configuration Professional monitoring is streamlined by using the VPN Status window, which displays a tree of VPN connections that are possible on the router, including all types of VPNs (remote access, site-to-site, and others).

Site-to-site tunnels are displayed as IPsec tunnels. This group displays statistics about each IPsec VPN that is configured on the router. Each row in the table represents one IPsec VPN. You will need to click Update to refresh the IPsec tunnel table and display the most current data from the router.

[Figure 14-19](#) showcases the Monitor Tunnel and Test Tunnel options, which are essential to the troubleshooting process.

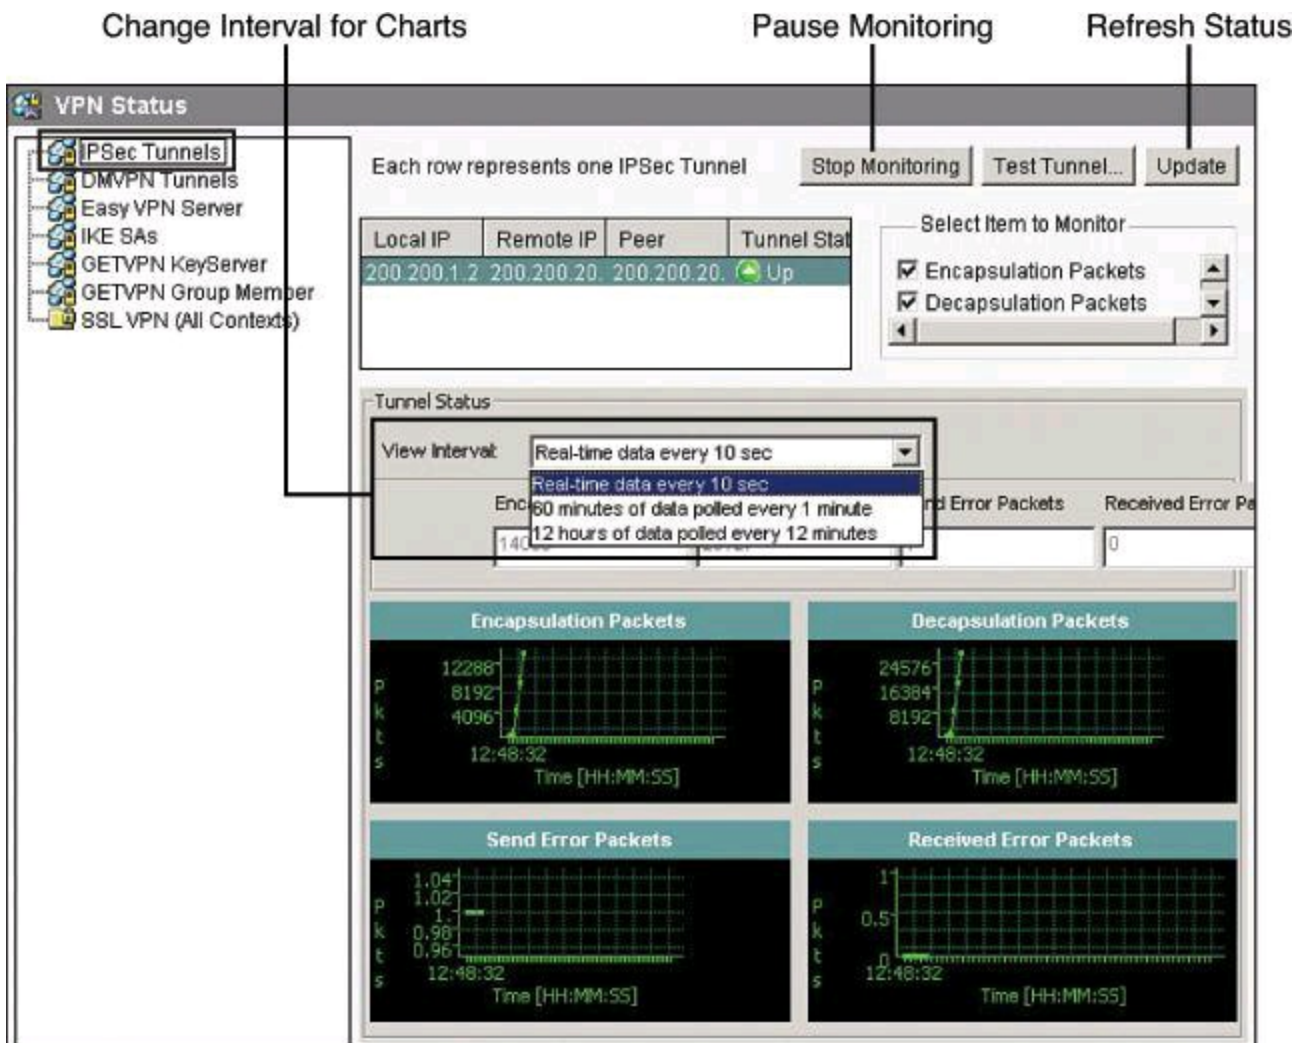


Figure 14-19. Monitoring VPN Status with CCP

IKE Policy Negotiation

The Test Tunnel option allows you to send simulated data through the VPN tunnel. You can click the Start button to start the test. The results are shown on the left side of [Figure 14-20](#).

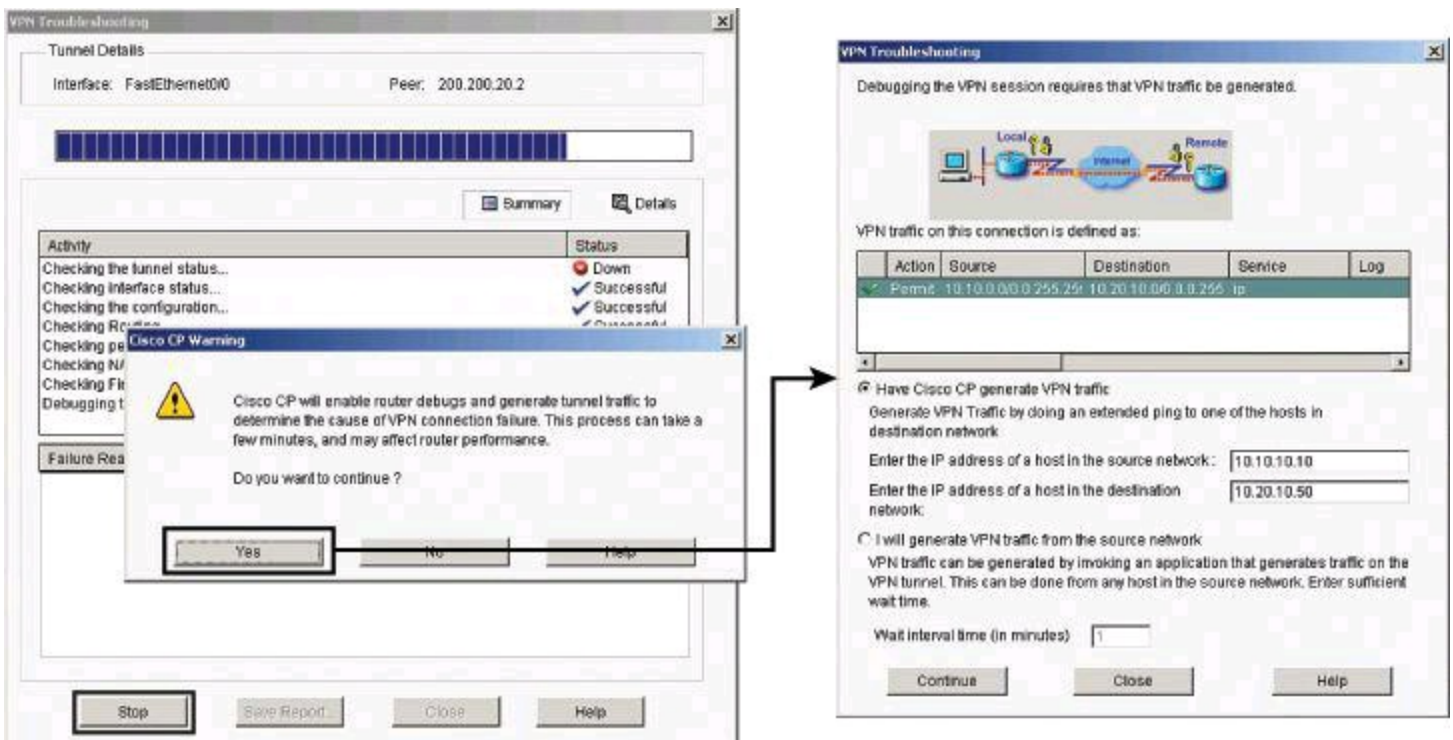
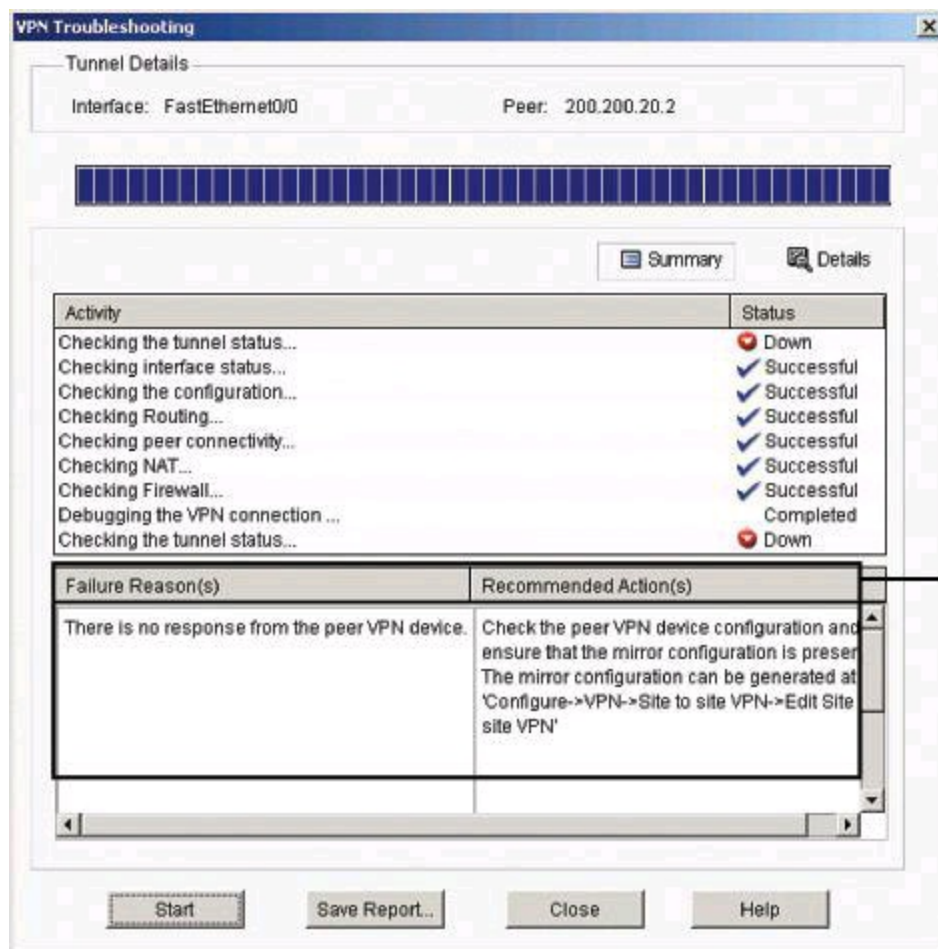


Figure 14-20. Testing Site-to-Site IPsec Tunnel

If the test fails, a popup window prompts you to confirm the initiation of a troubleshooting phase, used by CCP to generate simulated traffic and enable IPsec debugging, in order to automate the troubleshooting process. Clicking Yes to accept this debugging phase presents additional options to either allow CCP to simulate VPN traffic or let you define the profile of the simulated traffic by specifying a destination IP address.

VPN Troubleshooting

Figure 14-21 illustrates the report window of a failed VPN test. The icons shown in Figure 14-20 provide visual cues to determine the severity of the test and to define which of the multiple tests performed by CCP passed or failed. The panel at the bottom suggests the failure reason and spells out recommended actions to resolve the problem.



Failure Reason and Recommended Action

A green arrow means the VPN connection is up.

A red arrow means the VPN connection is down.

A blue checkmark means the VPN test was successful.

A red X means the VPN test failed.

Figure 14-21. VPN Troubleshooting Status Window

Monitoring IKE Security Association

Monitoring and troubleshooting can also be accomplished by using the Cisco IOS CLI. The **show crypto isakmp sa** command shows the status and settings of IKE Phase 1 SAs.

Notice the status of the tunnel under the Status column. Also notice the difference between the Status column and the State column. *State* refers to the progress of the negotiation within Phase 1, while *status* refers to VPN high availability (HA), showing a state of ACTIVE for SAs on the HA

active router, and STANDBY for the SAs on the standby router. This means that a status of ACTIVE is not related to active tunnels; in fact, the Status column is not an indication of health of each Phase 1 SA. A state of QM_IDLE is considered normal for an established Phase 1 tunnel.

[Table 14-3](#) and [14-4](#) list and describe the possible values for state and status.

Table 14-3. Values for IKE Security Association State

State	Description
MM_NO_STATE	The Phase 1 SA has been created, but nothing else has happened yet.
MM_SA_SETUP	The peers have agreed on parameters for the Phase 1 SA.
MM_KEY_EXCH	DH negotiation was successful, but the Phase 1 SA remains unauthenticated.
MM_KEY_AUTH	The Phase 1 SA has been authenticated.
QM_IDLE	The Phase 1 SA is idle, in quiescent state.

Table 14-4. Values for IKE Security Association Status

Status	Description
ACTIVE	The Phase 1 SA is terminated on the active device of an HA cluster.
STANDBY	The Phase 1 SA is terminated on the standby device of an HA cluster.

Monitoring IPsec Security Association

IPsec Phase 2 can also be monitored by using the **show crypto ipsec sa** command, as shown in [Example 14-4](#). Increasing encrypt and decrypt counters are a general indication of an operational IPsec Phase 2 tunnel.

Example 14-4. Output of the *show crypto ipsec sa* Command

[Click here to view code image](#)

```
RouterA# show crypto ipsec sa
interface: FastEthernet0/0
  Crypto map tag: SDM_CMAP_1, local addr 200.200.1.2

  protected vrf: (none)
  local ident (addr/mask/prot/port): (10.10.0.0/255.255.0.0/0/0)
  remote ident (addr/mask/prot/port): (10.20.10.0/255.255.255.0/0/0)
  current_peer 200.200.20.2 port 500
    PERMIT, flags={origin is acl,}
    #pkts encaps: 4030, #pkts encrypt: 4030, #pkts digest: 4030
    #pkts decaps: 4033, #pkts decrypt: 4033, #pkts verify: 4033
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0
```

Note

The **show crypto ipsec sa** command shows different information depending on whether it is issued from a Cisco IOS router or an ASA firewall:

- **Cisco IOS router:** IPsec SAs show the local and remote identities even if the tunnel is not up. The key item to look for is the presence of inbound and outbound Security Profile Indexes (SPI). Increasing encapsulation/decapsulation counters are also a way to confirm that actual traffic is going through the tunnel.
 - **ASA devices:** IPsec SAs do not show identities unless the tunnel is up. Also, the presence of encapsulated packets but no decapsulated packets (or vice versa) is usually a sign of bad routes on one end of the tunnel and/or an ACL in the middle of the flow that is blocking ESP or UDP 4500 (NAT-T) traffic.
-

Summary

The key points covered in this chapter are as follows:

- Planning an IPsec site-to-site VPN includes selecting the cipher suite and building monitoring and contingency planning into the IPsec site-to-site VPN configuration checklist.
- Cisco Configuration Professional provides a step-by-step wizard to guide configuration.
- Use the Monitor Tunnel options to verify configuration.
- Use the Test VPN option and CLI commands to test and troubleshoot.

References

For additional information, refer to these resources:

Carmouche, J. H. *IPsec Virtual Private Network Fundamentals* (Cisco Press, 2007).

Cisco Systems, Inc. “Cisco IOS IPsec,”
http://www.cisco.com/en/US/products/ps6635/products_ios_protocol_group_home.html.

Deal, R. *The Complete Cisco VPN Configuration Guide* (Cisco Press, 2005).

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, “[Answers to Chapter Review Questions](#).”

1. Which command is best suited to show that traffic is flowing through the VPN tunnel?
 - a. **show crypto map**
 - b. **show crypto ipsec transform-set**
 - c. **show crypto isakmp policy**
 - d. **show crypto ipsec sa**

- 2.** Which command would help troubleshoot a Phase I issue?
- debug crypto sa**
 - debug crypto phase 1**
 - debug crypto isakmp**
 - show crypto sa**
- 3.** Which command would be the most helpful to compare the configuration of peer, ACL, SA lifetime, and transform sets?
- debug crypto sa**
 - show crypto isakmp sa**
 - show crypto map sa**
 - show crypto map**
- 4.** You are deploying IPsec VPNs with the Cisco Configuration Professional wizard. Which of the following is true about traffic that is denied by the crypto ACL?
- The traffic will be denied access to the VPN.
 - The traffic will be dropped if it does not match the crypto policy.
 - The traffic is permitted if it matches the crypto policy.
 - The traffic is permitted in cleartext.
- 5.** Which of the following options is available in the VPN Status window of Cisco Configuration Professional?
- Create Mirror
 - Monitor Tunnel
 - Clear Connection
 - Add Policy
- 6.** Which of the following is considered a stable state for a Phase 1 security association in IPsec VPNs, using the **show crypto isakmp sa** command?
- QM_IDLE
 - QM_NO_STATE
 - QM_ACTIVE
 - QM_ESTABLISHED
- 7.** You issue the command **show crypto isakmp sa**. You discover that the IKE tunnel is down. Issuing the **show crypto ipsec sa** command will be helpful at troubleshooting the connectivity problem. True or false?
- True
 - False
- 8.** You use the Test Tunnel option of CCP. The test fails. What, if any, happens next?
- CCP returns a fail message and the Test Tunnel option closes.
 - CCP pings automatically the IP address of the peer to ensure that it is live.

- c. CCP presents a popup window prompting you to confirm the initiation of a troubleshooting phase.
- d. CCP opens a CLI window and issues **show crypto** commands to the router.

Chapter 15. SSL VPNs with Cisco ASA

Upon completing this chapter, you will be able to describe the Secure Sockets Layer (SSL) VPN operational framework. You will be able to demonstrate the deployment of basic SSL VPNs on Cisco Adaptive Security Appliances (ASA) using Cisco Adaptive Security Device Manager (ASDM) and the Cisco AnyConnect VPN Client. This chapter prepares you to meet these objectives:

- Describe the use cases and operational requirements of Cisco SSL VPNs
- Describe the protocol framework for SSL and TLS
- Describe a configuration that is based on SSL VPN deployment options and other design considerations
- Describe the steps to configure Cisco VPN clientless mode on Cisco ASA and demonstrate the configuration on Cisco ASDM
- Describe the steps to configure Cisco full-tunnel mode on Cisco ASA and demonstrate the configuration on Cisco ASDM using the Cisco AnyConnect VPN Client

Mobility and IT consumer market trends influence the need for comprehensive remote-access security policies. SSL VPNs are commonly used as a remote-access service. As such, SSL VPNs must integrate strong cryptography and standards-based components with deployment and operational efficiencies and endpoint security. This chapter describes the SSL protocol framework and describes the benefits of the Cisco SSL VPN solution. The chapter also demonstrates the configuration of clientless and full-tunnel SSL VPNs using Cisco ASDM and the Cisco AnyConnect VPN Client.

SSL VPNs in Borderless Networks

Remote-access and mobility services have gone through drastic changes in the past few years. There are three market transitions driving the network architectures of the future:

- **Mobility:** Over the next three years, there will be more than 1.3 billion new networked mobile devices. These devices will include mobile Internet devices, notebooks, smartphones, notebooks, tablets, and other devices.
- **Video:** Video has truly become pervasive in both consumer and business environments. No longer simply a means of communication, it has become a key factor in reducing training and travel costs. It is also beginning to reshape business models, especially in retail.
- **IT consumerization:** This is essentially what can be called the new workplace experience, in which consumer demand and new consumer devices influence technology decisions made by IT departments. Globalization has created a need for partners, customers, and employees to connect to each other and to the Internet over traditional boundaries from a variety of environments and devices.

These three market transitions are changing the IT environment and the strategic role that the network plays in delivery of new services. These market factors will influence the features that are found in remote-access products and technologies.

These transitions are also placing significant demands on the typical functions of the IT organization. Functional needs have remained similar: scalability, availability, performance, security,

manageability, and cost of ownership. But as new (and potentially untrusted) devices connect to the network from multiple (and potentially untrusted) places, and as the IT organization moves beyond the infrastructure it can directly control, these parameters have become much more complex. In fact, the parameters are no longer linear; they are multidimensional, and the IT organizations, therefore, are struggling to meet their requirements in terms of the underlying network infrastructure.

The main challenges to IT organizations in providing remote and mobile access to corporate resources include the following:

- They need to provide an improved access experience, providing everywhere connectivity and management whether through wired, wireless, or remote access.
- To support the increasing number of mobile workers, corporate security administrators must provide context-aware security and policy enforcement. They must extend the application experience, regardless of the locations of the end users, what devices they are using, what network they are on, and where the information they are accessing is located. Administrators must also be able to support a heterogeneous set of laptops and mobile devices to encourage choice for their clients—the end users. Finally, they must provide this security unobtrusively and agnostic of the access network to minimize end-user concerns.
- The days of the standard corporate device are over. Businesses must meet the first two challenges in a way that will allow them to embrace the wave of new devices—smartphones and tablets, in addition to the traditional laptops—that workers are using.

Cisco SSL VPN

The Cisco SSL VPN technology provides remote-access connectivity from almost any Internet-enabled location with a web browser and its native SSL encryption. Cisco SSL VPN provides the flexibility to support secure access for all users, regardless of the endpoint host from which they establish a connection. If application access requirements are modest, SSL VPN does not require a software client to be preinstalled on the endpoint host. This ability enables companies to extend their secure enterprise networks to any authorized user by providing remote-access connectivity to corporate resources from any Internet-enabled location.

Cisco SSL VPN currently delivers three modes of Cisco SSL VPN access: clientless, thin client, and full client, as shown in [Figure 12-2](#), earlier in this book. Cisco SSL VPNs allow users to access web pages and services, including the ability to access files, send and receive email, and run TCP-based applications without IPsec VPN client software. Cisco SSL VPNs are appropriate for user populations that require per-application or per-server access control, or access from desktops not owned by the enterprise.

In many cases, IPsec and Cisco SSL VPNs are complementary because they solve different problems. This complementary approach allows a single device to address all remote-access user requirements.

The primary benefit of Cisco SSL VPN is its ubiquity and versatility. SSL is available on computers, tablets, and smartphones. You use it when you shop or bank online.

An additional benefit of Cisco SSL VPN is that it is compatible with Dynamic Multipoint VPNs (DMVPN), Cisco IOS Firewalls, IPsec, intrusion prevention systems (IPS), Cisco Easy VPN, and Network Address Translation (NAT).

SSL VPNs and IPsec VPNs are complementary technologies that you can deploy together to better address the unique access requirements of diverse user communities. Both offer access to virtually any network application or resource. Cisco SSL VPNs offer additional features such as easy connectivity from desktops outside your company’s management, little or no desktop software maintenance, and user-customized web portals upon login.

[Table 15-1](#) shows the most significant differences between IPsec VPNs and Clientless SSL VPNs. IPsec excels in the number of applications that are supported, the strength of its encryption, the strength of its authentication, and its overall security. The areas in which Clientless SSL VPN excels are its ease of use and ease of deployment.

Table 15-1. Clientless SSL VPN Versus IPsec VPN

	Clientless SSL VPN	IPsec VPN
Applications	Web-enabled applications, file sharing applications, email.	All IP-based applications.
Encryption	Moderate. Key lengths range from 40 bits to 128 bits. (Note: If your browser supports AES 256, then the encryption strength is the same as IPsec.)	Stronger. Key lengths range from 56 bits to 256 bits.
Authentication	Moderate. One-way or two-way authentication.	Strong. Two-way authentication using shared secrets or digital certificates.
Ease of use	Very high.	Moderate. Can be challenging to nontechnical users.
Overall security	Moderate. Any device can connect.	Strong. Only specific devices with specific configurations can connect.
Summary	Anywhere access	Any application

When used for VPN connectivity, SSL does not require any special-purpose client software to be preinstalled on the system. Cisco SSL VPNs are also known as “clientless VPNs” because they rely on the existence of web browsers on the connecting host, a pre-installed client that already supports the protocol. This feature makes SSL VPNs capable of “anywhere” connectivity from company-managed desktops and non-company-managed desktops, such as employee-owned PCs, contractor or business partner desktops, and Internet kiosks. Any software that is required for application access across the SSL VPN connection is dynamically downloaded on an as-needed basis, minimizing desktop software maintenance.

Traditionally, when security was a primary issue, IPsec was the preferred choice. If support and ease of deployment were the primary issues, then SSL was considered. Nowadays, SSL VPN is usually considered the preferred choice, especially because it provides security comparable to IPsec. Cisco VPN products support both technologies in order to satisfy the largest number of customer

requirements.

SSL and TLS Protocol Framework

Transport Layer Security (TLS) and its predecessor, SSL, are cryptographic protocols that provide secure communications on the Internet for such things as web browsing, email, Internet faxing, instant messaging, and other data transfers. Originally developed by Netscape, SSL has been universally accepted on the World Wide Web to transfer information privately over the Internet. Virtually all online transactions and browser-based monetary interactions that occur on the Internet are secured by SSL.

As mentioned earlier, SSL has evolved from Netscape's original implementation. TLS is a standards-based alternative to SSL.

SSL and TLS provide confidentiality, integrity, and authentication services to the applications that use them. In that sense, these protocols encrypt and authenticate from the session layer up, including the presentation and application layers.

SSL is used to encrypt and authenticate the session layer and above. As such, it encrypts more than just HTTP (called HTTPS); it can also encrypt FTP (thus FTPS), POP (for POPs), LDAP (for LDAPS), wireless security (EAP-TLS), and others.

Cryptographically, SSL and TLS rely on public key infrastructure (PKI) and digital certificates for authenticating the VPN endpoints. In e-commerce environments, typically only the server side is authenticated and requires a digital certificate. Other applications are increasingly using mutual authentication, where both the client and the server use digital certificates to identify themselves.

[Figure 15-1](#) shows the encapsulation of SSL/TLS.

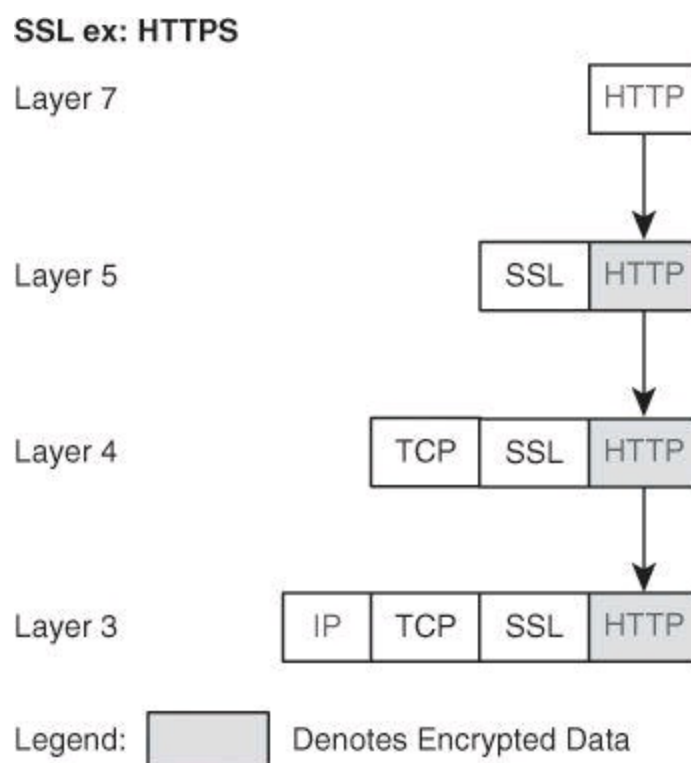


Figure 15-1. SSL/TLS Encapsulation

There are slight differences between SSL and TLS, but the protocols remain largely similar. The terms are sometimes used interchangeably, but interestingly, the protocols are not interoperable. Some implementations include provisions to switch to the other protocol if necessary for a given session, especially in the case of TLS, which incorporates mechanisms to switch to SSL if the client connection requires it.

TLS 1.0 originated from SSL 3.0, and later versions of both protocols show a similar development path. [Table 15-2](#) lists some of the differences.

Table 15-2. SSL and TLS

SSL	TLS
Developed by Netscape in the 1990s	Standard developed by IETF
Starts with security and proceeds directly to secured communications	Can start with an insecure “hello” and switch to secured communications only after the handshake is successful
Wider and more prevalent support on client-side applications	Granular and more pervasive support on server-side applications
More weaknesses identified in older SSL versions	Stronger implementation due to standards process and universal support

DTLS

TLS has a sister protocol called DTLS, for Datagram Transport Layer Security. The main difference between TLS and DTLS is that DTLS gets encapsulated in UDP at Layer 4, while TLS gets encapsulated in TCP, as shown in [Figure 15-1](#).

SSL Cryptography

The SSL and TLS protocols support the use of a variety of different cryptographic algorithms, or ciphers, for use in operations such as authenticating the server and client to each other, transmitting certificates, and establishing session keys. Symmetric algorithms are used for bulk encryption, asymmetric algorithms are used for authentication and the exchange of keys, and hashing is used as part of the authentication process. [Figure 15-2](#) illustrates some of the supported algorithms (in the figure, EC stands for Elliptic Curve).

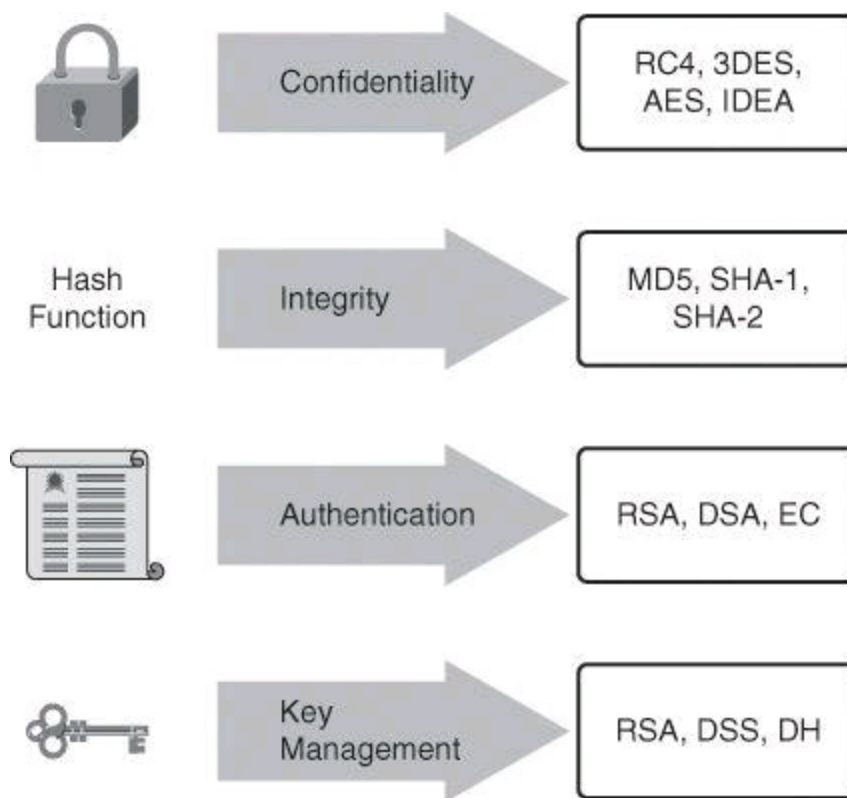


Figure 15-2. Supported SSL Cryptographic Cipher Suites

Note

The listed encryption algorithms are just examples for SSL and TLS support. Multiple other cipher suites are supported but not listed.

SSL Tunnel Establishment

[Figure 15-3](#) gives a simplified explanation of the key steps in establishing an SSL session:

Step 1. The user makes an outbound connection to TCP port 443.

Step 2. The security appliance responds with a digital certificate, which contains a public key that is digitally signed by a trusted certificate authority (CA). The ASA can also present a self-signed certificate. The client will have to accept the security warning mentioning that the validity of the signature cannot be confirmed, but the self-signed certificate is useable.

Step 3. The user computer generates a shared-secret, symmetric key that both parties will use.

Step 4. The shared secret is encrypted with the public key of the security appliance and transmitted to the router. The security appliance software is able to easily decrypt the packet using its private key. Now both participants in the session know the shared-secret key.

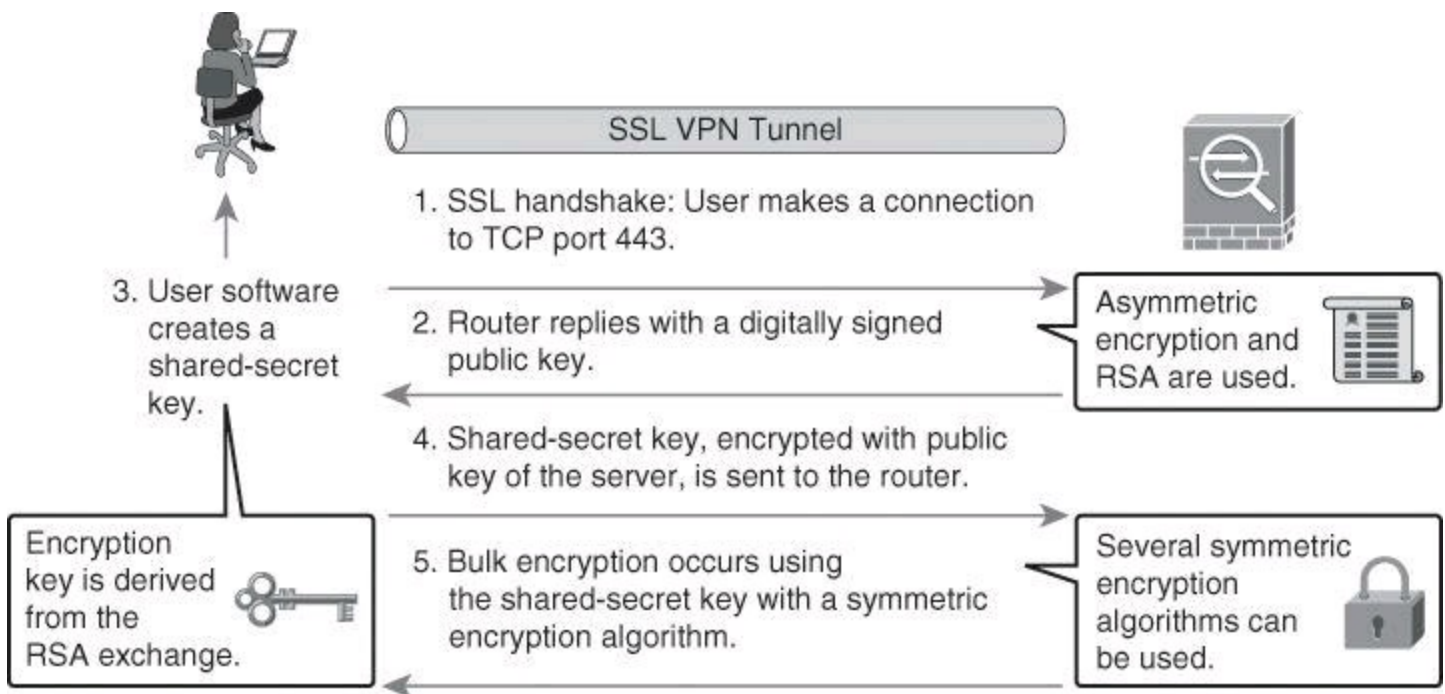


Figure 15-3. Major Steps of SSL VPN Tunnel Establishment

SSL utilizes encryption algorithms with key lengths that range from 40 to 256 bits.

SSL Tunnel Establishment Example

SSL and PKI, as explained in [Chapter 12](#), are complex topics. [Figure 15-4](#) and the following steps explain the essence of how an HTTPS session is established. This example is not meant to provide you with a detailed and accurate account of how an HTTPS session is set up, but rather to provide an understandable example without all the technicalities.

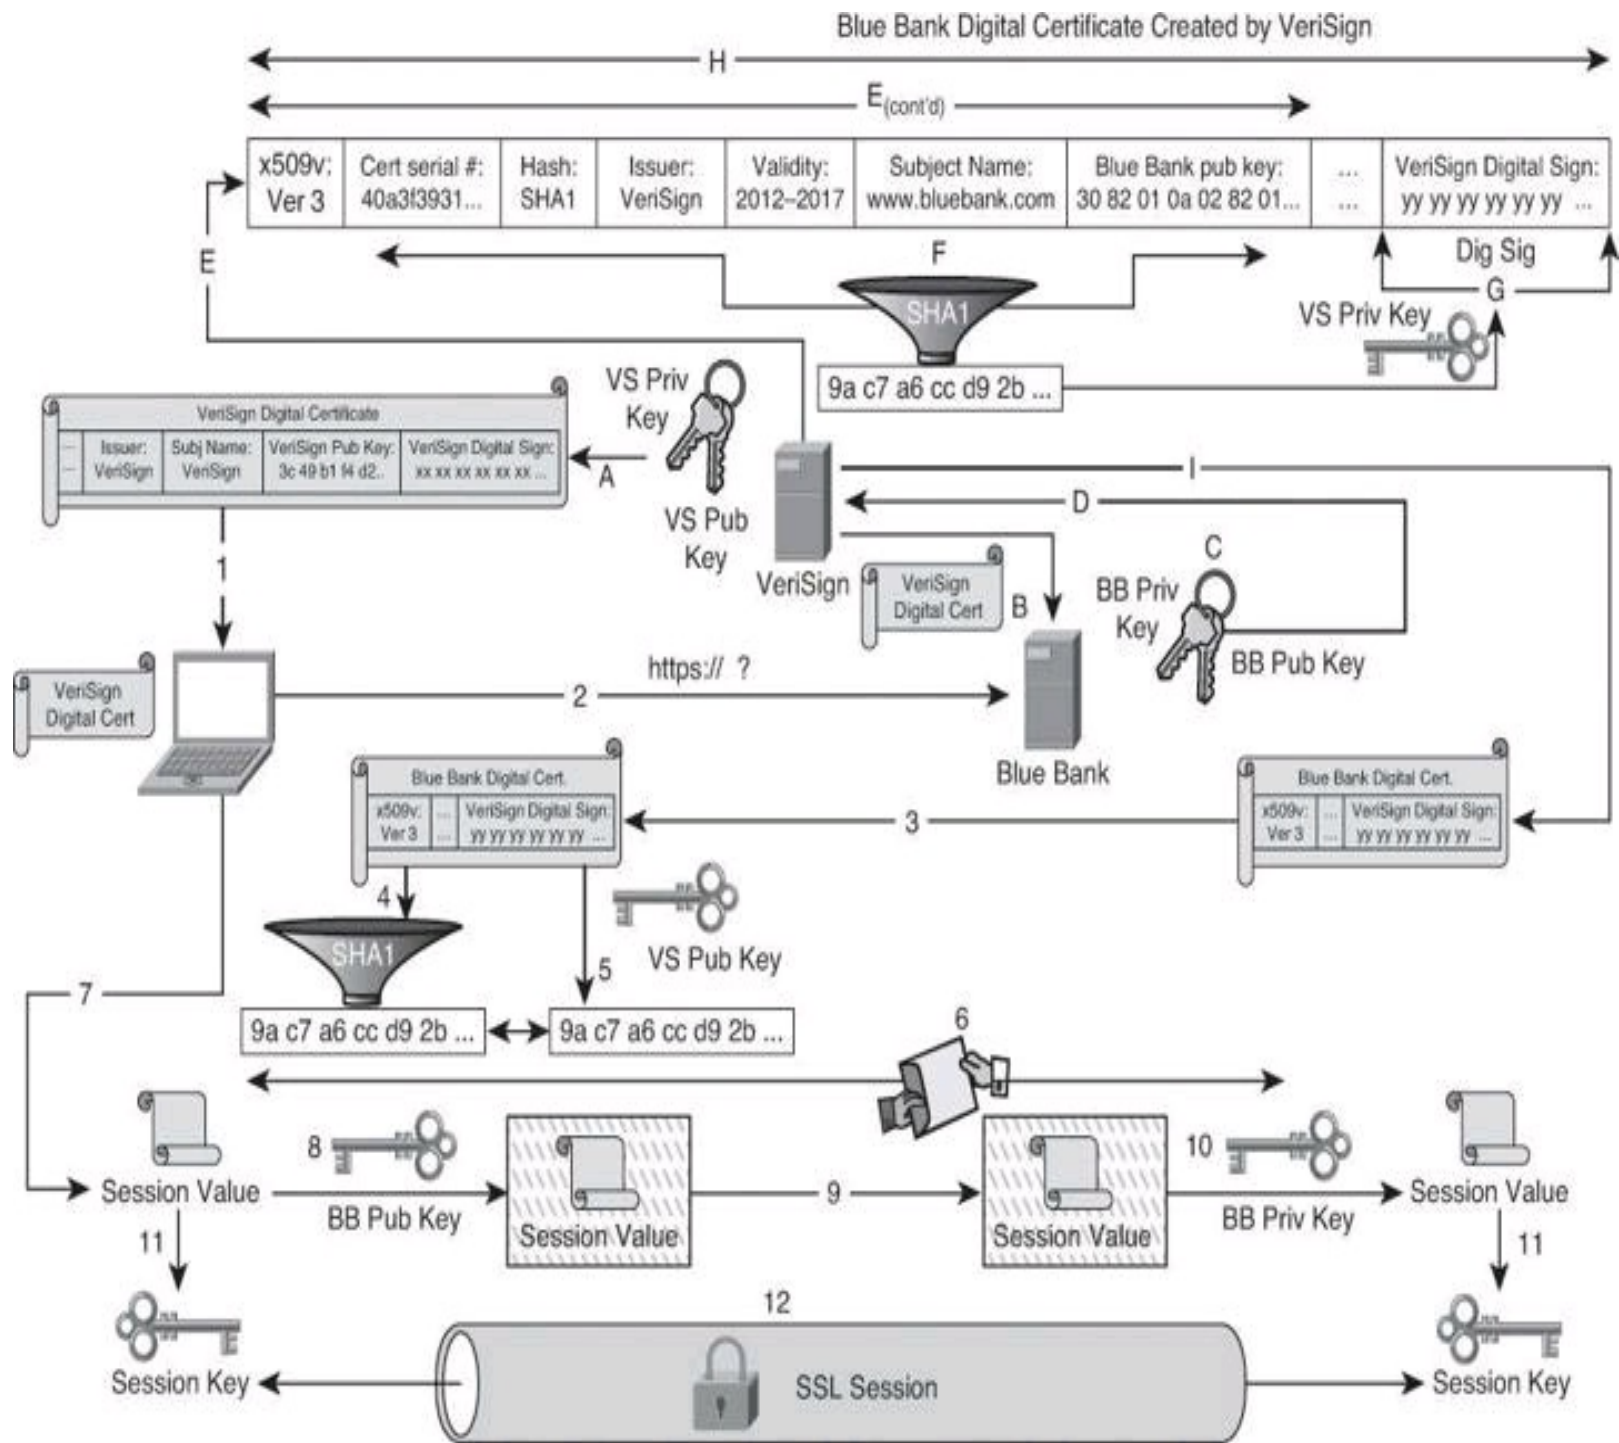


Figure 15-4. Example of an HTTPS Session

Steps A to I illustrate steps between the Blue Bank server and VeriSign.

Steps 1 to 11 illustrate steps between the HTTPS client and the Blue Bank server.

Step A. As a root CA, the VeriSign server has created a key pair. It used its own private key to sign its own public key—this is called a self-signed certificate. (Signing of certificates will be covered in Step E.) By the way, serious organizations such as VeriSign do not use their top root certificate to sign client certificates. (The hierarchy of CAs is beyond the scope of this book.)

Step B. Blue Bank wants to start offering an online banking service. It selects VeriSign as the public CA server (CA theory was covered in [Chapter 12](#)). The Blue Bank server acquires the digital certificate of VeriSign. (The creation of digital certificates will be

covered in Step D.) It is highly probable that the Blue Bank server would already have installed on it the VeriSign digital certificate. Popular desktops and server OSs, such as those from Microsoft, already have this CA certificate installed in the certificate store. One of the easiest ways to see the content of the certificate store, and thus to see the trusted root certificates installed on a server or a workstation, is to start Internet Explorer and navigate to **Tools > Internet Options**, click the **Content** tab, click the **Certificates** button, and then click the **Trusted Root Certification Authorities** tab.

Step C. Blue Bank generates a private and public key pair.

Step D. Blue Bank petitions VeriSign to validate its public key in order to generate a digital certificate.

Step E. VeriSign proceeds to turn the Blue Bank public key into a digital certificate by adding information to the public key. Some of the fields added to the public key, as shown in [Figure 15-4](#), are Subject Name (it could be Blue Bank URL), Validity (dates), and so forth.

Step F. VeriSign hashes some of the fields included with the information of the Blue Bank digital certificate and produces a digest.

Step G. VeriSign encrypts the digest and encrypts it with its private key. We now have a digital signature; so the definition of a digital signature is: *a hash encrypted with a private key*.

Step H. The Blue Bank digital certificate has been created. What is the Blue Bank digital certificate? It's a validated public key signed by the CA.

What is a validated public key? It's the public key, with fields added to it.

What is a signature? It's the hash result encrypted with the CA private key.

Step I. Blue Bank installs its digital certificate, which is Blue Bank's public key, signed by the CA.

Now, as a Blue Bank customer, you wish to open an online banking session, using your laptop.

Step 1. Your laptop is equipped with the root certificate of the CA used by Blue Bank, VeriSign. If the server on which you want to open an SSL session had its digital certificate issued by a CA that you are unfamiliar with, you would need to download the root certificate onto your computer. In our scenario, the laptop is running a popular OS and already has a copy of the VeriSign digital certificate in its certificate store.

Step 2. Your laptop browser queries the Blue Bank server to see if it accepts HTTPS sessions.

Step 3. Because Blue Bank accepts SSL sessions, the server sends a copy of its digital certificate to the laptop. Actually, sending a copy of the Blue Bank digital certificate happens only the first time that the particular laptop connects to the Blue Bank server. After that, the server and the client will only check whether they are using the right Digital Certificate Serial number (it's one of the fields in Step E). When Blue Bank puts a new certificate in production, either because the previous one has expired or because the key has been compromised, it will ask the laptop for a certificate number that the laptop hasn't heard about yet. In that case, the Blue Bank server will then push the new certificate to the

laptop.

Step 4. The client proceeds to validate the certificate. It checks for the validity date and the name of the issuer. Because the client already knows about the issuer of the certificate, in this case VeriSign, it proceeds to validate the signature. To do so, it first hashes the pre-agreed fields from the digital certificate to get a digest.

Step 5. Continuing with the validation of the Blue Bank digital certificate, the client decrypts the signature originally created by VeriSign's private key (Step G). To decrypt that signature, the client uses the VeriSign public key acquired in Step 1. If the hash digest calculated by the client in Step 4 matches the result found by decrypting the digital signature in Step 5, then the client knows two things for sure: the issuer (signing authority) was indeed VeriSign, and no one tampered with the content of the certificate. The authenticity of the certificate is proven by the fact that the client was successful at decrypting the signature using VeriSign's public key, which means that the signature was created in the first place by VeriSign's private key. And who has VeriSign's private key? Only VeriSign—and it's guarded ferociously.

Step 6. The ultimate goal of the client is to have an HTTPS session with the server, as shown in Step 12. This HTTPS session will use symmetric cryptography because it is much faster than asymmetric cryptography. The client and the server must at this stage agree on the encryption that they will use for the symmetric encrypted session in Step 12. Let's say for our scenario that the client and the server agree that the encryption algorithm will be RC4-128 (Rivest Cipher, 128-bit) and exchange other essential information.

Step 7. As part of the process of coming up with a symmetric encryption key that will be used in Step 12, the client generates a value (called the pre-master value, also referred as a session key in [Figure 15-4](#)) that will be used to finalize the symmetric key used by the server and client.

Step 8. The pre-master value (session value) needs to be shared confidentially with the server, so the client encrypts it with the server's public key, found with the Blue Bank digital certificate received in Step 3.

Step 9. The encrypted pre-master value (session value) is sent to the server.

Step 10. Using its private key, the Blue Bank server can decrypt the data it just received from the client and discover the pre-master value (session value) that is to be used in the fabrication of a symmetric encryption key.

Step 11. Both the client and the server now have a copy of the pre-master value (session value) and all the other values necessary to generate the same session key. They independently arrive at the same result, and that result is the symmetric key they will use for their SSL session.

Step 12. The SSL tunnel comes up. Only then will some browsers show you a padlock near the URL bar, and only then will Blue Bank present you with the login page where you will be required to provide your banking card number and password.

When the client gracefully logs out of the session or when the session times out, the client and the server both forget the session key that was generated in Step 11.

Cisco SSL VPN Deployment Options and Considerations

As shown in [Chapter 12](#) (in [Figure 12-2](#)), Cisco SSL VPNs offer different types of access options: SSL Clientless and SSL full client mode, referred to as Cisco AnyConnect SSL VPN. A third option exists as SSL Clientless thin-client, which provides some more functionality with the SSL Clientless. [Figure 15-5](#) shows the use case for the different ways to access the corporate network remotely.

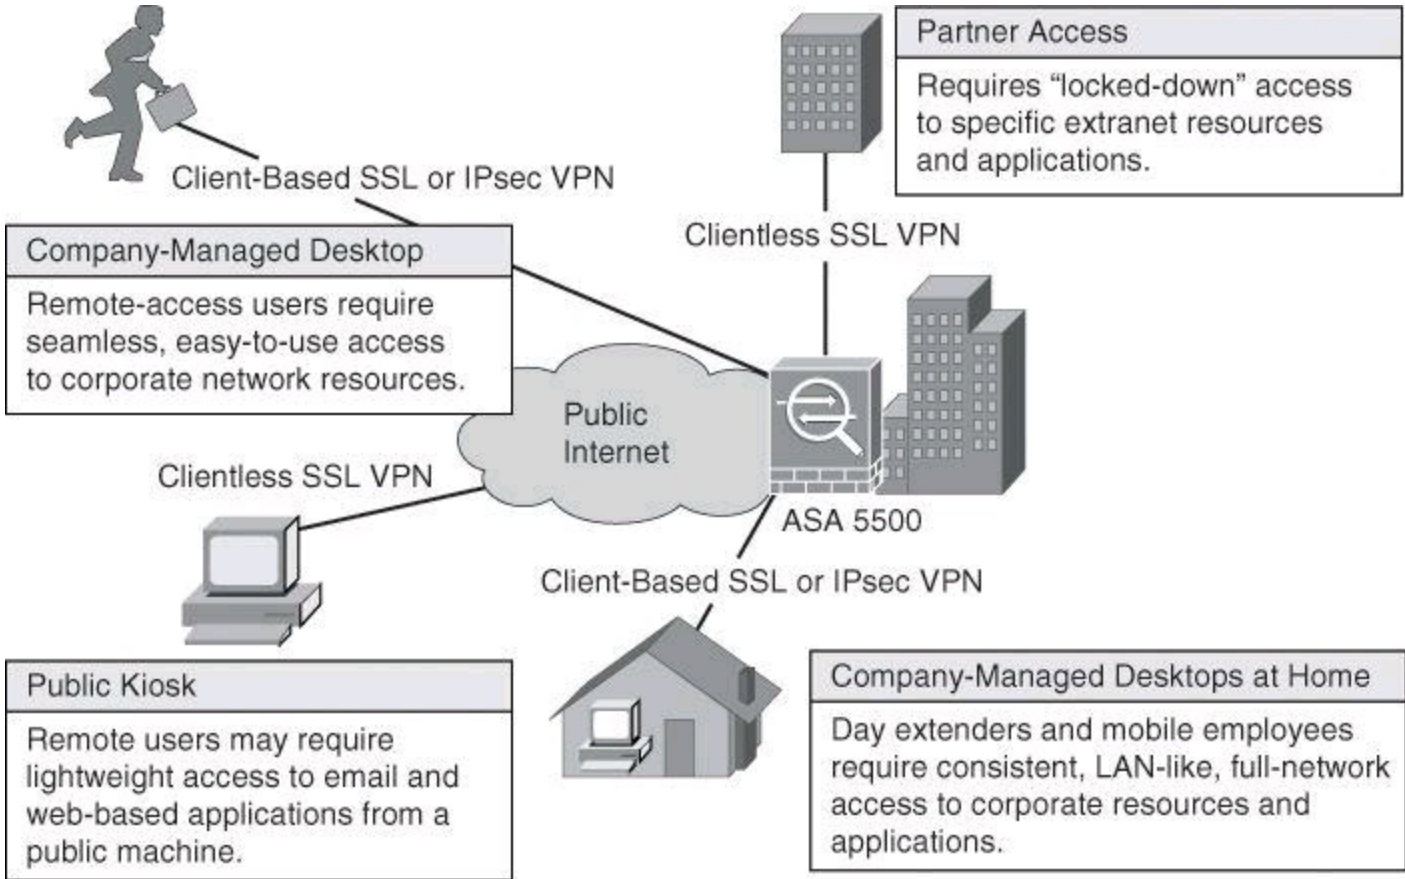


Figure 15-5. SSL VPN Access Scenarios

[Figure 15-5](#) shows the usage scenario for the following three types of SSL access:

- **Clientless:** This option enables a user to connect to the network with few requirements beyond a basic web browser. The access provides the ability to reach web servers or webified resources such as file shares.
- **Thin client:** A small applet or application (generally under 100 KB) is downloaded that provides access to a subset of application resources, generally TCP, and often an outbound and static port. This thin-client is part of the Clientless option.
- **SSL VPN client:** A larger client is delivered to the end user and is required when users need full application access. The applications that can be accessed are very similar to those available via IPsec VPN. This client is delivered via a web page (the device that the user is connecting to) and never needs to be manually distributed or installed. This client is referred to as Cisco AnyConnect.

The architecture includes VPN termination points (Cisco IOS routers or ASA) and the Cisco AnyConnect VPN Client.

[Figure 15-6](#) illustrates the differences between the two main deployment modes that are found in Cisco SSL VPN solutions:

- Clientless mode is easier to deploy but presents additional challenges in terms of

endpoint security and application access.

- Full network access supports a wider variety of applications but presents additional operational challenges in downloading and maintaining the client software on remote hosts.

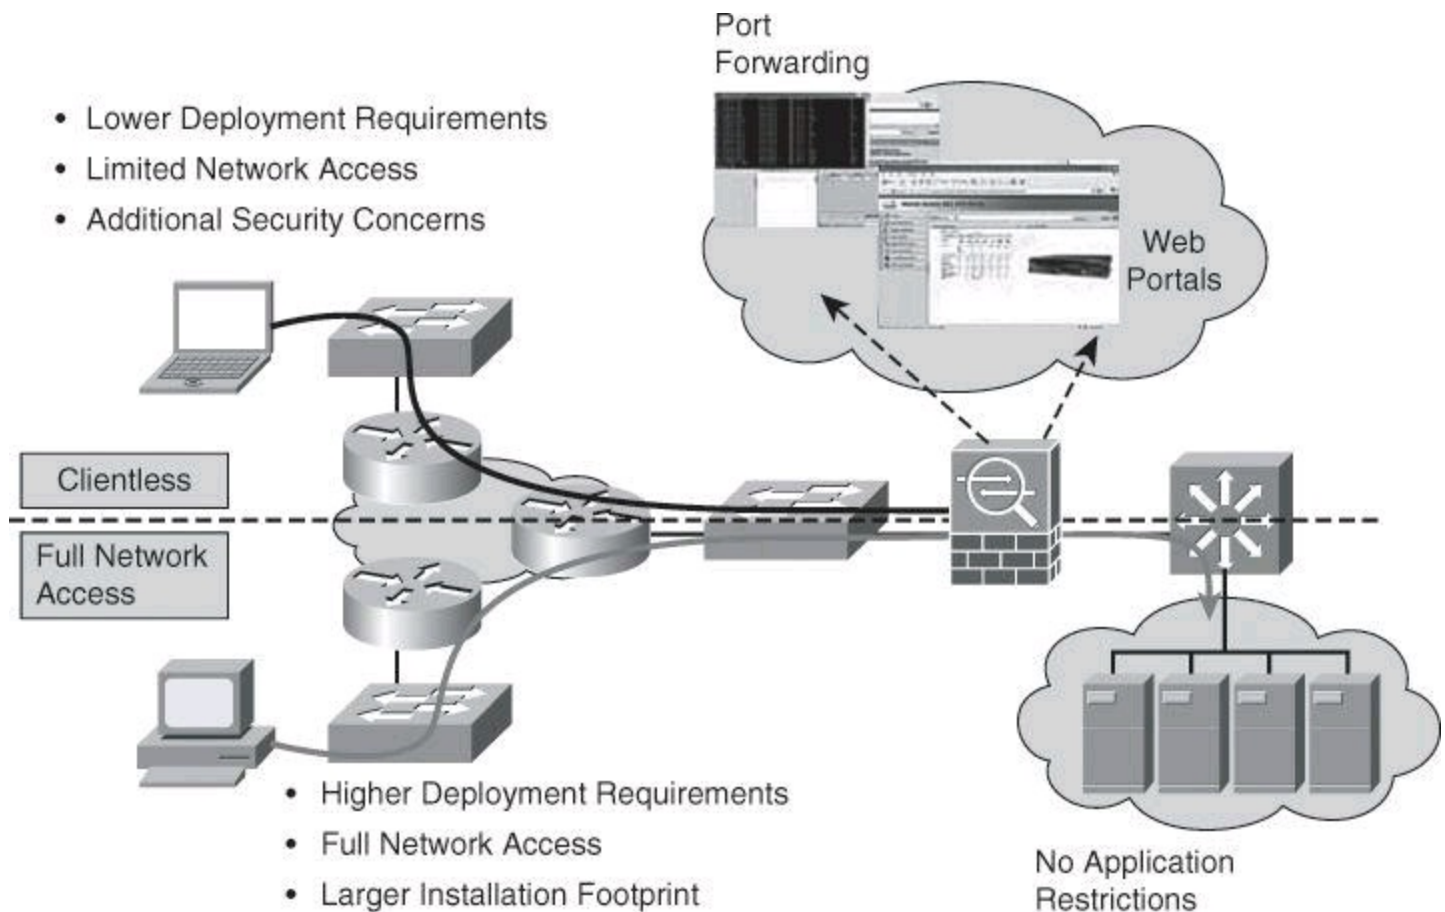


Figure 15-6. Two Main SSL Deployment Modes

SSL VPNs provide a flexible way of extending network resources and applications to remote users. When using a clientless SSL VPN deployment model, corporations have the additional flexibility of providing access to corporate resources even when the remote device is not corporately managed. In this deployment model, the Cisco ASA is used as a proxy device to network resources and provides a web portal interface for remote devices to navigate the network. Additionally, to access the SSL VPN network, the remote device system only requires a supported web browser with built-in SSL functionality. Although more flexible than client-based SSL VPNs, clientless SSL VPNs provide only limited network application or resource access and include additional security risks when using non-corporate-managed clients.

Full network access SSL VPNs operate much like standard IPsec VPN clients in the way that they provide network access. In comparison to the clientless SSL VPN deployment model, there are no network access restrictions for resources or applications. Full network access SSL VPNs require more planning for network deployment because they require a client to be installed on the remote systems. This requirement makes it difficult to deploy on non-corporate-managed systems because most SSL VPN clients require administrator privileges to install. The use of corporate-managed systems provides tighter control over endpoint security.

Cisco SSL VPN Client: Full Network Access

With an SSL VPN client (Layer 3), an SSL-VPN tunneling client is needed because users need full

application access. Users need to use applications such as Microsoft Outlook, Cisco IP Communicator, Lotus Notes, Lotus Sametime, PeopleCube Meeting Maker, Telnet, Secure Shell (SSH), X Window System, and so on, but without having to install and administer an IPsec VPN software package. Cisco SSL VPN clientless and port-forwarding modes are not capable of this level of functionality.

The Cisco AnyConnect VPN client, using secure SSL connections to the security appliance, provides remote users with full VPN tunneling to corporate resources. This client is delivered via a web page (hosted on the Cisco ASA to which the user is connecting) and never needs to be manually distributed or installed.

The following are among the many features of the Cisco AnyConnect VPN client:

- **Optimal gateway selection:** Determines and establishes connectivity to the most optimal network access point, eliminating the need for end users to determine the nearest location.
- **Mobility-friendly:** Designed for mobile users, it can be configured so that the VPN connection remains established during IP address changes, loss of connectivity, hibernation, or standby.
- **Broad operating system support:**
 - Windows 7 32-bit (x86) and 64-bit (x64)
 - Windows Vista 32-bit (x86) and 64-bit (x64), including Service Packs 1 and 2 (SP1/SP2)
 - Windows XP SP2+ 32-bit (x86) and 64-bit (x64)
 - Mac OS X 10.5 and 10.6.x
 - Linux Intel (2.6.x kernel)
 - Apple iOS 4 (requires an optional AnyConnect Mobile license)
 - Windows Mobile 5.0, 6.0, and 6.1 (Professional and Classic) (requires an optional AnyConnect Mobile license)
 - Android 2.1 and up
- **Wide range of deployment and connection options:**
 - Predeployment, including Microsoft Installer, or automatic headend deployment (administrative rights are required for initial installation) via ActiveX (Windows only) and Java
 - Connection modes:
 - Standalone via system icon
 - Browser initiated (web launch)
 - Clientless portal initiated
 - CLI initiated
 - Application programming interface (API) initiated
- **Ease of client administration:** Allows an administrator to automatically distribute software and policy updates from the headend security appliance, eliminating administration that is associated with client software updates.

- **Preconnection posture assessment (Premium license required):** In conjunction with Cisco Secure Desktop, Host Scan verification checking seeks to detect the presence of antivirus software, personal firewall software, and Windows service packs on the endpoint system before granting network access.
 - **Client firewall policy:** Added protection for split tunneling configurations, used in conjunction with Cisco Secure Mobility to allow for local access exceptions (for example, printing, tethered device support, and so on).
-

The Many Faces of AnyConnect

When we hear references to AnyConnect, most of us think of the SSL VPN full tunnel client mode. However, AnyConnect is now much more than that. AnyConnect Release 3.0 integrates many modules that offer functionality on wired, wireless, and remote-access networks; many of those modules were standalone clients in the past. As an example, in the past, looking under your Start menu, you might see a Cisco SSL VPN client, a Cisco 802.1x client, a Cisco ScanSafe client, a Cisco NAC client, a Cisco VPN IPsec client, and so forth. Those independent Cisco clients have been migrated to modules that now are integrated under the Cisco AnyConnect umbrella. So, Cisco AnyConnect now is your “network access unified client.” In this chapter, however, we will strictly see the SSL VPN full tunnel client mode aspect of AnyConnect.

SSL VPN on Cisco ASA in Clientless Mode

The following is the procedure for configuring a clientless SSL VPN using Cisco ASDM. The configuration process is demonstrated in this section following the presentation of the configuration scenario.

- Task 1.** Launch the Clientless SSL VPN Wizard from ASDM.
- Task 2.** Configure the SSL VPN interface.
- Task 3.** Configure user authentication.
- Task 4.** Configure user group policy.
- Task 5.** Configure a bookmark list.
- Task 6.** Verify the Clientless SSL VPN Wizard configuration.

Clientless Configuration Scenario

[Figure 15-7](#) shows a sample configuration topology that we will use for our scenario. All screen shots that follow are based on this scenario. We will use ASDM to configure the ASA firewall.

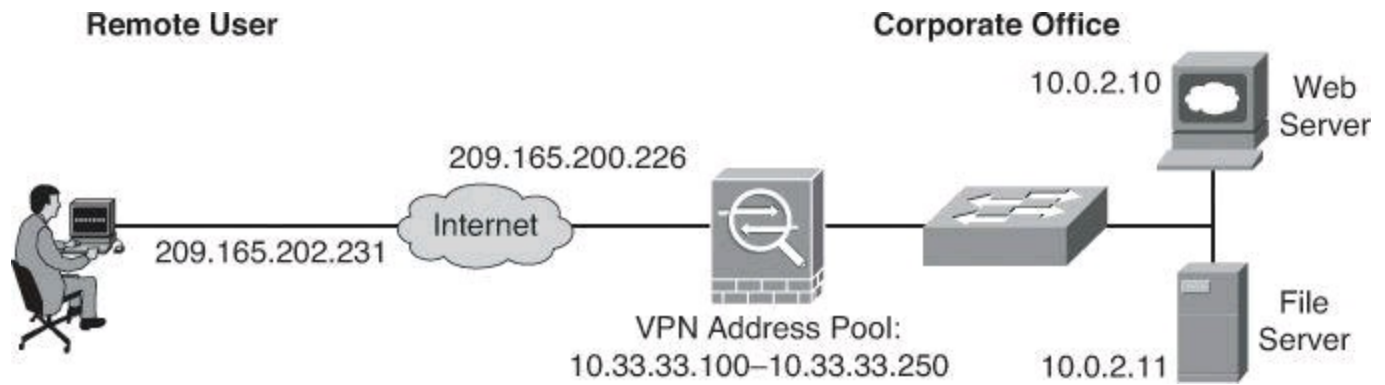


Figure 15-7. Clientless SSL VPN Configuration Topology

In this scenario, a clientless SSL VPN session is established from a remote client and terminates on the ASA. The user will land on the ASA default portal page, customized to include bookmarks to provide access to internal resources such as file servers.

Task 1: Launch the Clientless SSL VPN Wizard from ASDM

In Cisco ASDM, choose **Wizards > VPN Wizards > Clientless SSL VPN Wizard**, as shown in [Figure 15-8](#). When the wizard launches, the first page provides an introduction to the Cisco AnyConnect VPN client product. Although it's not labeled as such, the introduction page is the first step of the wizard.




Figure 15-8. Launching the Clientless SSL VPN Wizard from ASDM on Cisco ASA

Task 2: Configure the SSL VPN Interface

Step 2 of the SSL VPN Wizard is the SSL VPN Interface page. By default, the ASA will use a self-signed certificate to send to the client for authentication. If required, the ASA may be configured to use a certificate that is purchased from a well-known certificate authority, such as VeriSign, to connect clients. In the event that a certificate is purchased, you may select it in the Certificate drop-down menu of the Digital Certificate section, as described in the following steps.

To configure the SSL VPN interface, complete the following steps, demonstrated in [Figure 15-9](#):

Step 1. Enter a name for the clientless SSL VPN connection in the Connection Profile Name field.

Step 2. In the SSL VPN Interface drop-down menu, choose the interface to be used with the

clientless SSL VPN connection.

Step 3. (Optional) Select a third-party certificate that has been installed on the ASA for use in connecting SSL VPN clients. If no certificates were installed, as in our scenario, the ASA will use a self-signed certificate.

Step 4. Click **Next**.

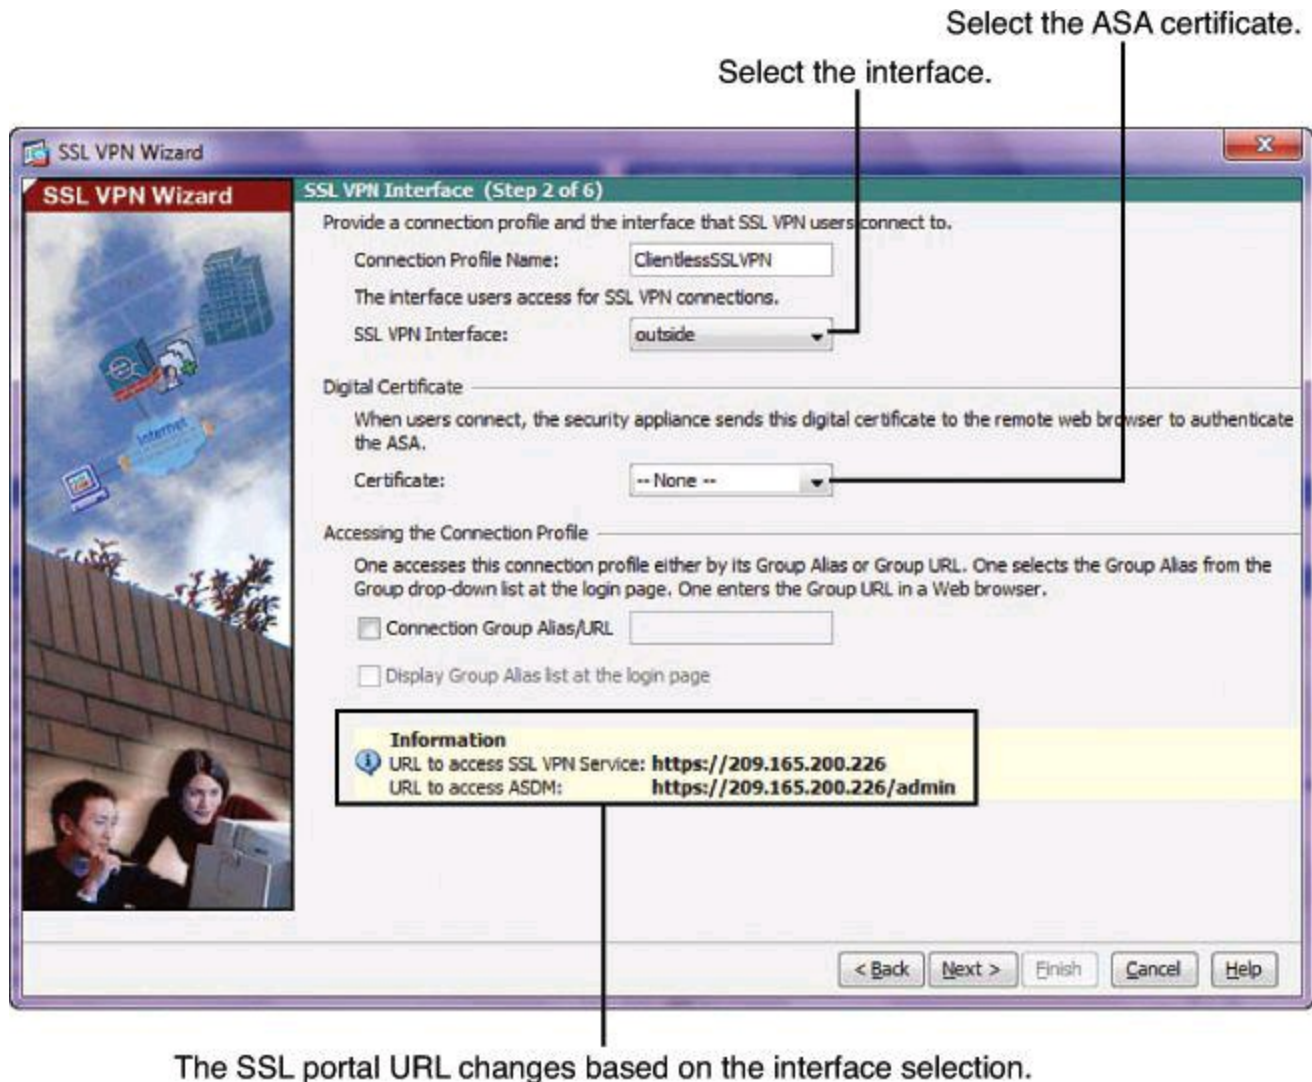


Figure 15-9. Configuring the SSL VPN Interface on Cisco ASA

Note

The SSL VPN Interface screen in the SSL VPN Wizard provides several links in the Information section. These links identify the URLs that need to be used for the SSL VPN service access (login) and for Cisco ASDM access (to access the Cisco ASDM software download).

Task 3: Configure User Authentication

User authentication may be managed by external authentication servers (such as RADIUS) or it may be managed locally by using the ASA local user database.

To configure user authentication using the local user database, complete the following steps, illustrated in [Figure 15-10](#):

Step 1. Click the **Authenticate Using the Local User Database** radio button.

Step 2. Enter a username and password for the desired user.

Step 3. Click **Add** to add the user to the local user database.

Step 4. Click **Next**.

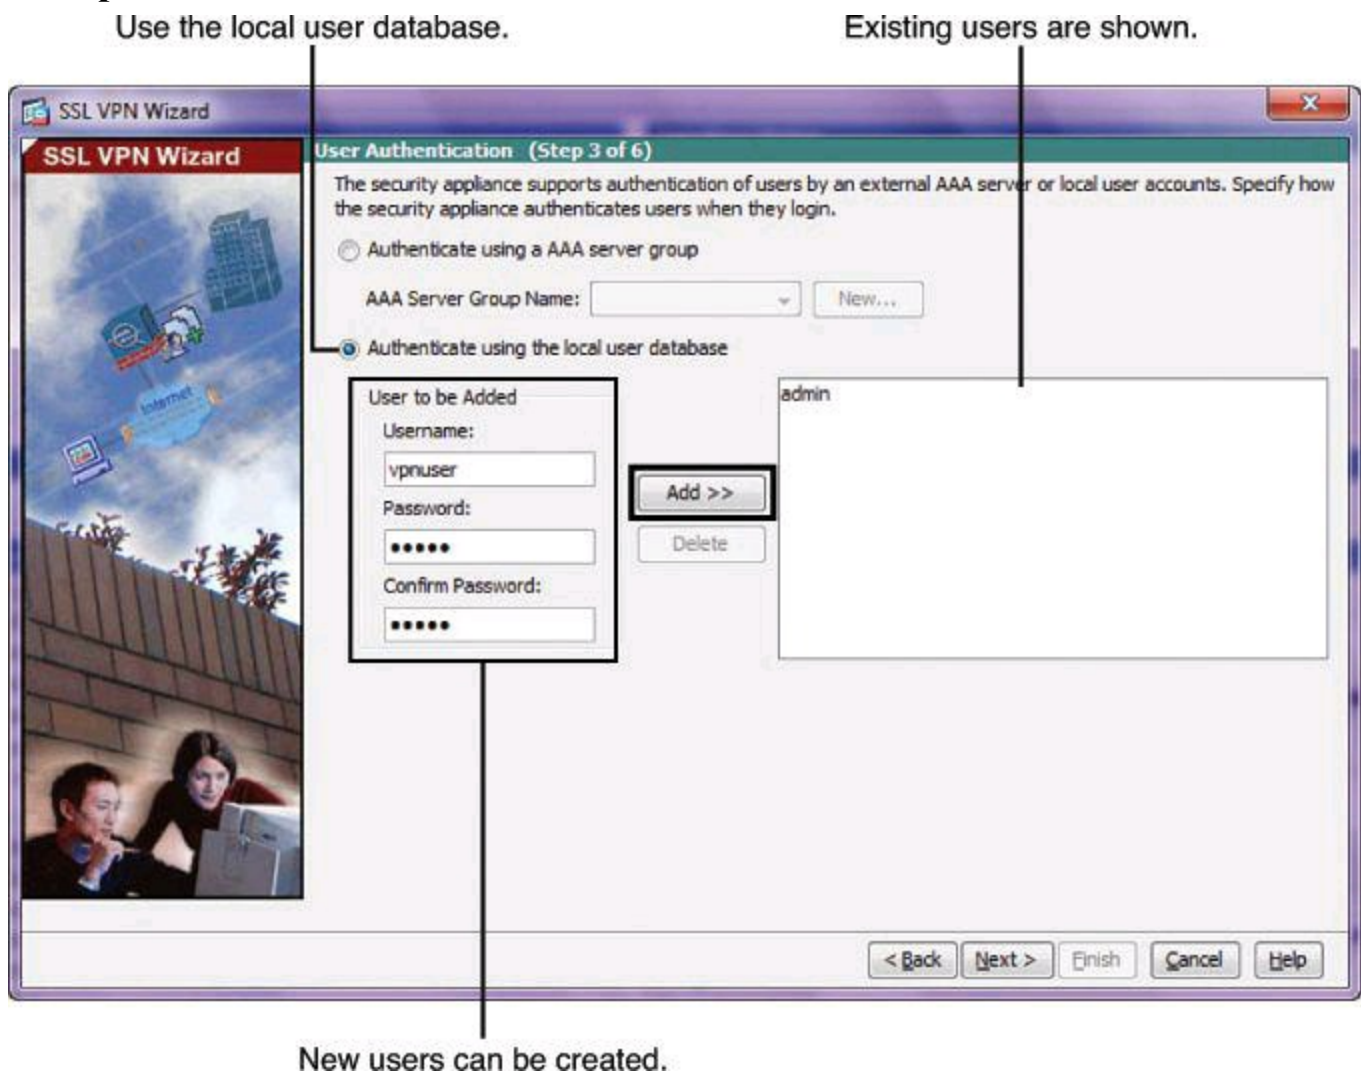


Figure 15-10. Configuring User Authentication for Cisco SSL VPN on Cisco ASA

Task 4: Configure User Group Policy

On the Group Policy wizard page, you may select an existing user group policy to modify or you may add a new user group policy for the clientless SSL VPN connection.

To create a new user group policy, complete the following steps, as shown in [Figure 15-11](#):

Step 1. Click the **Create New Group Policy** radio button.

Step 2. Enter a name for the new user group policy.

Step 3. Click **Next**.

Configure New Group Policy

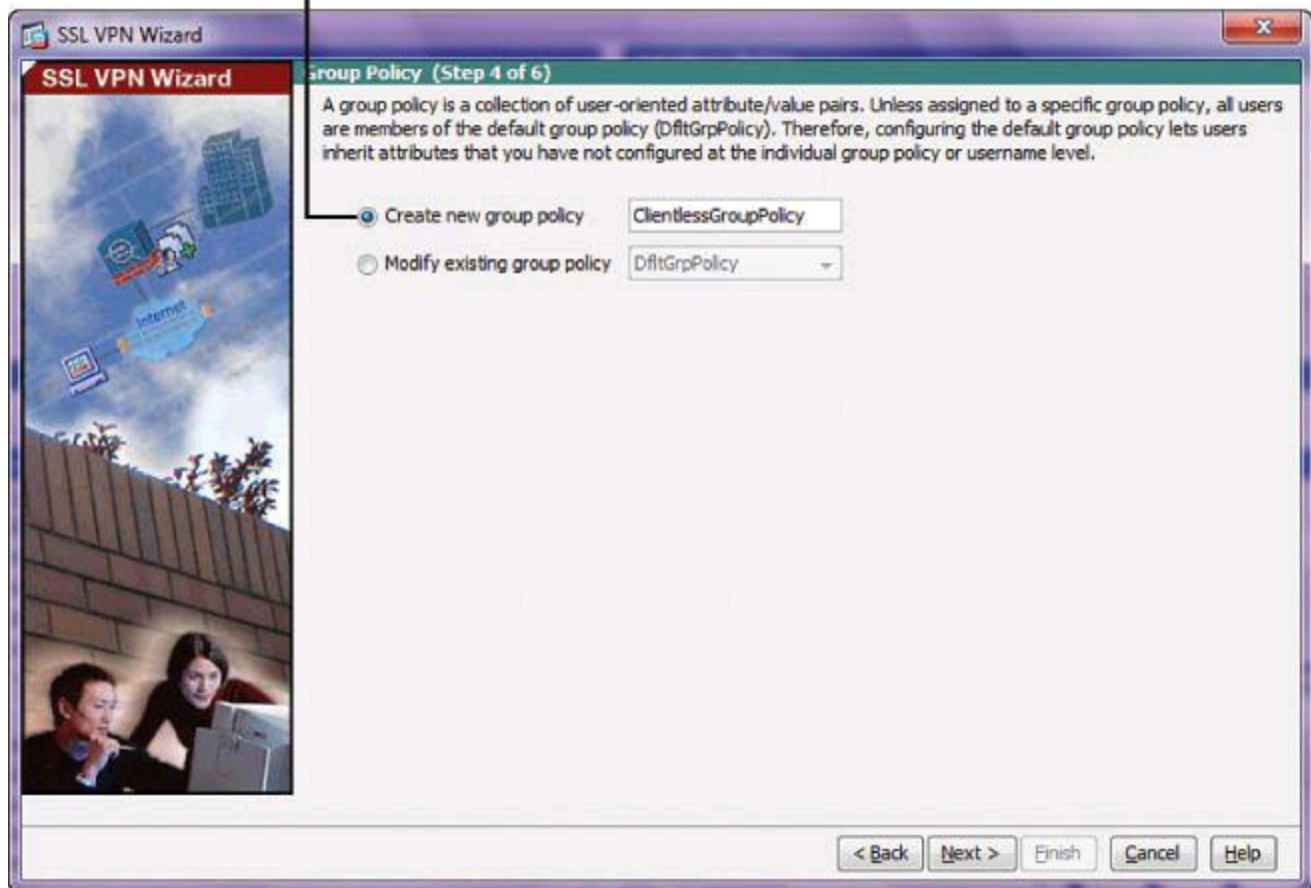


Figure 15-11. Configuring User Group Policy for Cisco SSL VPN on Cisco ASA

Note

By default, the created user group policy inherits its settings from the default group policy, DfltGrpPolicy. You may modify these settings after you have completed the SSL VPN Wizard by navigating to the **Configuration > Remote Access VPN > Clientless SSL VPN Access > Group Policies** submenu.

Task 5: Configure a Bookmark List

A bookmark list is a set of URLs that is configured to be used in the clientless SSL VPN web portal. By default, there are no configured bookmark lists and they must be configured by the network administrator.

To create a bookmark list and add bookmark entries to it, complete the following steps, as shown in [Figures 15-12](#) and [15-13](#):

Step 1. Click **Manage** to create a new bookmark list.

Step 2. In the Configure GUI Customization Objects dialog box, click **Add** to add a bookmark list that is based on the included template.

Step 3. In the Add Bookmark List dialog box, enter the bookmark list name in the corresponding field and click **Add** to create bookmarks for this list.

Step 4. The Add Bookmark dialog box opens, shown in [Figure 15-13](#). Enter a name for the

bookmark in the Bookmark Title field.

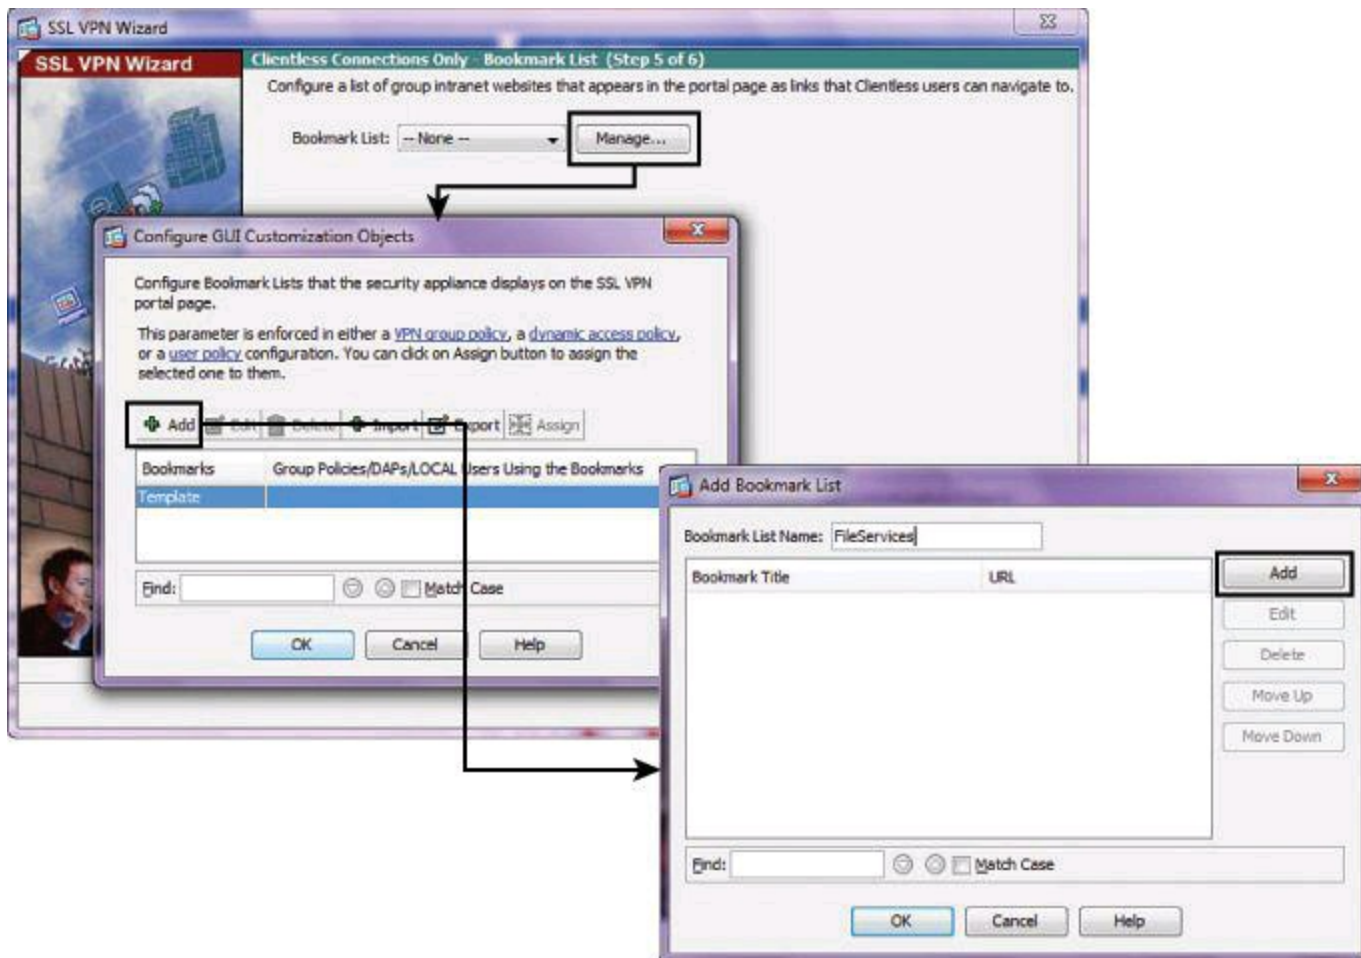


Figure 15-12. Configuring a Bookmark List for Cisco SSL VPN on Cisco ASA

Choose the appropriate protocol.

Define the URL.

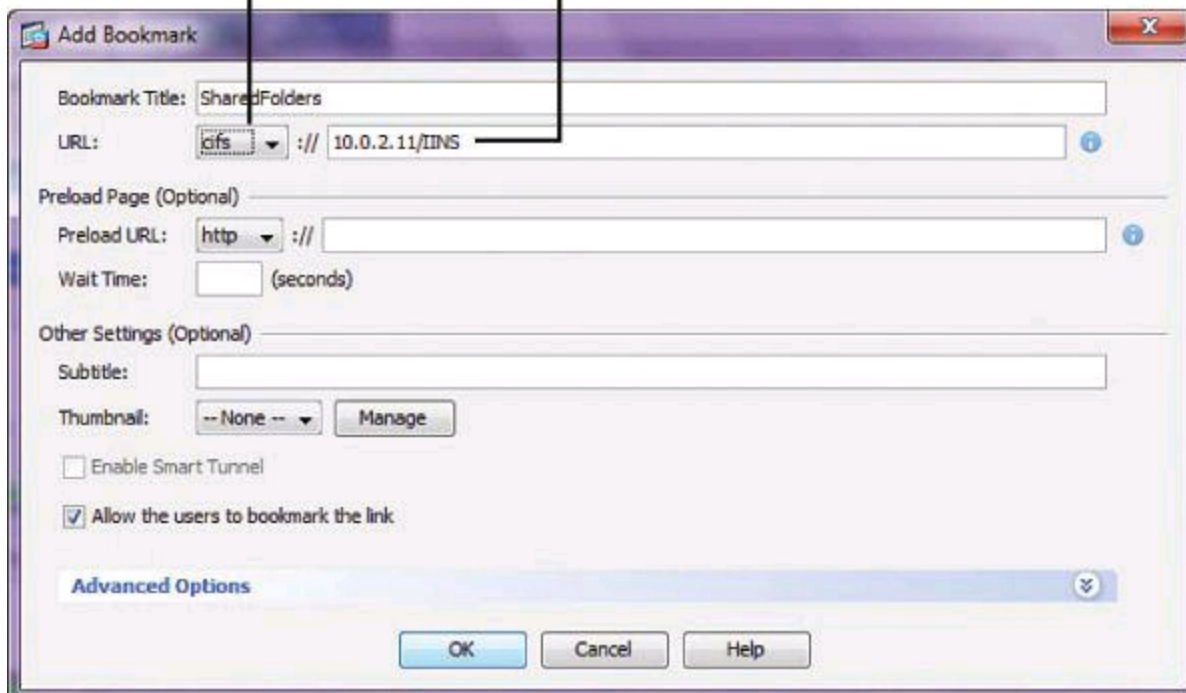


Figure 15-13. Creating a Bookmark List

Bookmark List

Think of the bookmark list as a bookmark container. Within that container, you will have listed multiple URLs for HTTP, HTTPS, FTP, and CIFS servers, which will be demonstrated next. Also, you could have destinations to SSL thin-client services listed in this bookmark list, such as RDP, Telnet, and SSH. For more details on thin clients, refer to *CCNP Security VPN 642-648 Official Cert Guide*, Second Edition, by Howard Hooper (Cisco Press, 2012).

Step 5. Enter the URL value for the bookmark as cifs. The following list of URL values is displayed in the drop-down menu:

- http
- https
- cifs
- ftp

Step 6. Enter the server destination IP address or hostname to be used with the bookmark entry.

Step 7. (Optional) Enter the URL endpoint browser that can fetch certain information that is sent along to the webserver or web application.

Step 7. (Optional) Enter the name in the Subtitle field. The subtitle will appear under the bookmark entry on the web portal.

Step 8. (Optional) Select the thumbnail image to be used with this bookmark entry in the Thumbnail field.

When you're finished creating a bookmark and click OK, you are brought back to the Add Bookmark List page. Click OK on the Add Bookmark List page to return to the wizard bookmark list page, where you click Next to return to the Summary page.

Note

To use thumbnails with bookmarks, the images must first be uploaded to the ASA.

Task 6: Verify the Clientless SSL VPN Wizard Configuration

Verify that the information that was entered in the SSL VPN Wizard is correct. As shown in [Figure 15-14](#), click **Finish** to finish the wizard and send the configuration to the ASA.

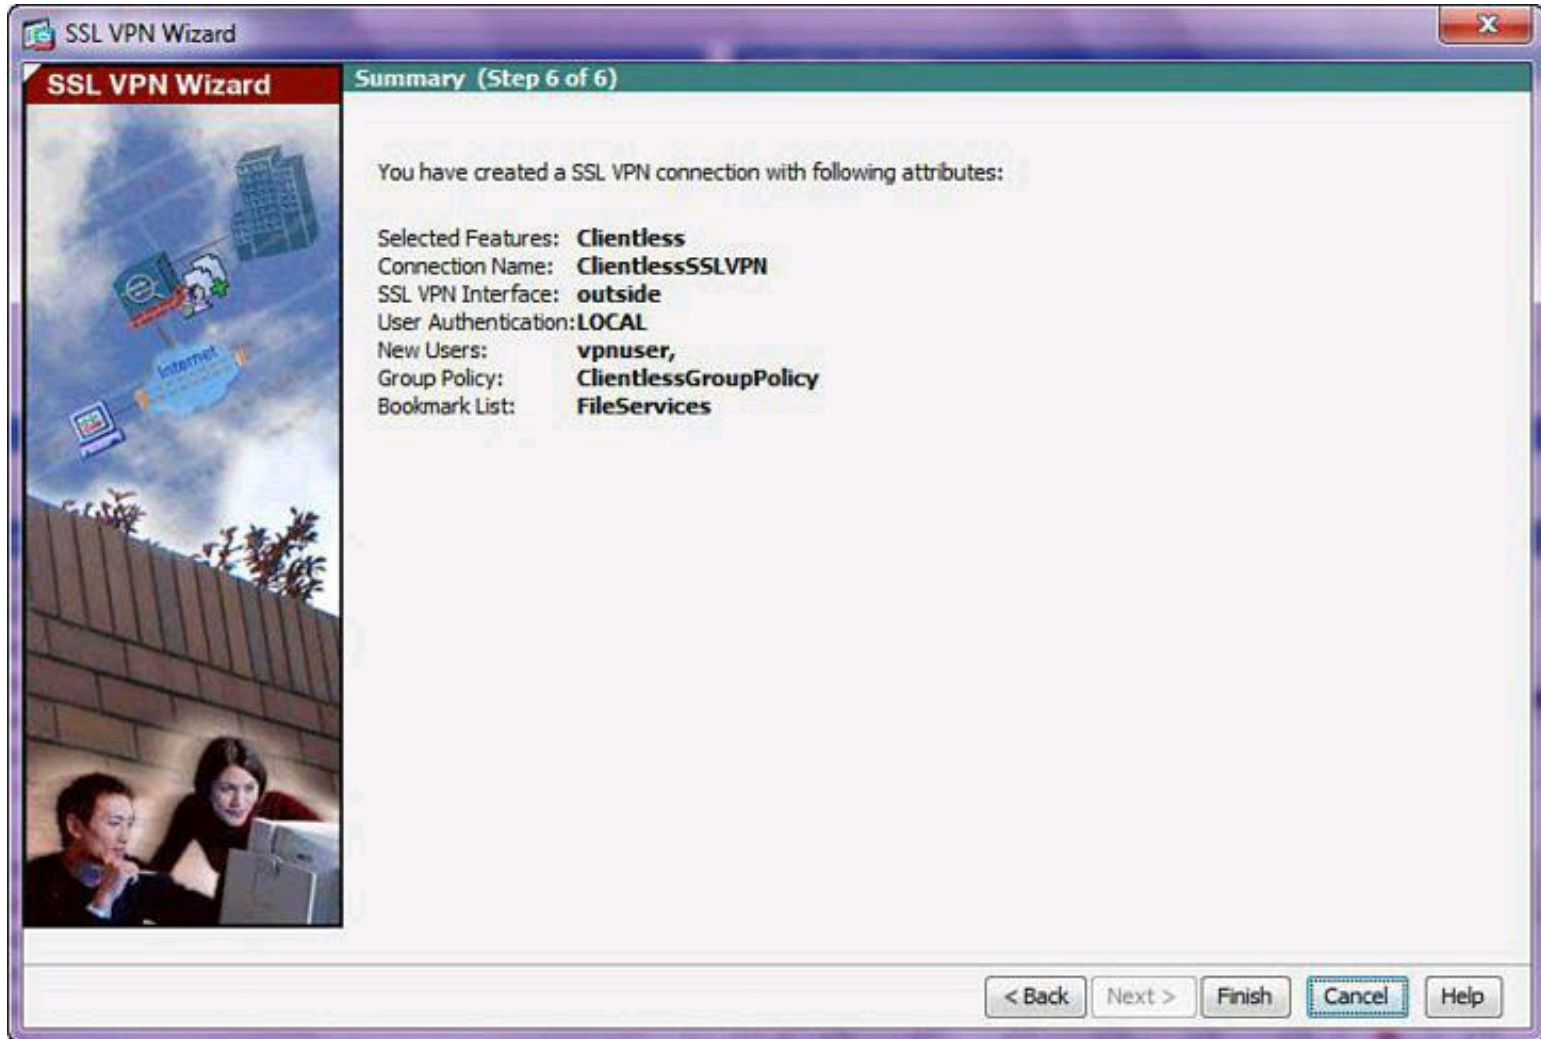


Figure 15-14. Verifying the SSL VPN Configuration Created by the ASDM SSL VPN Wizard

Log In to the VPN Portal: Clientless SSL VPN

For verification, open a compliant web browser and enter the login URL for the SSL VPN into the address field, as shown in the upper-left portion of [Figure 15-15](#). The address that you enter should be the configured address URL.

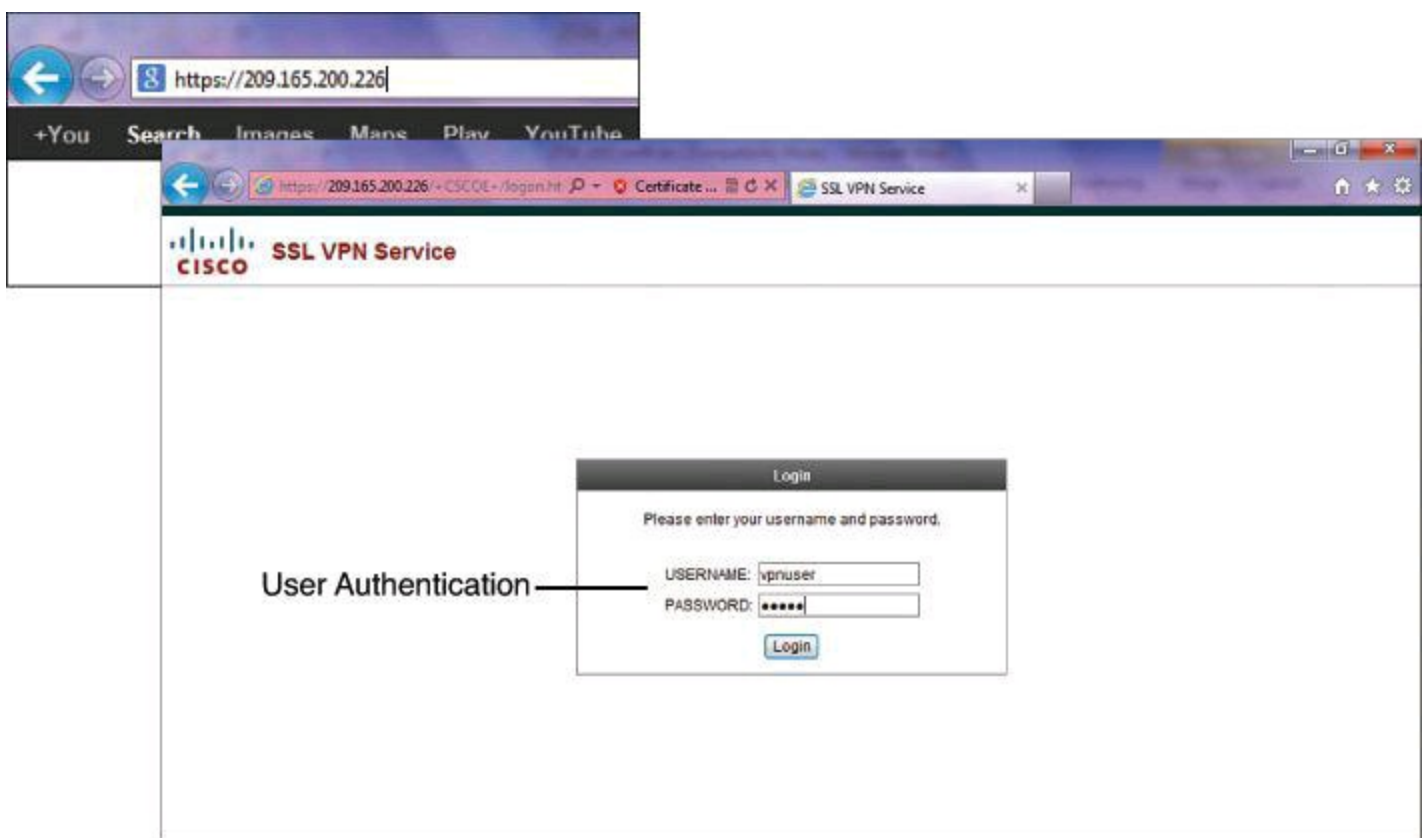


Figure 15-15. Logging In to the VPN Portal Using Clientless SSL VPN

Enter the previously configured username and password and click **Login**.

Note

The prefix to the path changes depending on whether you require authentication. The security appliance uses `/+CSCOE+` for objects that require authentication, and uses `/+CSCOU+` for objects that do not. The security appliance displays `/+CSCOE+` objects on the portal page only, while `/+CSCOU+` objects are visible and usable in either the login page or the portal page.

Once the user is logged in, the main portal page will be displayed to the user, as shown in [Figure 15-16](#). The default homepage displays configured web application bookmarks and Common Internet File System (CIFS) bookmarks. In the event that many bookmarks are available, one or the other type of bookmarks may be selected for display by using the buttons in the left navigation pane.

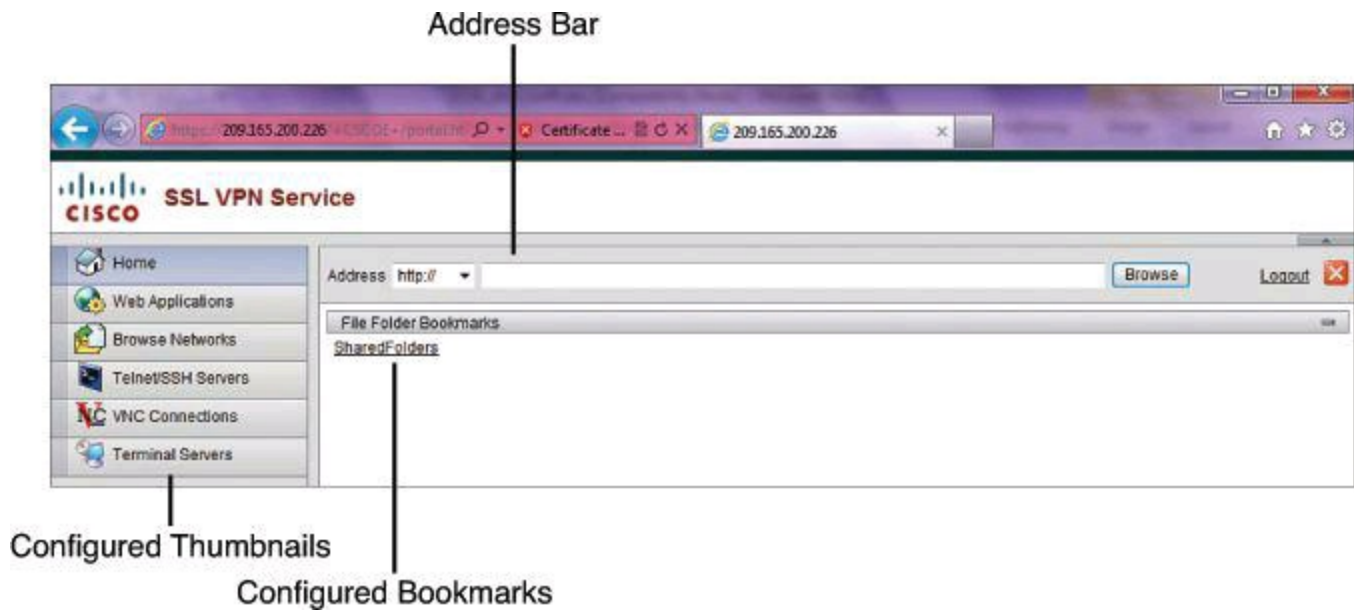


Figure 15-16. Resources Accessible in the Portal

To the left of each configured bookmark is the imported thumbnail image file that is used by the bookmark. If there are other sites that must be accessed but that are not configured as bookmarks, they may be accessed using the Address bar at the top of the portal page.

Portal Customization

The login page shown in [Figure 15-15](#), the portal page shown in [Figure 15-16](#), and the logout page (not shown) can all be customized with colors, corporate logos, and so forth. More information on this topic can be found in the Cisco Press book *CCNP Security VPN 642-648 Official Cert Guide, 2nd Edition*.

SSL VPN on ASA Using the Cisco AnyConnect VPN Client

Suppose that you need more functionality than what Clientless SSL VPN, via the portal, can give you. For example, let's say you are an engineer and need to run customized software that reaches the computer-aided design (CAD) server and prints on a plotter.

Because Clientless SSL VPN doesn't provide you with the adequate functionality, you wish to use the Full client. You will need to configure both the ASA and the client. Though the Remote Access client, Cisco AnyConnect VPN Client gives you a choice of IPsec (IKEv2 only) or SSL full client remote access; in our scenario, we will focus only on the SSL VPN full client mode.

There are three major phases to configuring SSL VPN full-tunnel mode using Cisco ASDM so that remote clients will connect using Cisco AnyConnect:

- Phase 1.** Configure Cisco ASA for Cisco AnyConnect.
- Phase 2.** Configure the Cisco AnyConnect VPN Client.
- Phase 3.** Verify VPN Connectivity with Cisco AnyConnect.

These phases are discussed in turn following the presentation of the configuration scenario.

Cisco AnyConnect Configuration Scenario

For this scenario, we will use the same topology as for the portal, as shown earlier in [Figure 15-7](#). We will add one element to this topology: an address pool for the VPN AnyConnect client. The VPN gateway, which in our example is the Cisco ASA, will provide IP addresses to connecting clients using the address pool 10.33.33.100 to 10.33.33.200.

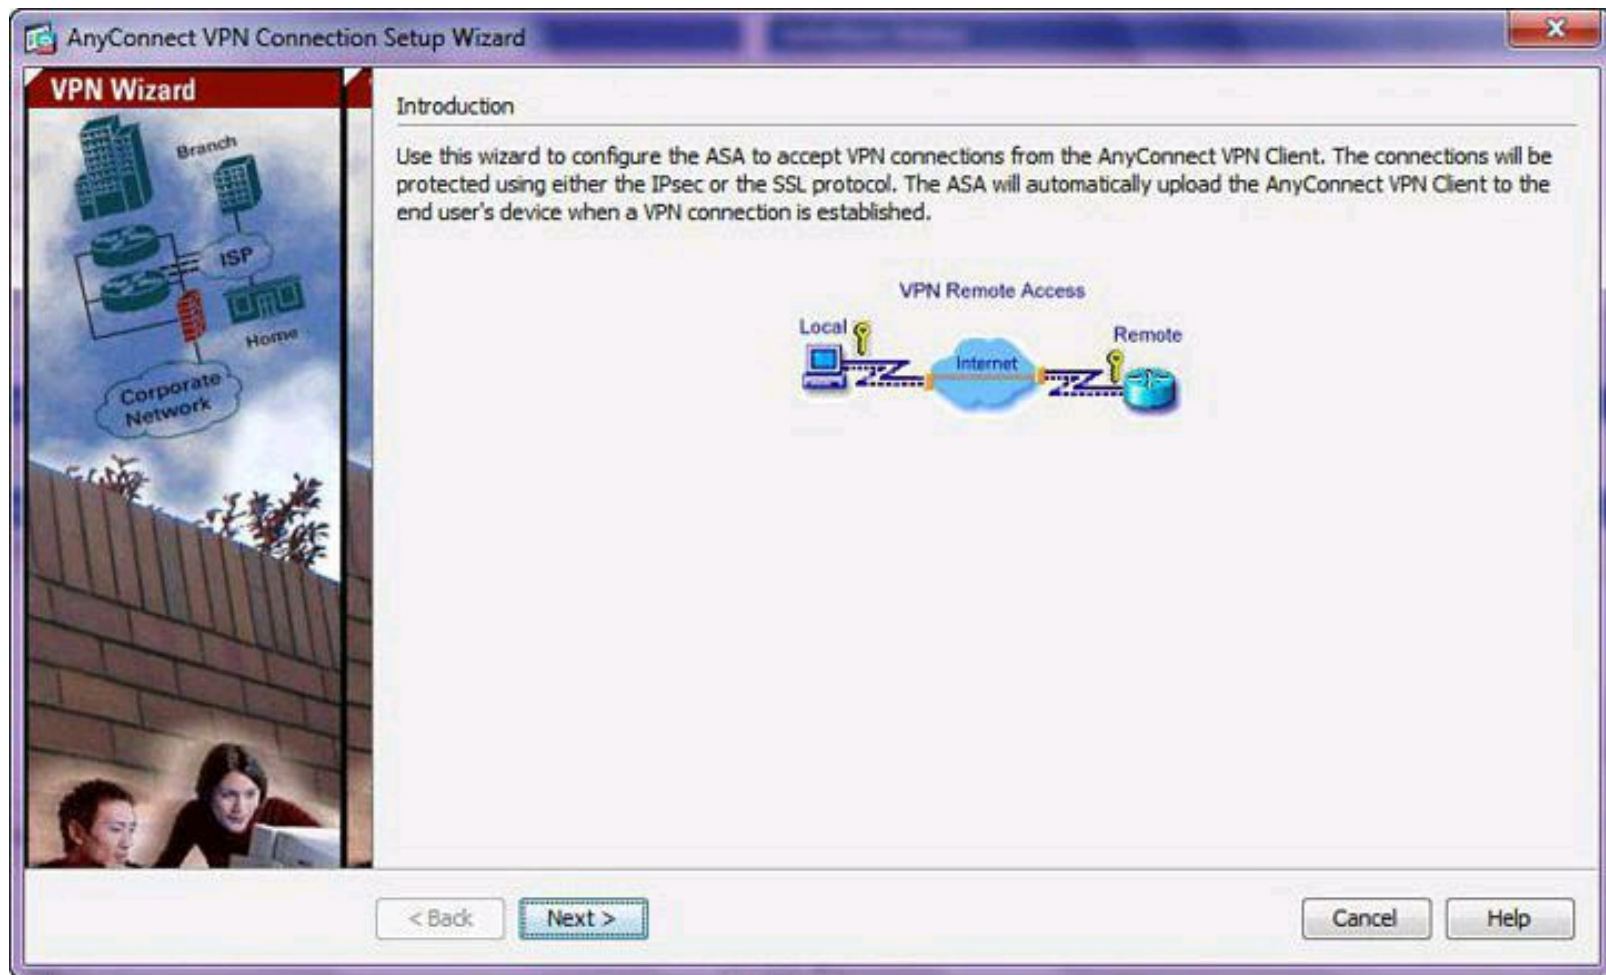


Figure 15-17. Launching the AnyConnect VPN Connection Setup Wizard from Cisco ASDM

In our scenario, we will first initiate a clientless session. After working from the web clientless portal, we will discover that a full VPN tunnel is required to progress in our work. From the clientless portal, we will click the AnyConnect option to proceed with the Cisco AnyConnect VPN Client download and installation. We will then start an SSL VPN full-tunnel session using the newly installed AnyConnect client. All screen shots that follow are based on this scenario. We will use ASDM to configure the Cisco ASA firewall.

Phase 1: Configure Cisco ASA for Cisco AnyConnect

Eight tasks are required for configuring the Cisco ASA for AnyConnect, which will be further outlined in this section:

1. Configure the connection profile.
2. Configure VPN protocols and the device certificate.
3. Configure the client image.
4. Configure the authentication methods.
5. Configure the client address management.

6. Configure the network name resolution servers.
7. Configure the network address translation exemption.
8. Configure the AnyConnect client deployment summary.

To configure the Cisco ASA, we will use the AnyConnect VPN Connection Setup Wizard. From Cisco ASDM, choose **Wizards > VPN Wizards > AnyConnect VPN Wizard** to open the Introduction page of the wizard, shown in [Figure 15-17](#).

Task 1: Connection Profile Identification

After you click **Next** on the Introduction page of the wizard, the first task is to configure the connection profile identification, used to identify the ASA to the remote-access users, as shown in [Figure 15-18](#).

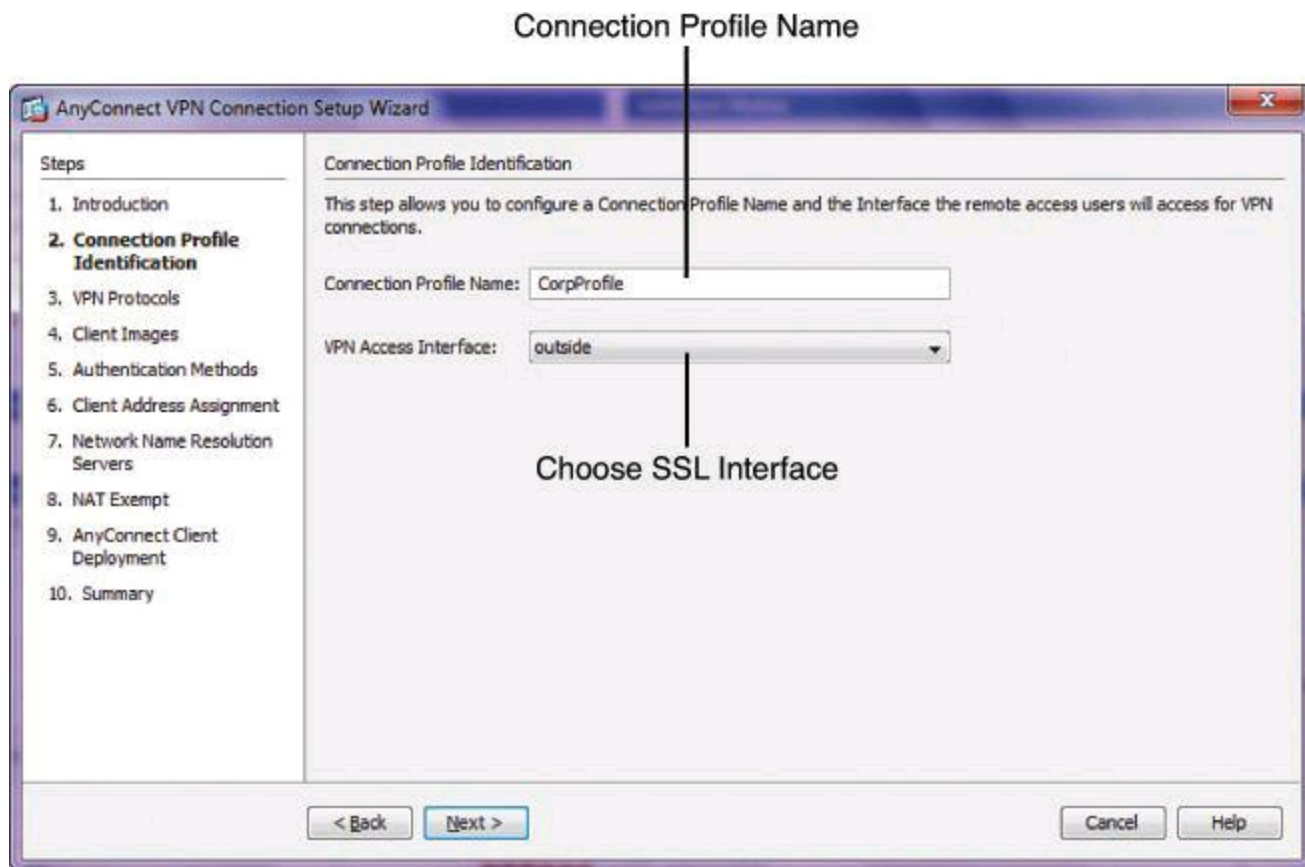


Figure 15-18. Setting the Connection Profile Identification

Complete the following steps:

Step 1. In the Connection Profile Name field, enter a name that remote-access users will access for VPN connections. Connection profiles appear as **tunnel-group** in the CLI. This connection profile name will be displayed in a drop-down list on the initial clientless session, allowing the user to log in using the profile and download the Cisco AnyConnect VPN Client. The drop-down list appears whenever you have created connection profiles in addition to the default connection profile named DefaultRAGroup—it is called group because what is called a connection profile within ASDM is called a Tunnel-Group at the CLI. So, connection profiles (found in ASDM) are also called Tunnel-Group (at the CLI).

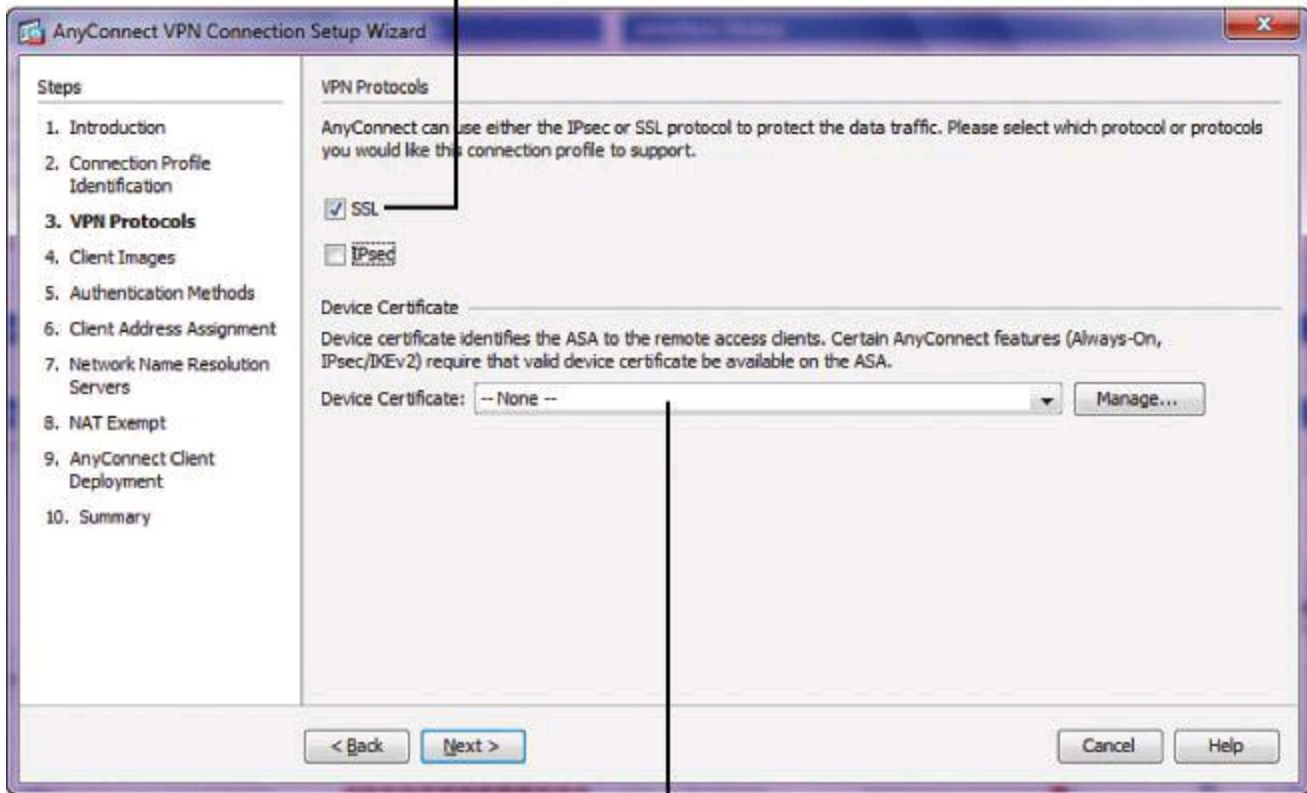
Step 2. In the VPN Access Interface drop-down menu, choose an interface that remote-access users will access for VPN connections.

Step 3. Click **Next** to continue to the next page of the AnyConnect VPN Connection Setup Wizard.

Task 2: VPN Protocols and Device Certificate

The next task is to specify the VPN protocol allowed for this connection profile, as shown in [Figure 15-19](#). The Cisco AnyConnect VPN Client defaults to SSL. If you enable IPsec as a VPN tunnel protocol for the connection profile, you must also create and deploy a client profile with IPsec enabled using the profile editor from Cisco ASDM. AnyConnect IPsec supports only IKEv2.

Select protocols. Selecting both is an option.



No certificate selection results in use of ASA self-signed certificate.

Figure 15-19. Selecting the VPN Protocol and the Certificate

Complete the following steps to define protocols and certificates:

Step 1. Check the check box for each of the desired protocols. In this example, only SSL is chosen.

Step 2. Choose an identity certificate from the Device Certificate drop-down list. This certificate identifies the ASA to remote-access clients. Clicking the Manage button opens the Manage Identity Certificates dialog box, which allows you to create new certificates and manipulate existing certificates.

Step 3. Click **Next** to continue to the next page of the AnyConnect VPN Connection Setup Wizard.

Task 3: Client Image

ASA can automatically upload the latest Cisco AnyConnect VPN Client package to the client device when it accesses the enterprise network. You can use a regular expression to match the user

agent of a browser to an image. You can also minimize connection setup time by moving the most commonly encountered operating system for your organization, for example Windows, to the top of the list.

To configure the location of the Cisco AnyConnect SSL VPN Client, complete the following steps:

Step 1. Click **Add** to add a new image definition, as shown in [Figure 15-20](#).

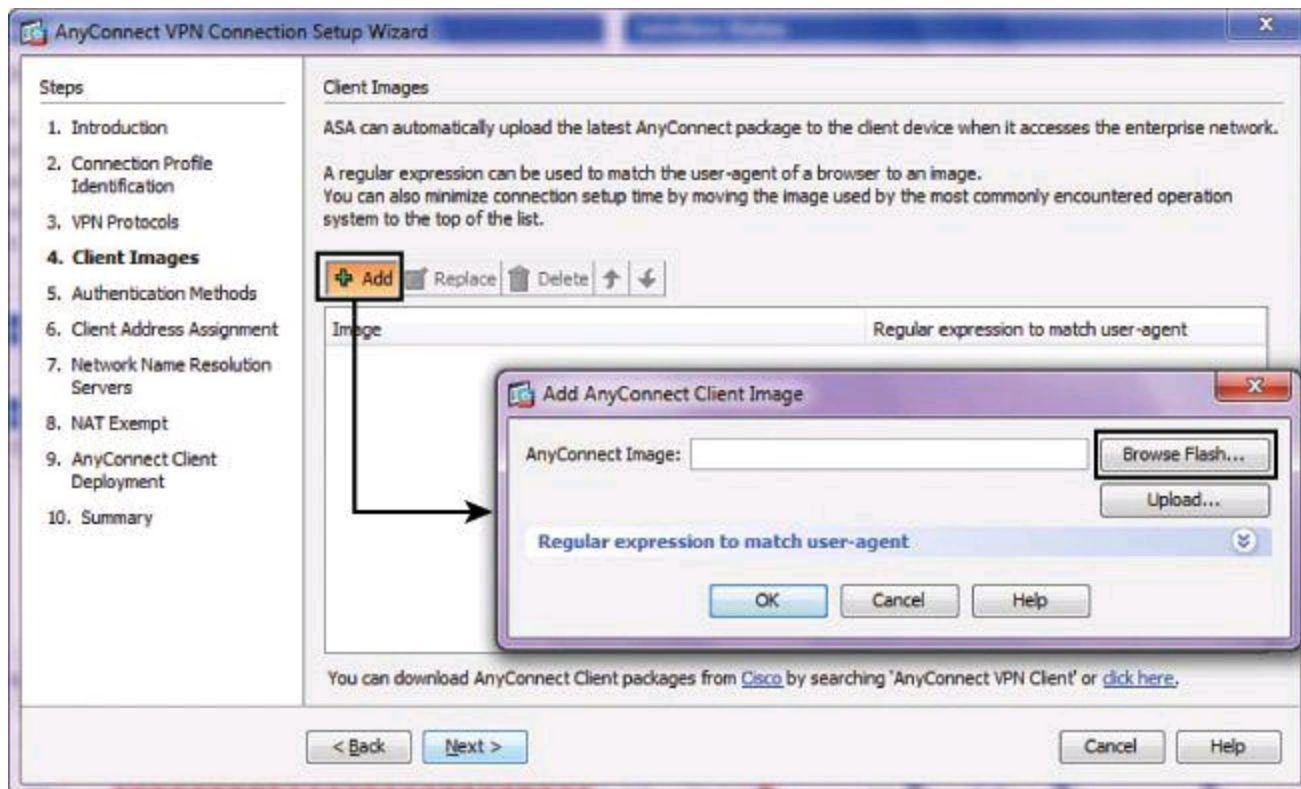


Figure 15-20. Browsing to Select the Client Image

Step 2. Click **Browse Flash** if the image file is already located on the Cisco ASA, or click **Upload** if you have a copy on the local machine to upload to the security appliance.

Step 3. The Browse Flash dialog box opens, as shown in [Figure 15-21](#), allowing you to navigate through files and folders in flash memory. Select the image of the Cisco AnyConnect VPN Client, and click OK to continue.

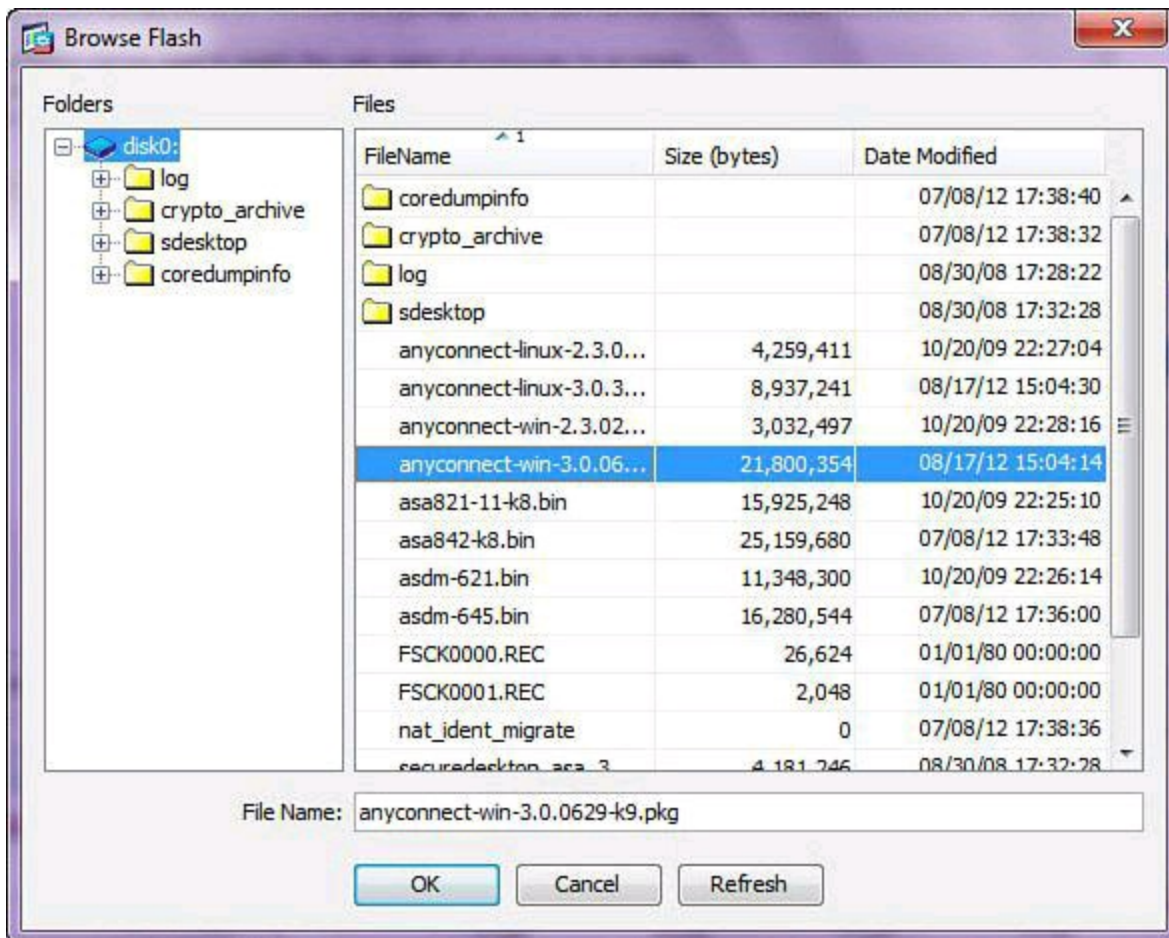


Figure 15-21. Selecting the Client Image

Step 4. Click **Next** to continue to the next page of the AnyConnect VPN Connection Setup Wizard.

Task 4: Authentication Methods

User authentication may be managed by external authentication servers (such as RADIUS) or it may be managed locally using the ASA local user database.

In [Figure 15-22](#), the requirements call for configuring user authentication using the local user database.

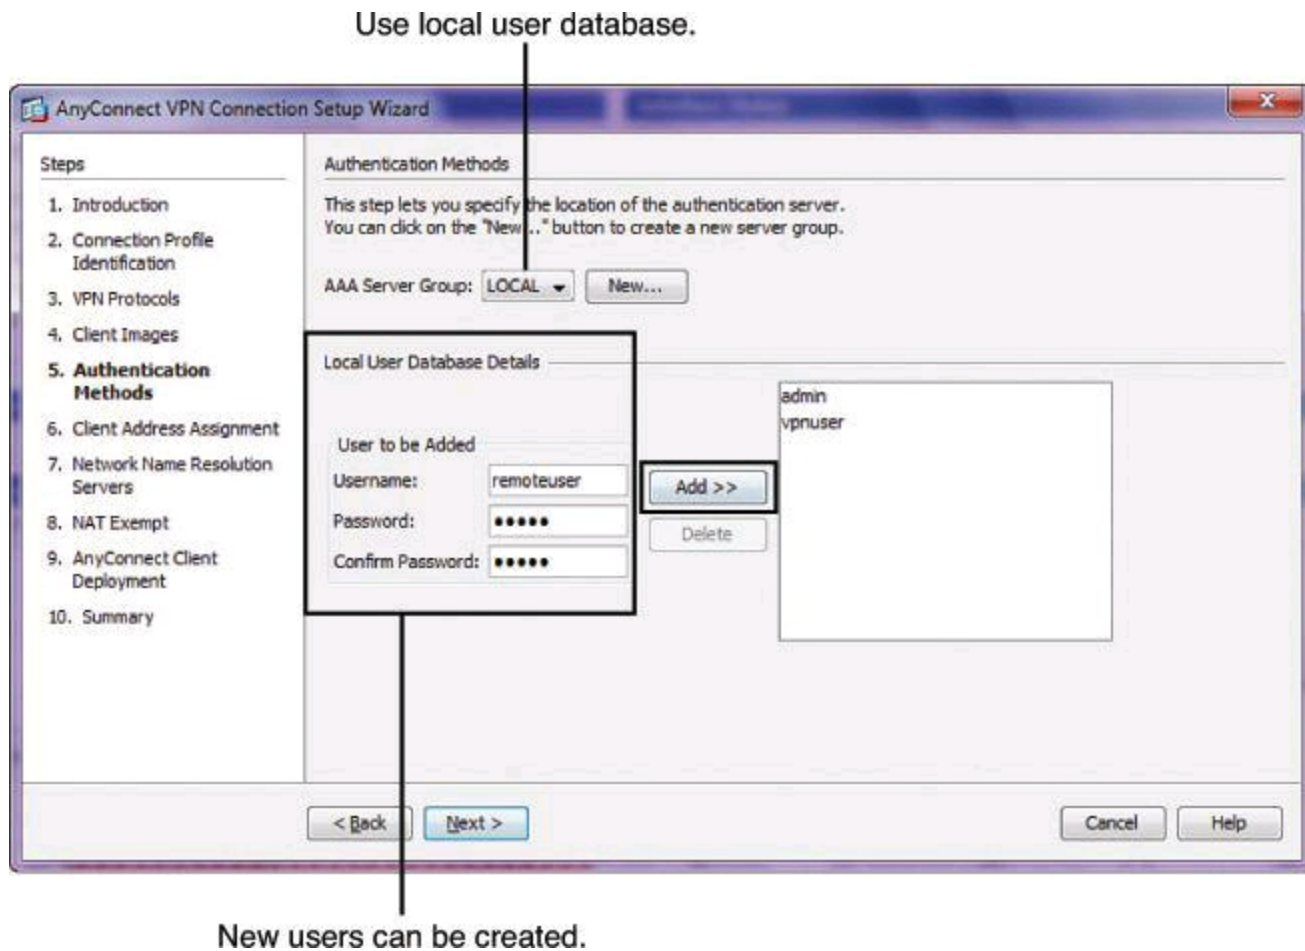


Figure 15-22. Selecting the Authentication Method

Complete the following steps:

- Step 1.** Choose the LOCAL value from the AAA Server Group drop-down menu.
- Step 2.** Configure a username and password for the desired user. Notice that existing users appear in the pane on the right.
- Step 3.** Click **Add** to add the user to the local user database. The username will move to the pane on the right.
- Step 4.** Click **Next** to continue to the next page of the SSL VPN Connection Setup Wizard.

Task 5: Client Address Assignment

SSL VPN clients receive new IP addresses when they connect to the ASA. Clientless connections do not require new IP addresses. Address pools define a range of addresses that remote clients can receive. The IP address pool configuration is required for successful client-based SSL VPN connectivity. Without an available IP address pool, the full-tunnel Cisco AnyConnect SSL VPN to the security appliance will fail.

Full Client Virtual Adapter

Later in this chapter, we will see that a remote-access user wishing to use a Full client—usually to have more functionality than what the clientless portal can offer—sees a virtual adapter added to the list of network adapters in his computer. We are used to seeing wireless and wired adapters listed under network properties. These network adapters, to fully join a network, usually use DHCP to receive an IP address, a DNS, and so forth.

AnyConnect installs a virtual adapter on your workstation, which enables your workstation to look and feel as if it's connected on the corporate network, though you are physically a distance away. This virtual adapter, similar to your wired and wireless adapters, requires and receives from the ASA an IP address, and a DNS server, so it becomes a client of your corporate LAN. So, for the next few tasks, we will configure for AnyConnect connectivity similar attributes that we would need to configure for our DHCP clients.

Using the options on the Client Address Assignment wizard page, you can select an existing IP address pool or click New to create a new pool, as shown in [Figure 15-23](#). For our particular scenario, complete these steps as follow:

Step 1. Click **New** to define the desired address pool.

Step 2. In the Add IP Pool dialog box, create the pool by configuring the following items:

- **Name:** A name to be associated with the IP address pool.
- **Starting IP Address:** Starting IP address of the range to be assigned to the client SSL VPN connections.
- **Ending IP Address:** Ending IP address of the range to be assigned to the client SSL VPN connections.
- **Subnet Mask:** Choose the desired subnet mask of the IP address pool from the drop-down menu.

Step 3. Click **OK** to accept the newly configured IP address pool.

Step 4. Click **Next** to continue to the next page of the SSL VPN Connection Setup Wizard.

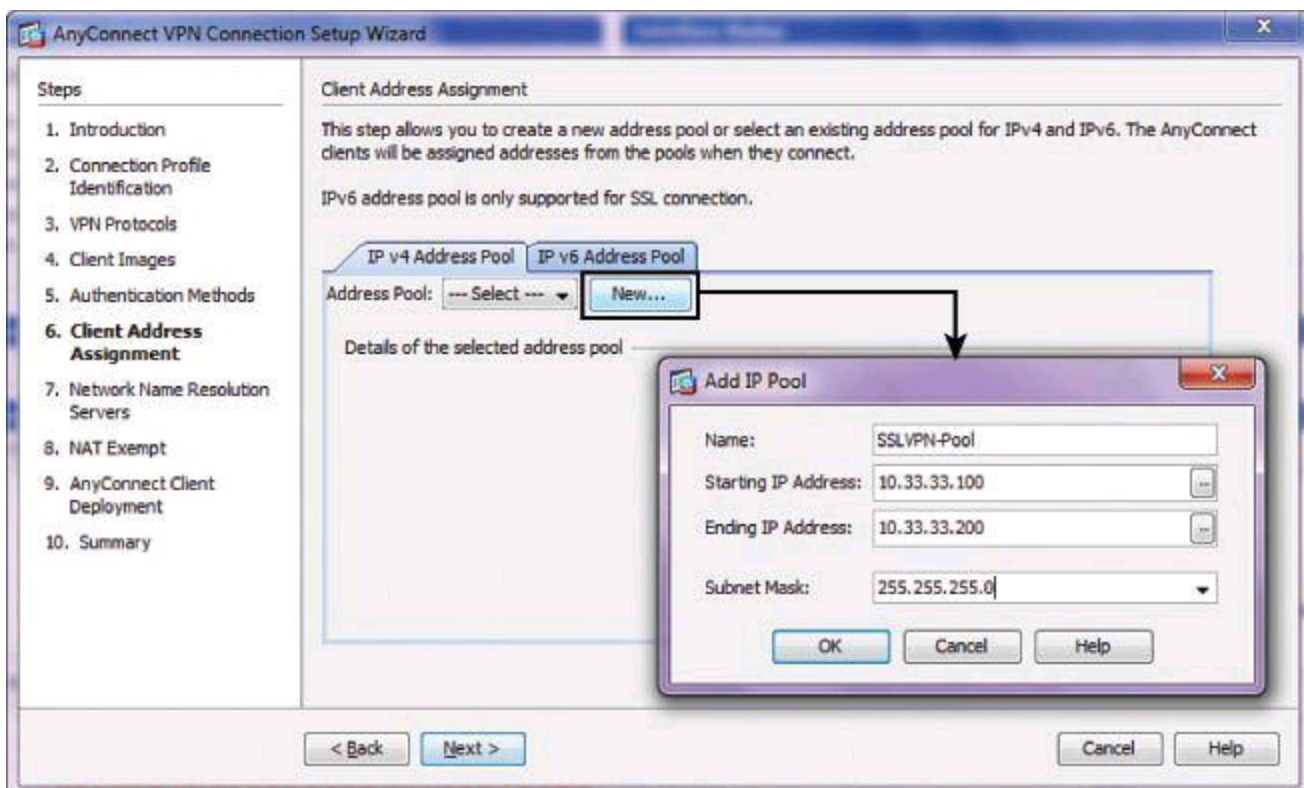


Figure 15-23. Selecting the Address Pool That Will Be Assigned to AnyConnect Clients

Task 6: Network Name Resolution Servers

The next step, shown in [Figure 15-24](#), lets you specify which domain names are resolved for the remote user when accessing the internal network. Simply complete these fields, if necessary:

- **DNS Servers:** Enter the IP address of the DNS server.
- **WINS Servers:** Enter the IP address of the WINS server.
- **Domain Name:** Type the default domain name.

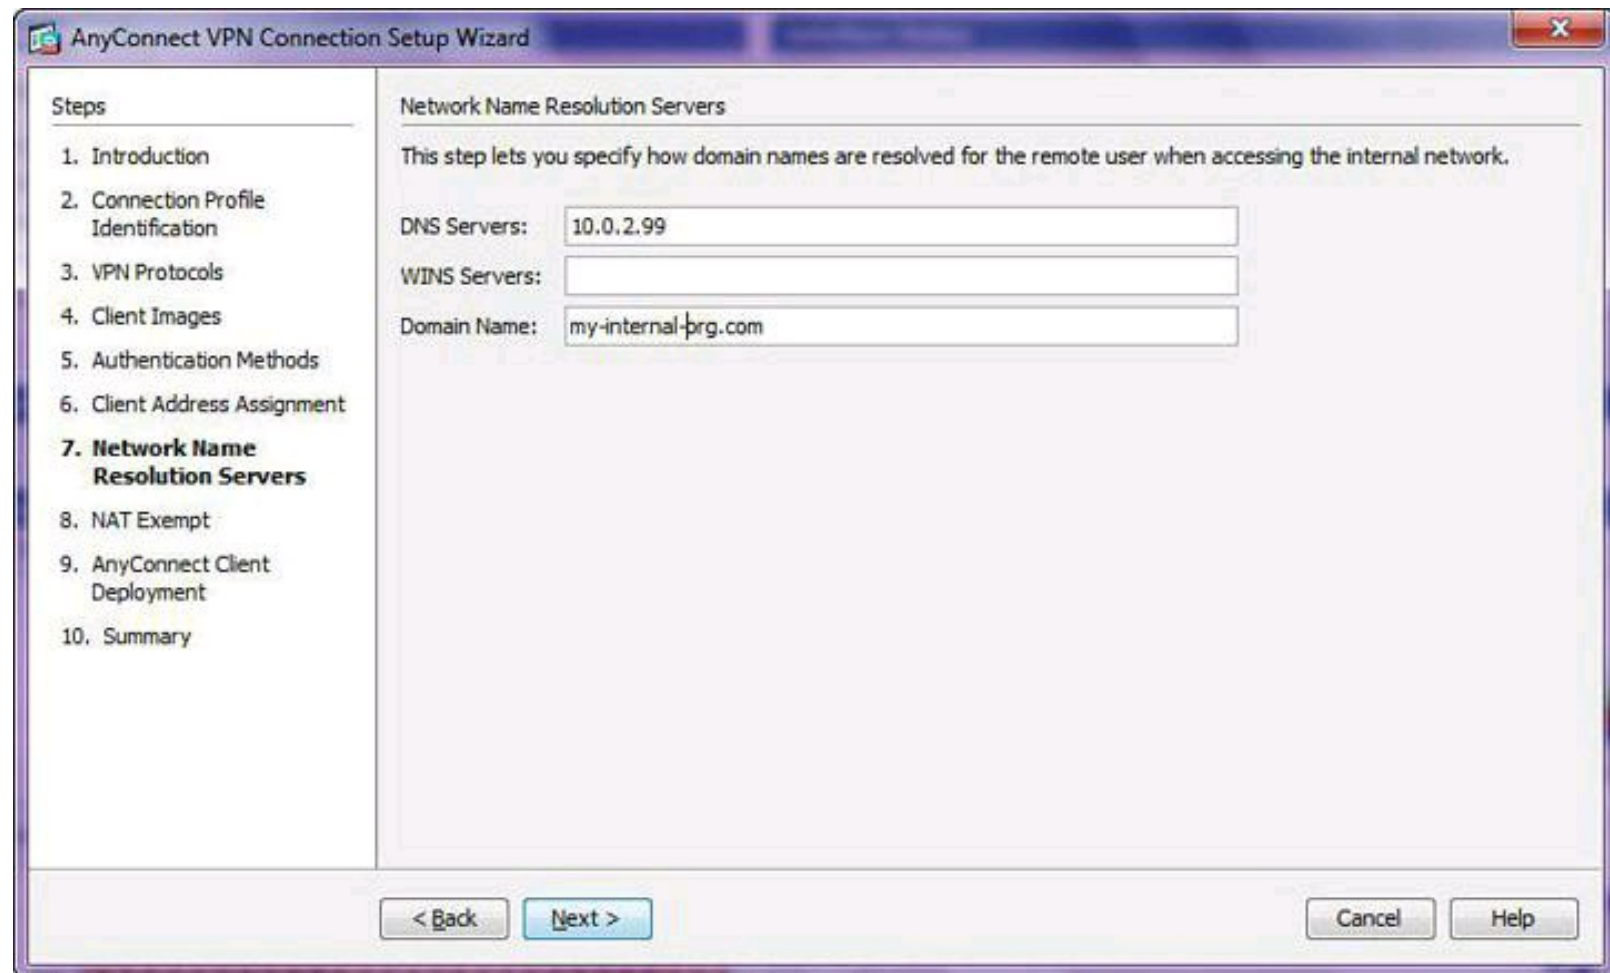


Figure 15-24. Selecting Network Name Resolution Servers

Click **Next** to continue to the next page of the SSL VPN Connection Setup Wizard.

Task 7: Network Address Translation Exemption

If NAT is enabled on the ASA, the VPN traffic must be exempt from this translation. Complete the following steps, shown in [Figure 15-25](#), to configure NAT exemption, which is necessary in our scenario:

Step 1. Check the **Exempt VPN Traffic from Network Address Translation** check box.

Step 2. Choose the interface connected to the internal network that needs to be exempted from NAT. In this example, the inside interface contains the hosts that the VPN clients will access.

Step 3. Type the host or subnet IP addresses that are to bypass NAT in the Local Network field. You can also click the ... button and select the addresses from the Browse Local Network dialog box. The traffic between Cisco AnyConnect VPN Clients and the internal networks defined here will be exempt from NAT.

Step 4. Click **Next** to move to the next page of the SSL VPN Connection Setup Wizard.

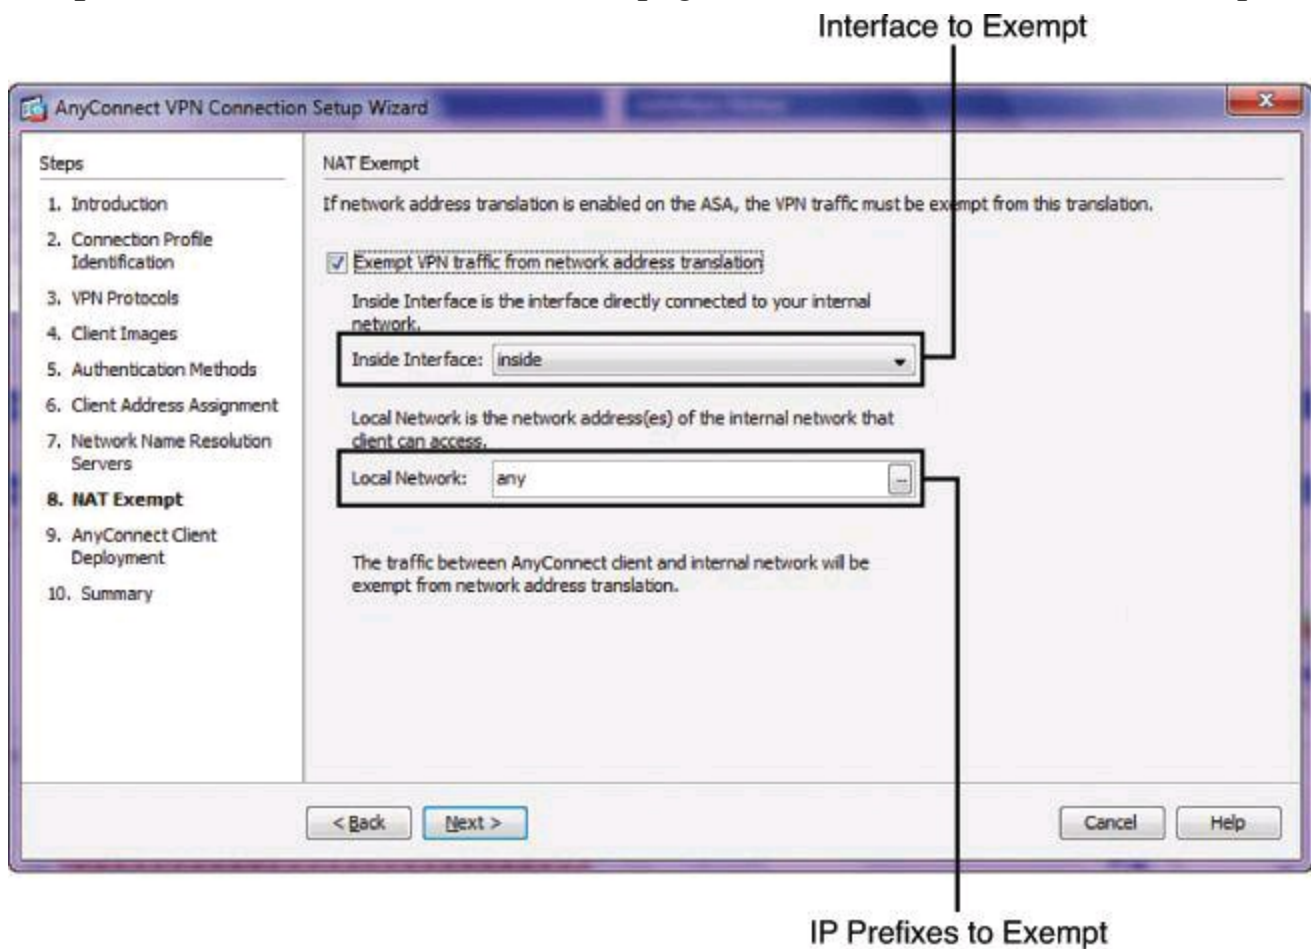


Figure 15-25. Configuring NAT Exemption

Task 8: AnyConnect Client Deployment Summary

The final two wizard pages are informational:

- **AnyConnect Client Deployment:** The message explains how you can install the Cisco AnyConnect VPN Client program to a client device using either of two methods. The first method is web launch, which installs automatically when accessing the ASA using a web browser. The second method is predeployment, which manually installs the Cisco AnyConnect VPN Client package.
- **Summary:** Provides a summary, shown in [Figure 15-26](#), of your selections from the previous wizard pages.

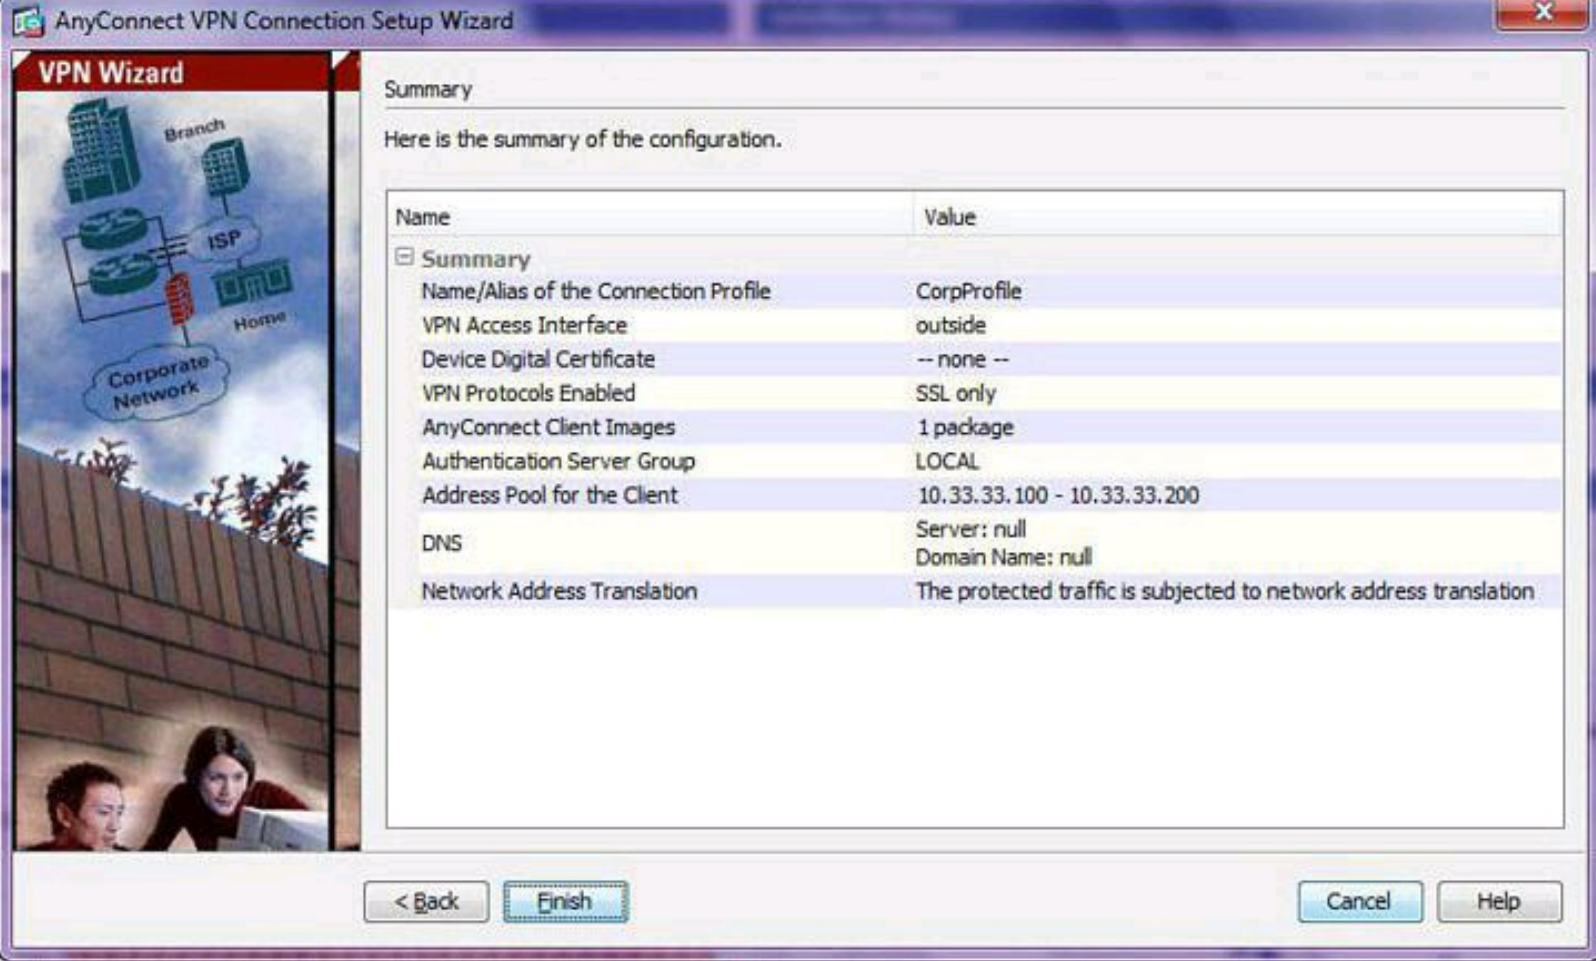


Figure 15-26. AnyConnect Wizard Summary

Click **Finish** to complete the SSL VPN Connection Setup Wizard.

Phase 2: Configure the Cisco AnyConnect VPN Client

To download the Cisco AnyConnect SSL VPN Client to the host system, you must complete the following steps:

Step 1. Open a compliant web browser and enter the login URL for the SSL VPN into the address field, as shown in [Figure 15-27](#).

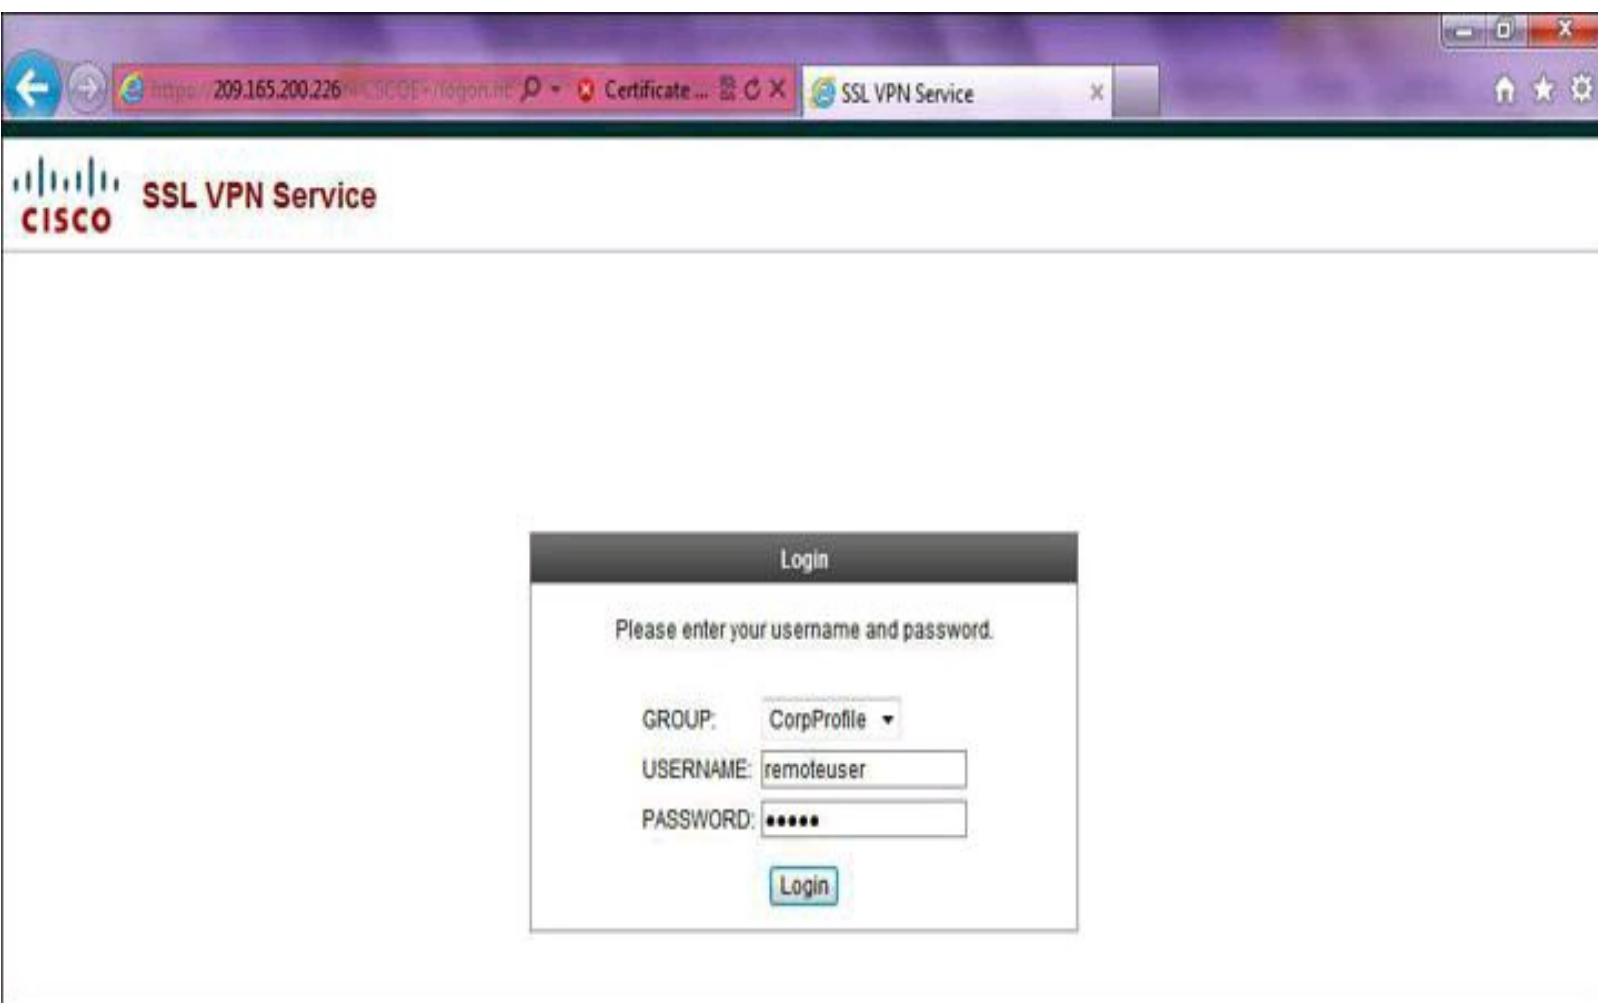


Figure 15-27. Connecting to the Portal to Eventually Request an AnyConnect Installation Download

Step 2. Ensure that the connection profile created for the client is selected in the Group drop-down menu, enter the previously configured username and password, and click **Login**.

Installing AnyConnect

AnyConnect can be preinstalled on the remote client using the standalone installation program that Cisco offers. Another option is to have the AnyConnect installation downloaded to the remote-access computer from the ASA using the AnyConnect version available in flash, as shown earlier in [Figure 15-21](#). This is the approach we will take in our scenario. Other variants of this option exist, such as having the AnyConnect installation pushed to your remote client as soon as the ASA receives an authenticated HTTPS session. Those variants are explained in detail in the Cisco Press book *CCNP Security VPN 642-648 Official Cert Guide*, Second Edition.

In our scenario, start with a clientless session but then choose to upgrade to a full client session by clicking the AnyConnect option on the left side of the portal page. A process will begin a series of compliance checks for the target system. The following items are checked on the host system:

- **Platform Detection:** The ASA first queries the client system in an attempt to identify the type of client connecting to the security appliance. Based on the platform that is identified, the proper software package may be auto-downloaded.

- **ActiveX:** Detects whether ActiveX is available on the host system for client download. For ActiveX to operate properly with the Cisco ASA, it is important that the security appliance is added as a trusted network site. ActiveX will be used for client download in the event that a web portal is not in use.
- **Java Detection:** Detects whether a supported version of Java is available on the host system for client download. Java will be used for client download in the event that a web portal page has been configured.

If all the preceding checks succeed, Cisco AnyConnect will be downloaded and installed automatically on your remote system. All along, the progress is shown, under the title WebLaunch, for the system inspection and the status of the download. Once the client completes the auto-download of the Cisco AnyConnect SSL VPN Client, the web session automatically launches the Cisco AnyConnect SSL VPN Client and attempts to log the user into the network using the same credentials that are supplied when logging into the web portal.

If the ActiveX or Java checks are not successful, the ASA will nevertheless offer you the chance to download manually the Cisco AnyConnect VPN Client located in its flash, as shown in [Figure 15-28](#).

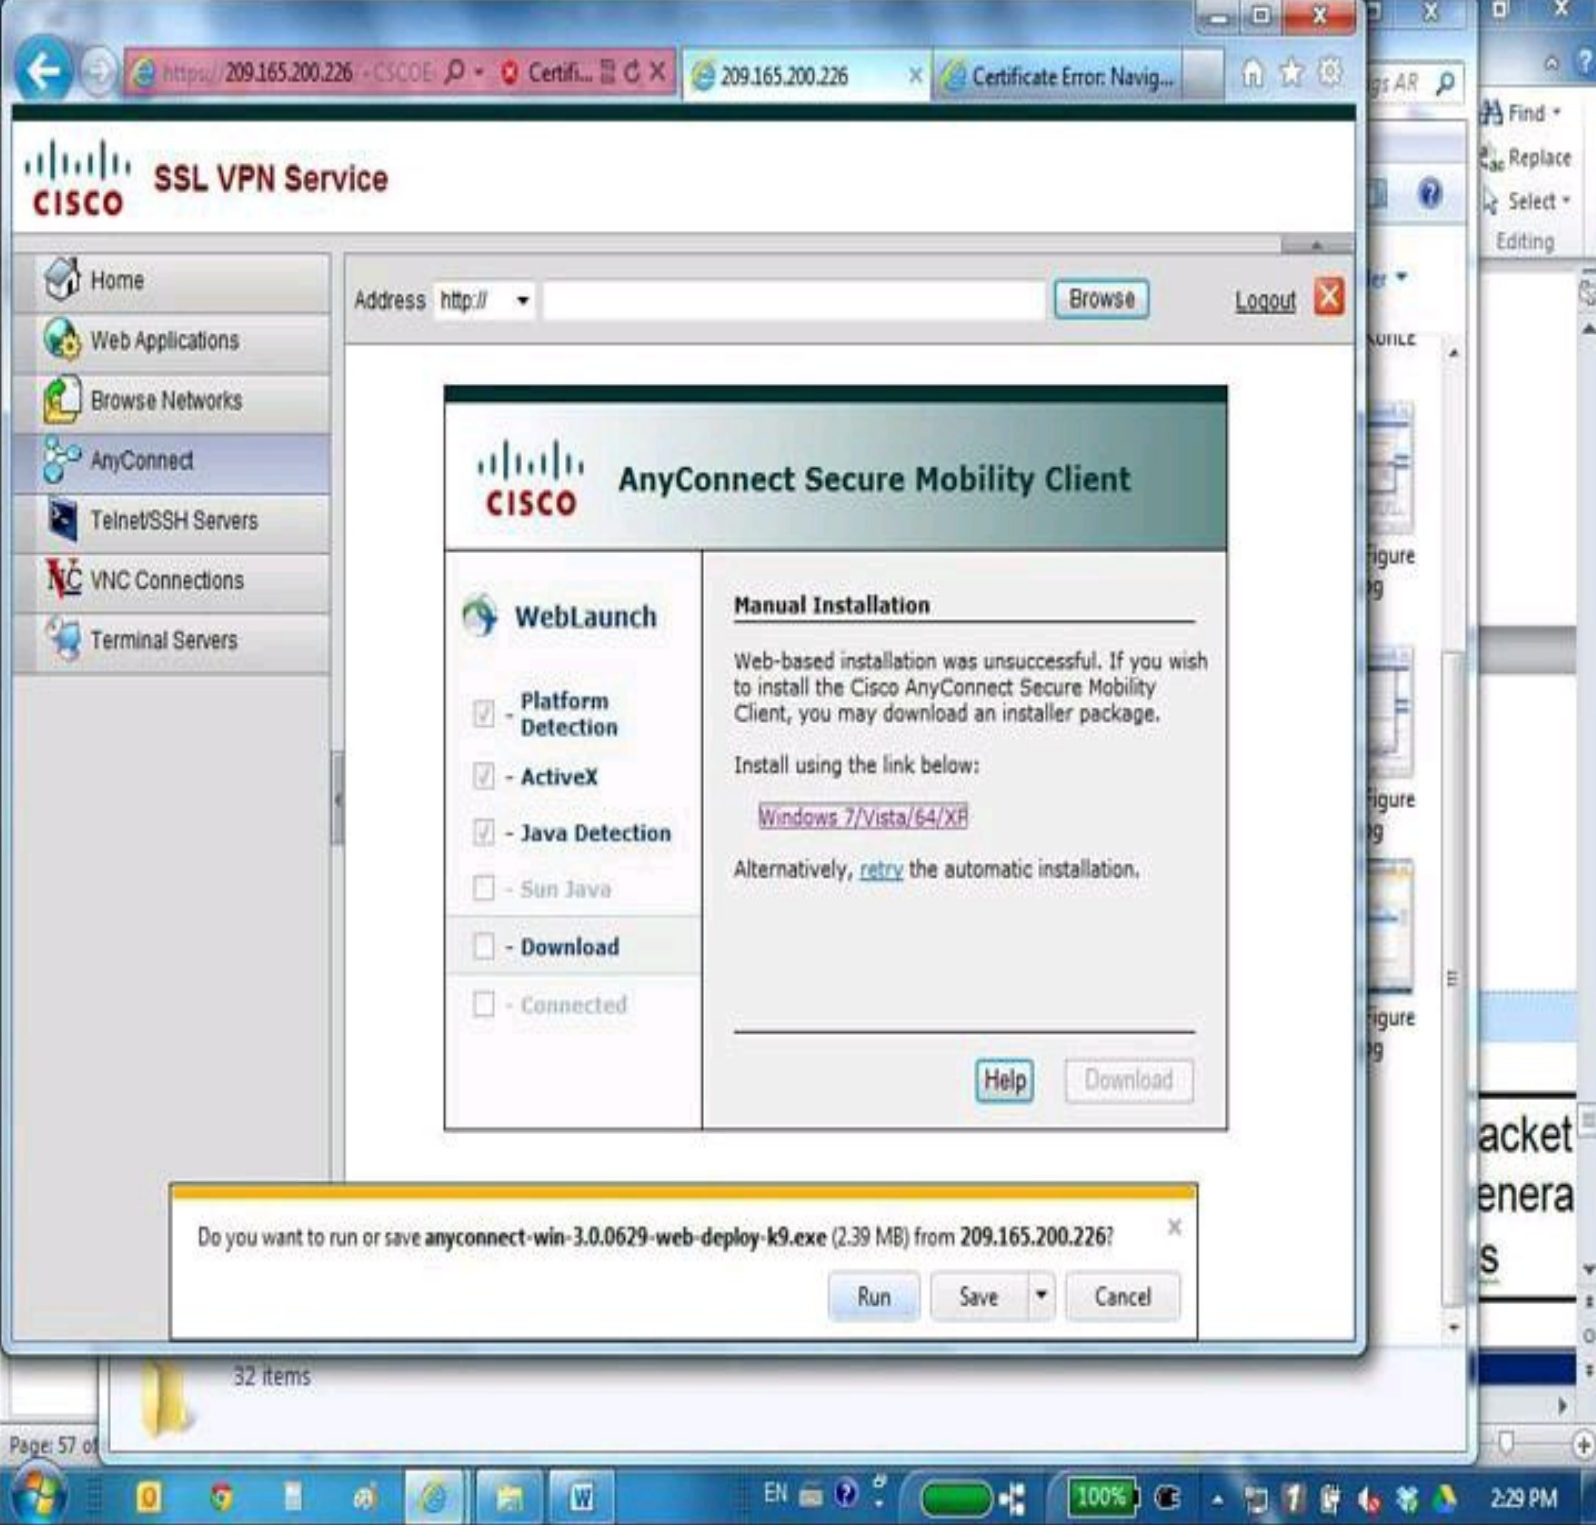


Figure 15-28. Cisco AnyConnect Installed from a VPN Clientless Session

After the download, the Cisco AnyConnect VPN Client is installed. You can then proceed to launch AnyConnect, as shown in [Figure 15-29](#), which will initiate the SSL full client VPN connection to your ASA. You will have concurrently both a clientless session and a full SSL session on the ASA.

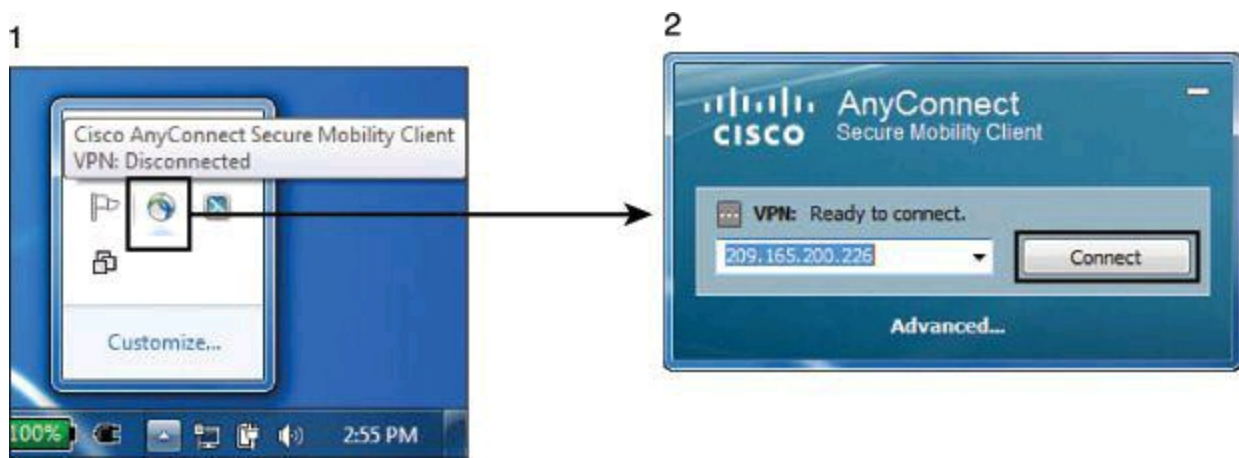


Figure 15-29. Starting Cisco AnyConnect Client

When the AnyConnect connection is established, an icon with a sphere and a lock will appear in the system tray, identifying that the client has successfully connected to the SSL VPN network, as shown in Action 1 of [Figure 15-30](#).

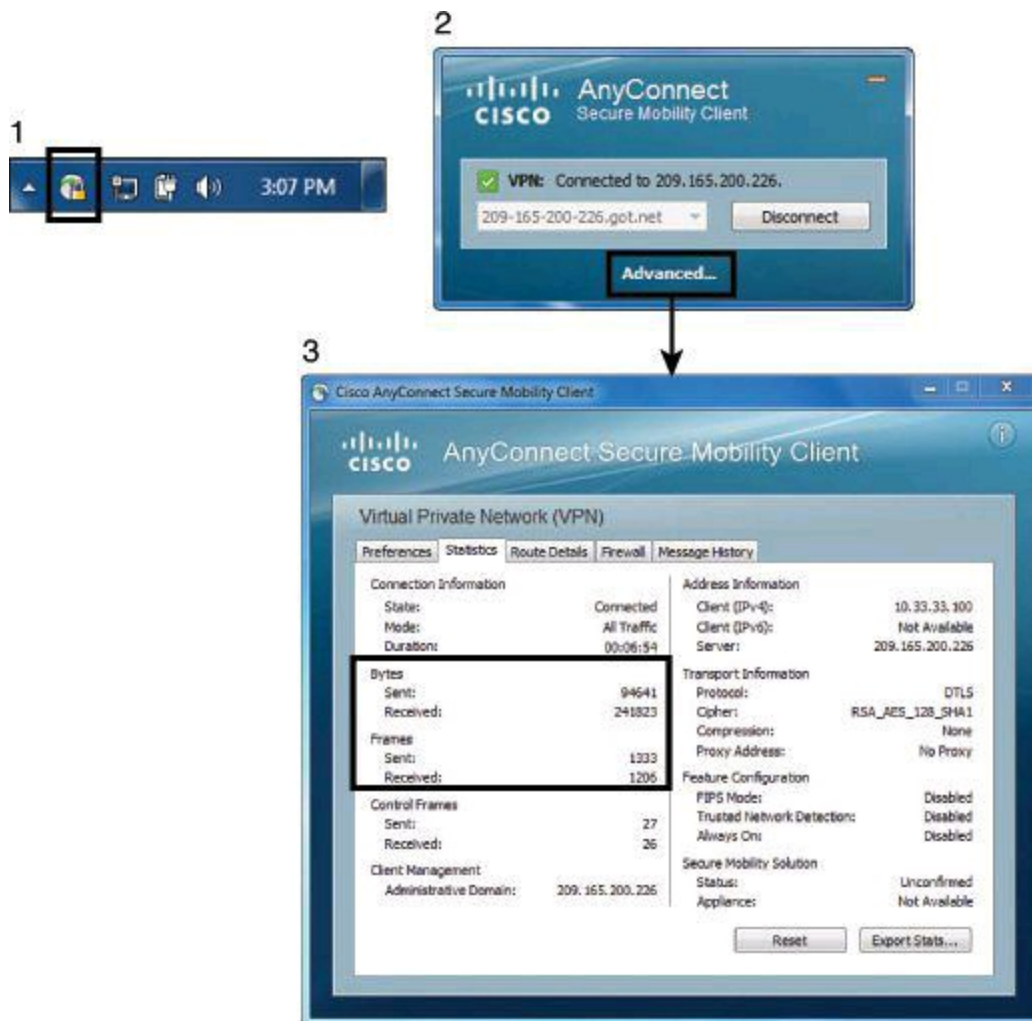


Figure 15-30. AnyConnect Icon in Computer System Tray and Details on the Connection

Phase 3: Verify VPN Connectivity with Cisco AnyConnect VPN Client

Additional connection statistics and information may be shown by clicking the icon in the system tray shown in [Figure 15-30](#), Action 1, and then selecting Advanced, as shown in [Figure 15-30](#), Action 2. This client interface may also be used to log the user out.

The Statistics tab on the Client interface presents more information, and allows you to review packet and byte counters, shown in Action 3, as well as the overall settings of the VPN (crypto suite, IP address obtained, split tunneling policy, and others).

In our scenario, future SSL VPN sessions may be launched through the web portal or through the installed Cisco AnyConnect SSL VPN Client.

Verifying VPN Connectivity from Cisco ASA

In Cisco ASDM, you can monitor established VPN sessions by navigating to **Monitoring > VPN > VPN Statistics > Sessions**. From the window, shown in [Figure 15-31](#), you can view session statistics for the ASA.

Filter by VPN Type (AnyConnect Client)

Type	Active	Cumulative	Peak Concurrent	Inactive
AnyConnect Client	1	1	7	1
SSL/TLS/DTLS	1	1	7	1
Clientless VPN	0	0	16	2
Browser	0	0	16	2

Filter By: AnyConnect Client -- All Sessions -- Filter

Username IP Address	Group Policy Connection Profile	Protocol Encryption	Login Time Duration	Bytes Tx Bytes Rx
remoteuser 10.33.33.100	GroupPolicy_CorpProfile CorpProfile	AnyConnect-Parent SSL-... RC4 AES 128	15:03:30 UTC Sat Aug 18 2012 0h:17m:41s	834070 146874

Logout By: -- All Sessions --
-- All Sessions --
Username
IP Address

Logout Sessions Refresh

Last Updated: 8/18/12 3:25:26 PM

Possibility for Bulk Logouts

Figure 15-31. Monitoring the VPN Sessions Status from ASDM

The top pane shows a summary of active, cumulative, peak concurrent, and inactive VPN sessions. This pane includes VPN sessions of all types, including IPsec, clientless SSL, and site-to-site sessions.

The bottom pane displays individual VPN sessions. You can filter the type of session to display by using the Filter By drop-down menu above the pane. This menu allows you to specify the type of sessions that the statistics in the bottom pane represent. You will need to click Filter to refresh the output according to the filter.

This bottom pane includes the username and IP addresses used in each session, as well as the connection profiles that were matched, the cipher suite in use for the session, and other session statistics such as duration and packet/byte counters.

Clicking Details to the right of the bottom pane provides additional details for each VPN session, as shown in [Figure 15-32](#).

The screenshot shows a 'Session Details' window with two main sections. The top section is a summary table with columns: Username, IP Address, Group Policy, Connection Profile, Protocol, Encryption, Login Time, Duration, Bytes Tx, and Bytes Rx. The bottom section is a 'Details' pane with tabs for 'Details' and 'ACL'. It contains a table with columns: ID, Type, Local Addr. / Subnet Mask / Protocol / Port, Remote Addr. / Subnet Mask / Protocol / Port, Encryption, Other, Bytes Tx, and Bytes Rx.

Username	IP Address	Group Policy	Connection Profile	Protocol	Encryption	Login Time	Duration	Bytes Tx	Bytes Rx
remoteuser	10.33.33.100	GroupPolicy_CorpProfile	CorpProfile	AnyConnect-Parent SSL-...	AES128	15:03:30 UTC Sat Aug 18 2012	0h:18m:27s	874171	149794

ID	Type	Local Addr. / Subnet Mask / Protocol / Port	Remote Addr. / Subnet Mask / Protocol / Port	Encryption	Other	Bytes Tx	Bytes Rx
	AnyConne...			none	Tunnel ID: 44.1 Public IP: 209.165.202.131 TCP Src Port 1704 TCP Dst Port 443 Authentication Mode: userPassword Idle Time Out: 30 Minutes Idle TO Left: 11 Minutes Client Type: AnyConnect Client Ver: 3.0.0629	4861	628

Figure 15-32. Detailed Information on Current VPN Session

Summary

The key points covered in this chapter are as follows:

- Market trends drive the need for effective remote-access security and present challenges to the IT organization.
- The SSL protocol uses the cryptology concepts presented in this chapter.
- Cisco SSL VPN solutions include clientless and full client tunnel modes of operation.
- Cisco SSL VPN clientless mode can be configured on Cisco ASA using Cisco ASDM.
- Cisco SSL VPN full client tunnel mode can be configured on Cisco ASA using Cisco ASDM and the Cisco AnyConnect VPN Client.

References

For additional information, refer to this resource:

CCNP Security VPN 642-648 Official Cert Guide, Second Edition (Cisco Press)

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in the Appendix, "[Answers to Chapter Review Questions](#)."

1. Which two current market trends create challenges to remote connectivity and Cisco SSL

VPN services?

- a. Consumerization
 - b. Bandwidth availability
 - c. Ubiquitous access to corporate resources
 - d. Increase in malware
 - e. Cost of VPN infrastructures
- 2.** Which option is the main difference between TLS and SSL?
- a. SSL operates at the transport layer.
 - b. TLS is standards-based.
 - c. SSL uses digital certificates.
 - d. TLS supports Suite B algorithms.
- 3.** When deploying Cisco SSL VPNs, which of the following options provides the most depth in application support?
- a. EasyVPN
 - b. Thin Client
 - c. Full Client
 - d. Clientless
- 4.** Which solution would you recommend for a user who requires access strictly to web-based applications?
- a. EasyVPN
 - b. Thin Client
 - c. Full Client
 - d. Clientless
- 5.** You wish to use the SSL VPN to access your network. However, the administrator has not installed a third-party certificate for Cisco SSL VPN sessions. What will happen to your SSL VPN session?
- a. The SSL VPN session will fail because it requires a server-side certificate.
 - b. The ASA will present a self-signed certificate.
 - c. A root certificate will need to be imported into ASA.
 - d. The SSL VPN will fail because the client automatically rejects untrusted certificates.
- 6.** Clientless SSL VPN is best suited for?
- a. Complex applications
 - b. Users who needs to print on a LAN printer
 - c. Web-enabled applications, file sharing, email
 - d. All IP-based applications
- 7.** Match the cryptographic algorithms with their function within SSL. Multiple algorithms

can be matched to one single function.

Function

- a.** Confidentiality
- b.** Integrity
- c.** Authentication
- d.** Key management

Algorithm

- 1.** RC4
- 2.** DSA
- 3.** SHA-2
- 4.** 3DES
- 5.** MD5
- 6.** DSS
- 7.** AES
- 8.** DH

Appendix A. Answers to Chapter Review Questions

Chapter 1

1. a, b, and e
2. a, d, and e
3. a. Preventative = f. Lock
b. Deterrent = e. Video surveillance
c. Detective = d. Motion sensor
4. a. White hat = p. Breaks security for nonmalicious reasons
b. Black hat = k. Unethical hacker
c. Gray hat = m. Ethically questionable hacker
d. Blue hat = i. Bug tester
e. Cracker = o. Synonymous with black hat hacker
f. Phreaker = l. Hacker of telecommunication systems
g. Script kiddy = j. Hacker with little skill
h. Hacktivist = n. Hacker with a political agenda
5. a. Escalate privilege = Step 4
b. Leverage the compromised system = Step 7
c. Perform footprint analysis = Step 1
d. Install back doors = Step 6
e. Enumerate applications and operating systems = Step 2
f. Gather additional passwords and secrets = Step 5
g. Manipulate users to gain access = Step 3
6. d, e, and g
7. a. Confidential = 4
b. Private = 3
c. Public = 1
d. Sensitive = 2
8. d
9. a. Owner = f. Ultimately responsible for the data
b. Custodian = e. Responsible on a day-to-day basis for the classified data
c. User = d. Responsible for using the data
10. b
11. b
12. a. Searching a network host and open ports = j. Port scanning
b. Capturing electrical transmission = g. Emanation capturing

- c. Hiding information within a transmission = h. Covert channel
- d. Intercepting traffic that passes over a physical network = e. Packet sniffing

- 13.** a. Operations and maintenance = Step 4
- b. Initiation = Step 1
 - c. Disposition = Step 5
 - d. Acquisition and development = Step 2
 - e. Implementation = Step 3

14. b

Chapter 2

- 1.** a. Context-aware enforcement = 4
- b. Cisco TrustSec = 1
 - c. Cisco SIO = 2
 - d. Cisco AnyConnect = 3

2. a

3. b and d

4. c

5. b and e

Chapter 3

1. a

2. c

3. a and d

4. c and e

- 5.** a. Management = f. Packets used to manage the network
- b. Data = d. User-generated packets
 - c. Control = e. Packets used for the creation and operation of the network itself

6. d and f

7. b

- 8.** a. Communities = g. Groups of devices that share common components
- b. Templates = e. Parameterized configuration files
 - c. Profiles = f. GUI views that allow role-based access control over Cisco Configuration Professional menus and options
 - d. Wizards = h. GUI tools to hide the complexity of commands

9. a

10. c

Chapter 4

1. c
2. a and d
3. d
4. a
5. c
6. d
7. a
8. b
9. c
10. b

Chapter 5

1. d
2. b, c, and d
3. a and c
4. b
5. c
6. a. = Step 3
b. = Step 4
c. = Step 1
d. = Step 2
7. c
8. a
9. b
10. b

Chapter 6

1. c
2. a and c
3. a
4. c
5. e
6. b
7. d
8. a and c
9. d
10. c

Chapter 7

- [1.](#) c
- [2.](#) b
- [3.](#) a and d
- [4.](#) c
- [5.](#) c

Chapter 8

- [1.](#) a
- [2.](#) b and c
- [3.](#) d
- [4.](#) d
- [5.](#) d
- [6.](#) a, b, and c
- [7.](#) c
- [8.](#) b, c, and d
- [9.](#) a
- [10.](#) a, b, and c

Chapter 9

- [1.](#) a and c
- [2.](#) c
- [3.](#) a. Packet-filtering firewalls = 1 Work primarily at the network level of the OSI model
b. Application layer gateways = 4 Were the first application layer firewalls
c. Stateful packet filters = 2 Are the most common firewalls
d. Application inspection firewalls = 3 Monitor sessions to determine the port numbers for secondary channels
- [4.](#) c and d
- [5.](#) b
- [6.](#) a. Static NAT = 3 Translation is one-to-one
b. Dynamic NAT = 2 Translation is many-to-many
c. Dynamic PAT = 4 Translation is many-to-one
d. Policy NAT = 1 Translation depends on both source and destination
- [7.](#) a. Service control = 4 Allow HTTP, allow HTTPS, deny everything else
b. Direction control = 1 Allow HTTP outbound, but not inbound
c. User control = 2 Allow campus VLANs HTTP access, deny it for wireless VLANs
d. Behavior control = 3 Open negotiated FTP ports after learning them during connection

Chapter 10

- [1.](#) b
- [2.](#) b
- [3.](#) d
- [4.](#) c
- [5.](#) a and d
- [6.](#) b
- [7.](#) b
- [8.](#) b
- [9.](#) d
- [10.](#) b
- [11.](#) d
- [12.](#) a
- [13.](#) c
- [14.](#) a
- [15.](#) b

Chapter 11

- [1.](#) b and d
- [2.](#) b
- [3.](#) a and d
- [4.](#) c
- [5.](#) a, b, and d
- [6.](#) a. Signature-based IPS = 3. Can produce false positives because certain normal network activity can be misinterpreted as malicious activity
 - b. Policy-based IPS = 4. Similar to implementing a restrictive firewall policy
 - c. Reputation-based IPS = 2. Typically implemented in the form of white lists or black lists
 - d. Anomaly-based IPS = 1. Normal behavior typically defined based on traffic patterns, traffic and protocol mix, traffic volumes, and other criteria
- [7.](#) d
- [8.](#) c
- [9.](#) b
- [10.](#) a
- [11.](#) a and d
- [12.](#) a

13. a

14. c

15. a

Chapter 12

1. a, d, and e

2. b

3. a, c, and d

4. b

5. c

6. d

7. a and e

8. b

9. d

10. b

11. a and c

12. d

13. d

14. a

15. b

16. c

17. a

18. d

19. c

20. c

21. c

22. d and e

Chapter 13

1. b

2. a, b, and c

3. a, d, and e

4. a. Group 5 = g. 4096

b. Group 2 = h. 163

c. Group 7 = i. 1024

d. Group 1 = j. 1536

e. Group 16 = k. 256

f. Group 19= 1. 768

5. b

6. a. ESP = 4. Confidentiality

b. IKE = 3. Negotiation

c. EDCH = 1. Key Exchange

d. EDCSA = 2. Authentication

7. b

8. b

9. b

10. a, e, and f

Chapter 14

1. d

2. c

3. d

4. d

5. b

6. a

7. b

8. c

Chapter 15

1. a and c

2. b

3. c

4. d

5. b

6. c

7. a. Confidentiality = 1, 4, 7

b. Integrity = 3, 5

c. Authentication = 2

d. Key management = 6, 8

Index

Numerics

3DES, [562-563](#)

6to4 tunneling, [284](#)

802.1Q tagging, [236-237](#)

configuring, [238-239](#)

2010/2011 CSI Computer Crime and Security Survey, [1](#)

A

AAA (authentication, authorization, and accounting), [186-205](#)

accounting policies, [213](#)

authentication, configuring on Cisco routers, [188-189](#)

authorization, configuring on Cisco routers, [190](#)

Cisco Secure ACS configuration, [198-205](#)

deploying, [127-128](#)

external database configuration, [208-214](#)

local database configuration, [191-198](#)

RADIUS, [206](#)

TACACS+, [205-206](#)

troubleshooting, [216-218](#)

accepting risk, [12](#)

access control, configuring outbound access control on Cisco ASA, [446-460](#)

access ports, configuring port security, [264-269](#)

accounting policies (AAA), [213](#)

ACL Editor, [349](#)

ACLs, [320-325](#)

configuring with CCP, [347-358](#)

developing, best practices, [345-347](#)

displaying, [342-343](#)

extended ACLs, configuring, [338-342](#)

filtering on Cisco NFP data plane, [128-129](#)

implicit deny any statement, [324](#)

in IPv6 environment, syntax, [362-363](#)

IPv6-based attacks, mitigating, [360](#)

mirrored crypto ACLs, [644](#)

monitoring with CCP, [356](#)

named ACLs, [324](#)

numbered ACLs, [324](#)

object groups, [343-345](#)

configuring with CCP, [357-358](#)

removing, [325](#)

standard ACLs

configuring, [335-337](#)

in IPv6 environment, [361](#)

traffic, controlling, [335-342](#)

wildcard bits, [331-334](#)

acquisition and development phase (SDLC), [65-66](#)

actions, [404, 407-408](#)

adding rules in CCP, [350-351](#)

addressing, IPv6, [286-292](#)

address representation, [285-286](#)

global unicast addressing, [287-288](#)

link-local addressing, [288-289](#)

multicast addressing, [289-290](#)

unicast addressing, [286-287](#)

administrative controls, [9](#)

administrative laws, [76](#)

adversaries, identifying, [20-21](#)

AES, [563-564](#)

age, classifying data by, [6](#)

AH (Authentication Header), [618](#)

alarms, [474-475](#)

IPS, [484-486](#)

ALE (annualized loss expectancy), calculating, [46-47](#)

anomaly-based IPS systems, [477-478](#)

anti-evasion techniques (IPS), [480-481](#)

antispoofing on Cisco NFP data plane, [129](#)

anycast addressing, [287-288](#)

application awareness, [313](#)

application border, [86](#)

application inspection firewalls, [382-383](#)

application layer firewalls, [374-378](#)

application layer, security controls, [309](#)

application-specific gateways, [313](#)

archiving, [485](#)

ARO (annualized rate of occurrence), calculating, [45](#)

assessing

liability, [76](#)

network security posture, [69-70](#)

assets, [3](#)

identifying, [53-54](#)

assigning VLANs to switch ports, [242-243](#)

associating CCP rules with interfaces, [352-353](#)

assumptions made regarding networks, [2](#)

asymmetric encryption algorithms, [565-567](#)

digital signatures, [583-587](#)

atomic signature engines, [501](#)

attacks

adversaries, identifying, [20-21](#)

availability attacks, [36](#)

back doors, [22](#)

blended threats, [39](#)

botnets, [37](#)

confidentiality breaches, [30-32](#)

covert channels, [33-35](#)

for cryptanalysis, [551-555](#)

DDoS, [15](#), [37-39](#)

DoS attacks, [37-39](#)

enumeration, [25](#)

fingerprinting, [25](#)

hackers, methodologies, [21-23](#)

hardware hacking, [16](#)

identity theft, [35-36](#)

IP spoofing, [25-27](#)

sequence prediction, [27](#)

IPS responses, [478](#)

IPv6 and IPv4, comparing, [296-298](#)

IPv6-based, [16](#)

mitigating with ACLs, [360](#)

Layer 2 protection, [250](#)

CAM table overflow attacks, mitigating, [259-260](#)

MAC address spoofing, mitigating, [260-261](#)

STP attacks, mitigating, [254-258](#)

VLAN hopping attacks, mitigating, [251-254](#)

man-in-the-middle attacks, [32-33](#)

memory scraping, [15](#)

motivation for, [13-14](#)

overt channels, [33-35](#)

password attacks, [28-30](#)

hashes, [29](#)

pharming, [35-36](#)

phishing attacks, [35-36](#)

social engineering, [22](#)

threats classification, [23-25](#)

Trojan horse attacks, [22](#)

trust exploitation, [28](#)

on websites, [15](#)

audience for security policies, [54-55](#)

authentication, [615-616](#)

of router access using AAA, [188-189](#)

authorization

in Cisco Secure ACS, [219-221](#)

configuring on Cisco routers, [190](#)

AV (asset value), [45](#)

availability attacks, [36](#)

availability of systems and data, [3](#)

awareness

application awareness, [313](#)

security awareness, [62-63](#)

B

back doors, [22](#)

IP spoofing, [25-27](#)

BackTrack 5, downloading, [72](#)

bastion hosts, [370](#)

best practices

for ACL development, [345-347](#)

for IPS systems, [492-494](#)

IPv6 networks, protecting, [300](#)

BID (bridge ID), [247](#)

birthday attacks, [553](#)

blended threats, [15](#), [39](#)

blind attacks, [32](#)

block ciphers, [547](#)

borderless data center component (Cisco Borderless Network Security Architecture), [90-91](#)

borderless end zone component (Cisco Borderless Network Security Architecture), [88-89](#)

borderless Internet component (Cisco Borderless Network Security Architecture), [89-90](#)

borderless networks, SSL VPNs, [670-672](#)

botnets, [37](#)

BPDU (bridge protocol data unit), [247](#)

BPDU Guard, [256-257](#)

breaches in confidentiality, [30-32](#)

broadcast storms, [245](#)

brute-force attacks, [29](#), [552](#)

building blocks of information security

security assumptions, [2](#)

security requirements, [2-3](#)

business continuity planning, [78](#)

business goals, need for network security, [12](#)

C

calculating

ALE, [46-47](#)

ARO, [45](#)

risk matrix, [48-49](#)

SLE, [45](#)

CAM table overflow attacks, mitigating, [259-260](#)

CAPEC (Common Attack Pattern Enumeration and Classification), [24](#)

CAs (certificate authorities), [590-593](#)

categorization of vulnerabilities, [47](#)

CCP (Cisco Configuration Professional), [131-142](#)

ACL Editor, [349](#)

ACLs

configuring, [347-349](#)

monitoring, [356](#)

Cisco AutoSecure features, [152-154](#)

Cisco IOS IPS, configuring, [507-524](#)

communities, [142-144](#)

creating, [143-144](#)

managing, [144](#)

content pane, [142](#)

initial configuration, [133-134](#)

logging, [354-355](#)

menu bar, [136](#)

navigation pane

Device Hardening folder, [140](#)

Firewall option, [141](#)

Interface Management option, [138-139](#)

IPS option, [141](#)

Router folder, [139-140](#)

Security folder, [140](#)

object groups, configuring, [357-358](#)

rules, [347-348](#)

adding, [350-351](#)

associating with interfaces, [352-353](#)

Security Audit feature, device hardening, [149-151](#)

site-to-site VPNs, configuring, [647-657](#)

SNMP options, enabling, [185](#)

status bar, [142](#)

templates, [145-146](#)

toolbar, [138](#)

user profiles, [147-148](#)

certificates, [590](#)

revocation methods, [599-600](#)

changes in workplace experience, effect on network security, [16](#)

Chapman, David, [393](#)

chosen-cyphertext attacks, [553](#)

chosen-plaintext attacks, [553](#)

CIA triad, [2-3](#)

ciphers, [540-549](#)

block and stream ciphers, [547-549](#)

one-time pad ciphers, [543-546](#)

polyalphabetic ciphers, [541](#)

substitution ciphers, [541](#)

ciphertext-only attacks, [552](#)

Cisco AnyConnect client, [97-98](#), [683](#)

installing, [702](#)

SSL VPN full-tunnel mode, configuring with ASDM, [692-707](#)

Cisco ASA 5500-X Series Firewalls, [491](#)

Cisco ASA (Adaptive Security Appliance), [427-460](#)

5500 Series, network services, [428-431](#)

Cisco ASDM, [436-442](#)

features, [438-442](#)

CLI, [434-435](#)

connection table, [430](#)

MPF, [443-446](#)

outbound access control, [446-460](#)

security levels, [432](#)

technologies used, [431-432](#)

Cisco ASDM, [436-442](#)

clientless SSL VPN, configuring, [683-691](#)

features, [438-442](#)

outbound access control on Cisco ASA, [446-460](#)

Cisco AutoSecure, [122](#), [152-154](#)

Cisco Borderless Network Security Architecture

borderless data center component, [90-91](#)

borderless end zone component, [88-89](#)

borderless Internet component, [89-90](#)

Cisco Borderless Network Services, [91-92](#)

Cisco SecureX, [93-98](#)

Cisco AnyConnect client, [97-98](#)

context-awareness, [94](#)

ISE, [98](#)

SIO, [94-95](#), [99-100](#)

TrustSec, [95-96](#)

Cisco Security Manager, [107-108](#)

cloud security, [100-101](#)

content security, [101](#)

data loss prevention, [101](#)

email security, [104-105](#)

policy management layer component, [91](#)

threat control and containment solutions, [98-99](#)

web security, [101-105](#)

Cisco Borderless Network Services, [91-92](#)

Cisco Common Classification Policy Language, Zone-Based Policy Firewall

actions, [407-408](#)

class maps, [405-406](#)

interzone policies, configuring, [411-422](#)

NAT services, configuring, [422-426](#)

policies, creating, policy maps, [405](#)

traffic flows, [409-410](#)

Cisco IOS devices

configuration files, [168-169](#)

multiple privilege levels, configuring, [170-171](#)

NTP, configuring, [177](#)

passwords

configuring, [163-166](#)

minimum length, setting, [165](#)

routers, configuring AAA, [186-205](#)

syslog, configuring, [178-182](#)

Cisco IOS IPS

configuring

with CCP, [507-524](#)

with CLI, [524-529](#)

features, [495-497](#)

signatures, [498-505](#)

managing, [500](#)

SMEs, [500-501](#)

tuning, [502-507](#)

Cisco IOS routers, configuring IPv6, [293-295](#)

Cisco IOS Zone-Based Policy Firewall. *See* [Zone-Based Policy Firewall](#)

Cisco IPS 4300 Series Sensors, [491](#)

Cisco IronPort, [104-105](#)

Cisco MPF (Modular Policy Framework), [443-446](#)

Cisco NFP (Network Foundation Protection), [112-118](#)

control plane, [118-122](#)

CoPP, [118-119](#)

CPPr, [119](#)

data plane, [128-131](#)

ACL filtering, [128-129](#)

antispoofing, [129-130](#)

Layer 2 protection, [131](#)

management plane, [123-128](#)

AAA, deploying, [127-128](#)

RBAC, [126-127](#)

secure management and reporting, [124-126](#)

Cisco routers

ACLs, [320-325](#)

displaying, [342-343](#)

implicit deny any statement, [324](#)

named ACLs, [324](#)

numbered ACLs, [324](#)

object groups, [343-345](#)

removing, [325](#)

wildcard bits, [331-334](#)

vty access, configuring, [338](#)

Cisco SAFE Blueprint, [41](#)

Cisco Secure ACS

AAA, configuring, [198-205](#)

authorization, [219-221](#)

configuring, [224-228](#)

non-Windows version, [203-204](#)

rule-based policies, [222](#)

for Windows, [201-202](#)

Cisco SecureX, [93-98](#)

Cisco AnyConnect client, [97-98](#)

context-awareness, [94](#)

ISE, [98](#)

SIO, [94-95](#)

TrustSec, [95-96](#)

Cisco Security Manager, [107-108](#)

Cisco SIO (Security Intelligence Operation), [99-100](#), [313-314](#)

Cisco threat control and containment strategies, [311-317](#)

fundamentals, [314-317](#)

technologies involved in, [312-314](#)

Cisco TrustSec, [95-96](#)

civil laws, [75](#)

Class C IP addresses, subnetting, [326-327](#)

class maps, [405-406](#)

classes, [403](#)

classful prefix length, [326](#)

classification

countermeasures classification, [8-9](#)

administrative controls, [9](#)

physical controls, [10](#)

technical controls, [9-10](#)

data classification, [4-7](#)

government classification schemes, [4-5](#)

private sector data classification, [5](#)

roles filled within, [7](#)

threat classification, [23-25](#)

vulnerabilities classification, [7-8](#)

CLI (command-line interface), Cisco ASA, [434-435](#)

clientless SSL VPN, configuring with ASDM, [683-691](#)

clientless VPNs, [672](#)

cloud computing

data location, [1](#)

effect on network security, [16](#)

security, [100-101](#)

COBIT (Control Objectives for Information and Related Technologies), [68](#)

collecting evidence, [75](#)

collisions, [568-569](#)

communities (CCP), [142-144](#)

creating, [143-144](#)

managing, [144](#)

comparing

fail-open and fail-close approaches, [493](#)

IKEv1 and IKEv2, [633-635](#)

IPS and IDS systems, [468-471](#)

IPv4 and IPv6 attacks, [296-298](#)

RADIUS and TACACS+, [205-206](#)

reporting and monitoring, [485](#)

symmetric encryption algorithms, [560](#)

compartmentalization, [40](#)

compliance regulations, [50-53](#)

data breach disclosure, [52](#)

examples of, [51-52](#)

globalization, [52](#)

computer crime investigations, [74-75](#)

Computer Security Institute

2010/2011 CSI Computer Crime and Security Survey, [1](#)

confidential data classification, [5](#)

within public sector, [5](#)

confidentiality, [2](#)

breaches in, [30-32](#)

configuration files, securing for Cisco IOS devices, [168-169](#)

configuring

AAA

with Cisco Secure ACS, [198-205](#)

with external database, [208-214](#)

local database configuration, [191-198](#)

ACLs

with CCP, [347-349](#)

extended ACLs, [338-342](#)

standard ACLs, [335-337](#)

Cisco ASA

outbound access control, [446-460](#)

Cisco IOS devices

multiple privilege levels, [170-171](#)

NTP, [177](#)

passwords, [163-166](#)

syslog, [178-182](#)

Cisco IOS IPS

with CCP, [507-524](#)

with CLI, [524-529](#)

Cisco routers, vty access, [338](#)

Cisco Secure ACS, [224-228](#)

clientless SSL VPN with ASDM, [683-691](#)

inter-VLAN routing, [243-244](#)

IPv6 on Cisco IOS routers, [293-295](#)

port security, [261-269](#)

Role-Based CLI Access, [171-174](#)

ROM monitor, [167-168](#)

site-to-site VPNs with CCP, [647-657](#)

SSH, [161-162](#)

trunking, 802.1Q tagging, [238-239](#)

Zone-Based Policy Firewall

interzone policies, [411-422](#)

NAT services, configuring, [422-426](#)

connection table (Cisco ASA), [430](#)

connection-oriented design, [26](#)

connectivity, VPN security, [105-106](#)

consumer devices, exploits, [15](#)

content pane (CCP), [142](#)

content security with Cisco Borderless Network Security Architecture, [101](#)

context-awareness of SecureX products, [94](#)

control plane, Cisco NFP, [118-122](#)

CoPP, [118-119](#)

CPPr, [119](#)

controlling

threats, design guidelines, [308](#)

traffic with ACLs, [335-342](#)

controls

categorization of, [11](#)

COBIT, [68](#)

CoPP (Control Plane Policing), [118-119](#)

countermeasures, [3-4](#)

classification of, [8-9](#)

administrative controls, [9](#)

physical controls, [10](#)

technical controls, [9-10](#)

covert channels, [31](#), [33-35](#)

CPPr (Control Plane Protection), [119](#)

Cisco AutoSecure, [122](#)

routing protocol integrity, [121](#)

traffic classes, identifying, [120-121](#)

crackers, [20](#)

cracking passwords, [29-30](#)

creating

communities with CCP, [143-144](#)

VLANs, [240-242](#)

cribs, [553](#)

criminal laws, [75](#)

crypto maps, [646](#)

cryptoanalysis, [539](#)

attacks used for, [551-555](#)

cribs, [553](#)

cryptography, [539](#)

ciphers, [540-549](#)

encryption, [549-551](#), [612-613](#)

symmetric encryption algorithms, [558-561](#)

history of, [540](#)

Suite B cryptographic standard, [611-612](#)

cryptology, [538-539](#)

custodian role in classification system, [7](#)

CVE (Common Vulnerabilities and Exposures), [8](#)

CVSS (Common Vulnerability Scoring System), [8](#)

D

data breach disclosure, [52](#)

data classification, [4-7](#)

Cisco NFP

ACL filtering, [128-129](#)

government classification schemes, [4-5](#)

methods of classification, [6](#)

private sector data classification, [5](#)

roles filled within, [7](#)

data collection, collecting evidence, [75](#)

data integrity, [614](#)

data location, [1](#)

data loss prevention with Cisco Borderless Network Security Architecture, [101](#)

data plane, Cisco NFP, [128-131](#)

antispoofing, [129-130](#)

Layer 2 protection, [131](#)

DDoS attacks, [15](#), [37-39](#)

Stacheldracht, [39](#)

deep packet inspection, [382-383](#)

defense in depth, [39-42](#)

deploying

AAA, [127-128](#)

IPS systems, [488-491](#)

NAT, [389-390](#)

SSL VPNs, [679-683](#)

DES (Data Encryption Standard), [560-562](#)

designing secure networks

principles, [39-42](#)

defense in depth, [41-42](#)

SDLC

acquisition and development phase, [65-66](#)

disposition phase, [67](#)

implementation phase, [66-67](#)

initiation phase, [65](#)

models and frameworks, [67-68](#)

operations and maintenance phase, [67](#)

desirable encryption algorithm features, [554-555](#)

detective controls, [11](#)

deterrent controls, [11](#)

developing ACLs, best practices, [345-347](#)

device borders, [85](#)

Device Hardening option (CCP), [140](#)

devices, hardening with Security Audit Wizard, [149-151](#)

DH (Diffie-Hellman) algorithm, [579-583](#), [613-614](#)

DHCPv6, [292](#)

Diffie, Whitfield, [579](#)

digital signatures, [575-579](#), [583-587](#)

diminishing returns of security investments, [72](#)

disaster recovery, [77-78](#)

displaying ACLs, [342-343](#)

disposition phase (SDLC), [67](#)

distributed security intelligence, [309-310](#)

DMZ (demilitarized zone), [28](#)

DoS attacks, [37-39](#)

downloading BackTrack 5, [72](#)

DTLS (Datagram Transport Layer Security), [673-674](#)

dual stack, [284](#)

dumpster diving, [31](#)

Dynamic NAT, [389](#)

E

ECDSA (Elliptical Curve Digital Signature Algorithm), [615](#)

EF (exposure factor), [45](#)

egress traffic, [321](#)

email security with Cisco Borderless Network Security Architecture, [104-105](#)

emanations capturing, [31](#)

encryption, [549-551](#)

asymmetric encryption algorithms, [565-567](#)

digital signatures, [583-587](#)

desirable algorithm features, [554-555](#)

digital signatures, [575-579](#)

SEAL, [565](#)

symmetric encryption algorithms, [558-561](#)

3DES, [562-563](#)

AES, [563-564](#)

DES, [560-562](#)

end-user policies, [57](#)

Enigma machine, [540](#)

enumeration, [25](#)

errdisable recovery feature (port security), [262-263](#)

ESP (Encapsulating Security Payload), [619-620](#)

ethics, [75-76](#)

EUI-64 interface ID assignment, [291-292](#)

event monitoring, IPS systems, [485-486](#)

evidence, collecting, [75](#)

evolution of threats in information security, [306-307](#)

examples

of compliance regulations, [51-52](#)

of IPv6 attacks, [298-299](#)

of subnetting, [327-328](#)

of VLSM, [329-330](#)

exploits, [3-4](#)

on consumer devices, [15](#)

exploitation of trust, [17-18](#)

virtualization exploits, [15](#)

extended ACLs, configuring, [338-342](#)

external database configuration, AAA, [208-214](#)

external security policy audience, [55](#)

F

fail-open versus fail-close approaches, [493](#)

false negatives, [474](#)

false positives, [474](#)

features

- of Cisco ASDM, [438-442](#)
- of Cisco IOS IPS, [495-497](#)
- of desirable encryption algorithms, [554-555](#)
- of IPv6, [278-284](#)
- of Zone-Based Policy Firewall, [400-401](#)

fingerprinting, [25](#)

FIPS (Federal Information Processing Standard) 140 publication, [583](#)

Firewall option (CPP), [141](#)

firewalls

- application inspection firewalls, [382-383](#)
- application layer firewalls, [374-378](#)
- bastion hosts, [370](#)
- Cisco ASA, [427-460](#)
 - CLI, [434-435](#)*
 - security levels, [432-434](#)*
- Cisco ASA 5500-X Series Firewalls, [491](#)
- common properties of, [368-369](#)
- in layered defense strategy, [370-371](#)
 - policies, [391-392](#)*
- limitations of, [369-370](#)
- network access control, [369](#)
- rules, design guidelines, [392-394](#)
- stateful packet-filtering firewalls, [378-382](#)
- static packet-filtering firewalls, [372-374](#)
- transparent firewalls, [383-384](#)
- Zone-Based Policy Firewall, [398-403](#)
 - actions, [407-408](#)*
 - class maps, [405-406](#)*
 - features, [400-401](#)*
 - interzone policies, [399](#)*
 - NAT services, configuring, [422-426](#)*
 - policy maps, [405](#)*
 - traffic flows, [409-410](#)*
 - zones, [402-403](#)*

footprinting, [21](#)

full-tunnel mode (SSL VPN), configuring with ASDM, [692-707](#)

G

gateways, application-specific, [313](#)

gathering intelligence

distributed approach, [309-310](#)

global correlation, [486-487](#)

GetMAC, [22](#)

global correlation, [486-487](#)

global unicast addressing, [287-288](#)

EUI-64 interface ID assignment, [291-292](#)

manual interface assignment, [291](#)

globalization, effect on compliance regulations, [52](#)

governing policies, [56-57](#)

government classification schemes, [4-5](#)

GRE (generic routing encapsulation), [534](#)

guidelines, [60](#)

firewall rules design, [392-394](#)

for OOB and in-band architecture management, [176](#)

for threat control design, [308](#)

H

hackers

identifying, [20-21](#)

methodologies, [21-23](#)

tools, [21-22](#)

hacktivists, [21](#)

hardening Cisco IOS devices with Security Audit Wizard, [149-151](#)

hardware hacking, [16](#)

hashes, [29, 167](#)

hashing, [568-575](#)

collisions, [568-569](#)

HMAC, [573-575](#)

MD5, [572](#)

SHA-1, [572](#)

SHA-2, [573](#)

headers, IPv6, [279-280](#)

Hellman, Martin, [579](#)

HMAC (Hashed Message Authentication Code), [184, 573-575](#)

honeypots, [478](#)

I

ICMP sweeps, [25](#)

ICMPv6, [280-281](#)

identifying

adversaries, [20-21](#)

assets, [53-54](#)

traffic classes for CPPr, [120-121](#)

identity theft, [35-36](#)

IDS (intrusion detection systems)

comparing with IPS systems, [468-471](#)

honeypots, [478](#)

limitations of, [471](#)

sensors, alarms, [474-475](#)

ignoring risk, [12](#)

IKEv1

modes, [624](#)

phases, [625-631](#)

IKEv2, [632](#)

implementation phase (SDLC), [66-67](#)

implicit deny any statement, [324](#)

in-band management, [124](#)

guidelines, [176](#)

inbound ACLs, [323](#)

incident response, [72-73](#)

computer crime investigations, [74-75](#)

data collection, [75](#)

information security, building blocks of

security assumptions, [2](#)

security requirements, [2-3](#)

ingress traffic, [321](#)

initial configuration, CCP, [133-134](#)

initiation phase (SDLC), [65](#)

inline mode, [471](#)

inside local addresses, [385](#)

inside source address translation (NAT), [387](#)

installing Cisco AnyConnect client, [702](#)

integrity, [3](#), [614](#)

intelligence gathering, global correlation, [486-487](#)

intent evolution in risk management, [13-14](#)

interface, CCP

menu bar, [136](#)

toolbar, [138](#)

Interface Management option (CCP), [138-139](#)

internal security policy audience, [55](#)

inter-VLAN routing, configuring, [243-244](#)

interzone policies, [399](#)

configuring, [411-422](#)

investigating computer crimes, [74-75](#)

IP spoofing, [25-27](#)

sequence prediction, [27](#)

IPS (intrusion prevention systems), [468](#). See also [Cisco IOS IPS](#)

alarms, [484-485](#)

event monitoring, [485-486](#)

anomaly-based, [477-478](#)

anti-evasion techniques, [480-481](#)

architectures, [494](#)

attack responses, [478](#)

best practices, [492-494](#)

comparing with IDS systems, [468-471](#)

deploying, [488](#)

fail-open versus fail-close approaches, [493](#)

global correlation, [486-487](#)

IPv6-aware, [484](#)

management consoles, [471](#)

policy-based, [477](#)

reputation-based, [478](#)

risk-based, [481-483](#)

RR, [484](#)

sensors, promiscuous mode, [471](#)

signature-based, [475-477](#)

technologies used, [475-476](#)

TR, [484](#)

IPS option (CCP), [141](#)

IPsec

AH, [618](#)

ESP, [619-620](#)

IKE

IKEv1 modes, [624](#)

IKEv1 phases, [625-631](#)

modes of operation, [620-622](#)

services for transition to IPv6, [636](#)

site-to-site VPNs

configuring with CCP, [647-657](#)

planning and preparation checklist, [643](#)

verifying configuration, [658-661](#)

VPN connections, monitoring, [661](#)

IPv4

attacks, comparing with IPv4, [296-298](#)

transition to IPv6, [283-284](#)

IPv6

ACLs

mitigating attacks with, [360](#)

standard ACLs, [361](#)

syntax, [362-363](#)

addressing, [286-292](#)

address representation, [285-286](#)

global unicast addressing, [287-288](#)

link-local addressing, [288-289](#)

multicast addressing, [289-290](#)

unicast addressing, [286-287](#)

attacks, comparing with IPv4, [296-298](#)

configuring on Cisco IOS routers, [293-295](#)

DHCPv6, [292](#)

EUI-64 interface ID, [292](#)

features, [278-284](#)

headers, [279-280](#)

ICMPv6, [281](#)

NDP, [280-281](#)

need for, [275-277](#)

stateless address autoconfiguration, [280-281](#), [292](#)

tunneling, [284](#)

IPv6-aware IPS systems, [484](#)

IPv6-based attacks, [16](#)

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol), [284](#)

ISE (Cisco Identity Services Engine), [98](#), [204-205](#)

ISL (Inter-Switch Link), [236](#)

ISO 27000 standards, [68](#)

ITIL (Information Technology Infrastructure Library), [68](#)

J-K

Jefferson, Thomas, [540](#)

key management, [555](#)

key length issues, [556-557](#)

keyspaces, [556](#)

keyspaces, [556](#)

known plaintext attacks, [552](#)

L

Layer 2 protection, mitigating attacks, [250](#)

CAM table overflow attacks, [259-260](#)

MAC address spoofing, [260-261](#)

STP attacks, [254-258](#)

VLAN hopping attacks, [251-254](#)

layered defense strategies

firewalls, policies, [391-392](#)

role of firewalls in, [370-371](#)

least privilege principle, [40](#)

legislation

effect on network security, [18-19](#)

regulatory compliance, [50-53](#)

data breach disclosure, [52](#)

examples of, [51-52](#)

globalization, [52](#)

levels of risk, [43](#)

liability, assessing, [76](#)

lifecycle approach to risk management, [49-50](#)

likejacking, [15](#)

limitations of firewalls, [369-370](#)

line timeouts, configuring on Cisco routers, [165](#)

link-local addressing, [288-289](#)

local database configuration (AAA), [191-198](#)

location border, [86](#)

logging

enabling with CCP, [354-355](#)

syslog, configuring, [178-182](#)

M

MAC address spoofing, mitigating, [260-261](#)

MAEC (Malware Attribute Enumeration and Characterization), [25](#)

malicious threats, [23](#)

malware, effect on network security, [18](#)

man-in-the-middle attacks, [32-33](#)

management consoles, [471](#)

management plane

Cisco NFP, [123-128](#)

AAA, deploying, [127-128](#)

secure management and reporting, [124-126](#)

Cisco NFP, RBAC, [126-127](#)

managing

communities with CCP, [144](#)

signatures, [500](#)

manual global unicast address assignment, [291](#)

MD5 hashing, [166](#), [572](#)

mediated access principle, [40](#)

meet-in-the-middle attacks, [554](#)

memory scraping, [15](#)

menu bar, CCP, [136](#)

messages, steganography, [539](#)

methodologies of hackers, [21-23](#)

tools, [21-22](#)

methods of classification, [6](#)

MIBs, [182](#)

Microsoft EPDump, [22](#)

Microsoft Remote Procedure Call Dump, [22](#)

minimum length passwords, setting for Cisco IOS routers, [165](#)

mirrored crypto ACLs, [644](#)

misconfiguration of Zone-Based Policy Firewall actions, [411](#)

mitigating Layer 2 attacks, best practices, [250](#)

money muling, [18](#)

monitoring

ACLs with CCP, [356](#)

IPsec VPN connections, [661](#)

versus reporting, [485](#)

motivation for attacks, [13-14](#)

MPF (Cisco Modular Policy Framework), [443-446](#)

MTD (maximum tolerable downtime), [78](#)

multicast addressing, [289-290](#)

multiple privilege levels, configuring on Cisco IOS devices, [170-171](#)

multi-string signature engines, [501](#)

N

NAC (Network Admission Control), [200](#)

named ACLs, [324](#)

NAT (network address translation), [384-390](#)

for Cisco ASA 5500 series, [428-431](#)

deployment modes, [389-390](#)

inside source address translation, [387](#)

PAT, [386-387](#)

NAT overload, [389](#)

native VLANs, [237](#)

NAT-TP, [284](#)

navigation pane, CCP

Device Hardening folder, [140](#)

Firewall option, [141](#)

Interface Management option, [138-139](#)

IPS option, [141](#)

Router folder, [139-140](#)

Security folder, [140](#)

NDP (Neighbor Discovery Protocol), [280-281](#)

need for network security, [12](#)

business goals, [12](#)

risk management, [12-13](#)

negative alarms, [475](#)

NetCat, [22](#)

network borders

application border, [86](#)

device border, [85](#)

location border, [86](#)

network infrastructure, threats against, [112-113](#)

network security

assumptions made regarding networks, [2](#)

compliance regulations, [50-53](#)

need for, [12](#)

business goals, [12](#)

risk management, [12-13](#)

podcasts regarding, [17](#)

posture assessment, [69-70](#)

requirements, [2-3](#)

testing, [70-72](#)

techniques, [71](#)

tools, [71-72](#)

trends affecting, [16-19](#)

changes in workplace experience, [16](#)

cloud computing, [16](#)

exploitation of trust, [17-18](#)

malware, [18](#)

regulatory compliance, [18-19](#)

NFP (Cisco Network Foundation Protection), [112-118](#)

control plane, [118-122](#)

CoPP, [118-119](#)

CPPr, [119](#)

data plane, [128-131](#)

ACL filtering, [128-129](#)

antispoofing, [129-130](#)

Layer 2 protection, [131](#)

management plane, [123-128](#)

AAA, deploying, [127-128](#)

RBAC, [126-127](#)

secure management and reporting, [124-126](#)

nonblind attacks, [32](#)

non-malicious threats, [23](#)

non-Windows version, Cisco Secure ACS, [203-204](#)

NTP (Network Time Protocol), [124](#)

configuring, [177](#)

numbered ACLs, [324](#)

NVD (National Vulnerability Database), [8](#)

O

object groups, [343-345](#)

configuring with CCP, [357-358](#)

One-Step Lockdown (CCP), [152](#)

OOB (out-of-band) management, [124](#)

guidelines, [176](#)

operations and maintenance phase (SDLC), [67](#)

OSI (Open Systems Interconnection) model, [26](#)

application layer, security controls, [309](#)

Layer 2, attack mitigation best practices, [250](#)

outbound access control, configuring on Cisco ASA, [446-460](#)

outbound ACLs, [323](#)

Outbreak Intelligence, [102](#)

overlapping physical controls, [11](#)

overt channels, [31](#), [33-35](#)

OWASP (Open Web Application Security Project), [24](#)

owner role in classification system, [7](#)

P

packet filtering

ACLs, [320-325](#)

developing, best practices, [345-347](#)

implicit deny any statement, [324](#)

named ACLs, [324](#)

numbered ACLs, [324](#)

object groups, [343-345](#)

removing, [325](#)

wildcard bits, [331-334](#)

stateful packet-filtering firewalls, [378-382](#)

static packet-filtering firewalls, [372-374](#)

packet sniffing, [31](#)

packets, classes, [403](#)

password attacks, [28-30](#)

hashes, [29](#)

passwords

configuring on Cisco IOS devices, [163-166](#)

hashes, [167](#)

PAT (Port Address Translation), [386-387](#)

PDA's, exploits, [15](#)

personal association, classifying data by, [6](#)

PFS (Perfect Forward Secrecy), [632](#)

pharming, [31](#), [35-36](#)

phishing attacks, [15](#), [31](#), [35-36](#)

phreakers, [20](#)

physical controls, [10](#)

detective controls, [11](#)

deterrent controls, [11](#)

preventive controls, [11](#)

ping sweeps, [25](#)

PKI (Public Key Infrastructure), [587-602](#)

CAs, [590-593](#)

certificates, [590](#)

revocation methods, [599-600](#)

standards, [593-599](#)

planning considerations for secure management and reporting, [175](#)

podcasts regarding network security, [17](#)

policy management layer component (Cisco Borderless Network Security Architecture), [91](#)

policy maps, [405](#)

Policy NAT, [389](#)

policy-based IPS systems, [477](#)

political motivation for attacks, [14](#)

polyalphabetic ciphers, [541](#)

port scans, [25](#)

port security

configuring, [261-269](#)

errdisable recovery feature, [262-263](#)

violation modes, [262](#)

PortFast, [255](#)

positive alarms, [475](#)

preventive controls, [11](#)

principles of secure network design, [39-42](#)

defense in depth, [41-42](#)

private sector data classification, [5](#)

procedures, [61](#)

promiscuous mode, [468](#), [471](#)

protecting IPv6 networks, best practices, [300](#)

public data classification, [5](#)

public-key encryption, [583-587](#)

PVRST+, verifying, [248](#)

Q

QoS (quality of service), [277](#)

qualitative risk analysis, [44](#)

quantitative risk analysis formula, [45-47](#)

SLE, calculating, [45](#)

R

RADIUS, [206](#)

RBAC (role-based access control), [116](#), [126-127](#)

regulatory compliance, [50-53](#)

data breach disclosure, [52](#)

effect on network security, [18-19](#)

globalization, [52](#)

remote-access VPNs, [537](#)

removing ACLs, [325](#)

reporting versus monitoring, [485](#)

reputation-based IPS systems, [478](#)

requirements for network security, [2-3](#)

restricted data, [5](#)

revocation methods (certificates), [599-600](#)

Rijndael cipher, [563-564](#)

risk analysis, [44-48](#). *See also* [risk management](#)

ALE, calculating, [46-47](#)

ARO, calculating, [45](#)

building blocks of, [47-50](#)

categorization of vulnerabilities, [47](#)

qualitative, [44](#)

quantitative risk analysis formula, [45-47](#)

SLE, calculating, [45](#)

risk matrix, calculating, [48-49](#)

threats, [48](#)

risk management, [3-4](#)

disaster recovery, [77-78](#)

incident response, [72-73](#)

intent evolution, [13-14](#)

levels of risk, [43](#)

liability, assessing, [76](#)

lifecycle approach, [49-50](#)

need for network security, [12-13](#)

reduction of risk, [17](#)

threats, [14-16](#)

risk-based IPS systems, [481-483](#)

Rivest ciphers, [564](#)

Role-Based CLI Access, configuring, [171-174](#)

roles of security policies, [61-62](#)

ROM monitor, configuring, [167-168](#)

root bridge election (STP), [246](#)

Root Guard, [257-259](#)

Router folder (CCP), [139-140](#)

routers, AAA, [186-205](#)

access, authenticating, [188-189](#)

authorization, configuring, [190](#)

routing protocols, maintaining integrity with CPr, [121](#)

RPO (Recovery Point Objective), [78](#)

RR (risk rating), [484](#)

RSA algorithm, [585-587](#)

RSTP (Rapid Spanning Tree Protocol), [246](#)

verifying, [248](#)

RTO (Recovery Time Objective), [78](#)

rule-based policies (Cisco Secure ACS), [222](#)

rules, CCP, [347-348](#)

adding, [350-351](#)

associating with interfaces, [352-353](#)

S

SafeScan Web Security, [102](#)

SBU (sensitive but unclassified) data classification, [5](#)

Scherbius, Arthur, [540](#)

script kiddies, [21](#)

SDEE (Security Device Event Exchange), [486](#)

SDKs (software development kits), [22](#)

SDLC (system design lifecycle), [64-68](#)

acquisition and development phase, [65-66](#)

disposition phase, [67](#)

implementation phase, [66-67](#)

initiation phase, [65](#)

models and frameworks, [67-68](#)

operations and maintenance phase, [67](#)

SEAL (Software-optimized Encryption Algorithm), [565](#)

secret data classification, [5](#)

secure management access

passwords, configuring on Cisco IOS devices, [163-166](#)

Role-Based CLI Access, configuring, [171-174](#)

ROM monitor, configuring, [167-168](#)

SSH, configuring, [161-162](#)

secure management and reporting

OOB and in-band architecture management, [176](#)

planning considerations, [175](#)

secure network lifecycle, [64-68](#)

acquisition and development phase, [65-66](#)

disposition phase, [67](#)

implementation phase, [66-67](#)

initiation phase, [65](#)

models and frameworks, [67-68](#)

operations and maintenance phase, [67](#)

SecureX, [93-98](#)

context-awareness, [94](#)

ISE, [98](#)

SIO, [94-95](#), [99-100](#)

Security Audit feature (CCP), [140](#)

One-Step Lockdown, [152](#)

Security Audit feature, device hardening, [149-151](#)

Security folder (CCP), [140](#)

security levels (Cisco ASA), [432-434](#)

security models, [184](#)

security policies, [53-63](#)

assets, identifying, [53-54](#)

audience for, [54](#)

end-user policies, [57](#)

governing policy, [56-57](#)

guidelines, [60](#)

procedures, [61](#)

reasons for having, [54](#)

roles filled within, [61-62](#)

roles of, [54-55](#)

security awareness, [62-63](#)

standards, [60](#)

technical policies, [57-59](#)

SEM (security event management), [482](#)

sensitive data classification, [5](#)

sensors

alarms, [474-475](#)

inline mode, [471](#)

promiscuous mode, [471](#)

limitations of, [471](#)

sequence prediction, [27](#)

service signature engines, [501](#)

SHA-1 hashing, [572](#)

SHA-2 hashing, [573](#)

SIEM (Security Information and Event Management) ecosystem partners, [311](#)

signature-based IPS systems, [475-477](#)

signatures, [498-505](#)

managing, [500](#)

SMEs, [500-501](#)

tuning, [502-507](#)

SIM (security information management), [482](#)

SIO (Cisco Security Intelligence Operations), [94-95](#), [99-100](#)

site-to-site VPNs, [536-537](#)

building blocks of, [643-646](#)

configuring with CCP, [647-657](#)

crypto maps, [646](#)

planning and preparation checklist, [643](#)

verifying configuration, [658-661](#)

SLE (single loss expectancy), [45](#)

calculating, [45](#)

SMEs (signature microengines), [500-501](#)

SNMP (Simple Network Management Protocol), [182-185](#)

MIBs, [182](#)

SNMPv3, [184-185](#)

social engineering, [22, 31](#)

social networking

exploitation of trust, [17-18](#)

likejacking, [15](#)

SSH (Secure Shell), configuring, [161-162](#)

SSL (Secure Sockets Layer)

and TLS, [673-674](#)

tunnel establishment, [675-679](#)

SSL VPNs

in borderless networks, [670-672](#)

clientless SSL VPN, configuring with ASDM, [683-691](#)

deployment options, [679-683](#)

Stacheldrucht, [39](#)

standard ACLs

configuring, [335-337](#)

in IPv6 environment, [361](#)

standards, [60](#)

ISO 27000 standards, [68](#)

PKI, [593-599](#)

stateful packet-filtering firewalls, [378-382](#)

stateless address autoconfiguration, [280-281, 292](#)

static packet-filtering firewalls, [372-374](#)

Static PAT, [390](#)

status bar (CCP), [142](#)

steganography, [539](#)

Sternberg, David, [23](#)

STP (Spanning Tree Protocol), [244-248](#)

attacks, mitigating, [254-258](#)

BPDU, [247](#)

BPDU Guard, [256-257](#)

designated port selection, [247](#)

PortFast, [255](#)

root bridge election, [246](#)

Root Guard, [257-259](#)

RSTP, [246](#)

stream ciphers, [548-549](#)

string signature engines, [501](#)

Stuxnet worm, [13](#)

subnetting, [326-328](#)

Class C IP addresses, [326-327](#)

classful prefix length, [326](#)

example of, [327-328](#)

VLSM, [328-330](#)

substitution ciphers, [541](#)

Suite B cryptographic standard, [611-612](#)

switches, VLANs, [234-235](#)

creating, [240-242](#)

inter-VLAN routing, configuring, [243-244](#)

symmetric encryption algorithms

3DES, [562](#)

AES, [563-564](#)

DES, [560-562](#)

syntax, IPv6-based ACLs, [362-363](#)

syslog, configuring on Cisco IOS devices, [178-182](#)

T

TACACS+, [205-206](#)

AAA configuration example, [215-216](#)

troubleshooting, [216-218](#)

TCP session hijacking, [33](#)

TCP/IP, IP spoofing, [26](#)

technical controls, [9-10](#)

technical policies, [57-59](#)

technologies

in Cisco ASA, [431-432](#)

in IPS systems, [475-476](#)

technologies involved in Cisco threat control and containment, [312-314](#)

Telnet, configuring vty access, [338](#)

templates, CCP, [145-146](#)

Teredo tunneling, [284](#)

testing network security, [70-72](#)

techniques, [71](#)

tools, [71-72](#)

threat vectors, [306-307](#)

threats, [3](#)

blended threats, [15](#), [39](#)

Cisco threat control and containment strategies, [311-317](#)

classification of, [23-25](#)

controlling, design guidelines, [308](#)

evolution of in information security, [306-307](#)

to network infrastructure, [112-113](#)

in risk analysis, [48](#)

vulnerabilities, classification of, [7-8](#)

TLS (Transport Layer Security) and SSL, [673-674](#)

toolbar, CCP, [138](#)

tools used by hackers, [21-22](#)

top secret data classification, [5](#)

TR (threat rating), [484](#)

traffic, controlling with ACLs, [335-342](#)

traffic classes

actions, [404](#)

identifying for CPPr, [120-121](#)

traffic flow for Zone-Based Policy Firewall, [409-410](#)

training, security awareness, [62-63](#)

transition to IPv6, [283-284](#)

IPsec services, [636](#)

transparent firewalls, [383-384](#)

transport mode (IPsec), [621](#)

transposition ciphers, [542-543](#)

trends affecting network security, [16-19](#)

exploitation of trust, [17-18](#)

malware, [18](#)

trends in information security threats, [306-307](#)

Trojan horse attacks, [22](#)

troubleshooting TACACS+, [216-218](#)

true negatives, [474](#)

true positives, [474](#)

trunking, [235-237](#)

802.1Q tagging, [236-237](#)

configuring, [238-239](#)

ISL, [236](#)

verifying, [239-240](#)

trust exploitation, [28](#)

TrustSec, [95-96](#)

tuning signatures, [502-507](#)

tunnel mode (IPsec), [621-622](#)

tunneling, [284](#)

GRE, [534](#)

Turing, Alan, [553](#)

U

unclassified data, [4](#)

unicast addressing, [286-287](#)

useful life, classifying data by, [6](#)

user profiles (CCP), [147-148](#)

user role in classification system, [7](#)

V

value, classifying data by, [6](#)

verifying

RSTP, [248](#)

site-to-site VPN configuration, [658-661](#)

trunks, [239-240](#)

Vigenère cipher, [541](#)

vigilantism as motivation for attacks, [14](#)

violation modes (port security), [262](#)

virtualization exploits, [15](#)

VLAN hopping attacks, mitigating, [251-254](#)

VLANs, [234-235](#)

assigning to switch port, [242-243](#)

creating, [240-242](#)

inter-VLAN routing, configuring, [243-244](#)

native VLANs, [237](#)

trunking, [235-237](#)

802.1Q tagging, [236-237](#)

ISL, [236](#)

VLSM (variable-length subnet masking), [328-330](#)

example of, [329-330](#)

VPNs, [534-537](#)

Cisco AnyConnect client, [97-98](#)

SSL VPN full-tunnel mode, configuring with ASDM, [692-707](#)

clientless VPNs, [672](#)

IPsec, monitoring connections, [661](#)

remote-access, [537](#)

security, [105-106](#)

site-to-site, [536-537](#)

building blocks of, [643-646](#)

configuring, [647-657](#)

crypto maps, [646](#)

verifying configuration, [658-661](#)

SSL VPNs

in borderless networks, [670-672](#)

deployment options, [679-683](#)

uity access, configuring, [338](#)

vulnerabilities, [3-4](#)

classification of, [7-8](#)

W

WASC TC (Web Application Security Consortium Threat Classification), [24](#)

weakest link concept, [40](#)

web security with Cisco Borderless Network Security Architecture, [101-105](#)

websites, attacks on, [15](#)

wildcard bits, [331-334](#)

Windows operating system, Cisco Secure ACS, [201-202](#)

wiretapping, [31](#)

workplace experience, effect on network security, [16](#)

worms, Stuxnet worm, [13](#)

X-Y-Z

zero day attacks, [39](#)

zone pairs, [402-403](#)

Zone-Based Policy Firewall, [398-403](#)

actions, [407-408](#)

class maps, [405-406](#)

features, [400-401](#)

interzone policies, [399](#)

configuring, [411-422](#)

NAT services, configuring, [422-426](#)

policy maps, [405](#)

traffic flows, [409-410](#)

zones, [402-403](#)

zones, [402-403](#)



ciscopress.com: Your Cisco Certification and Networking Learning Resource

The screenshot shows the ciscopress.com website with a top navigation bar including Home, Products, Safari Bookshelf, Authors, Chapters & Articles, Promotions, Certification Info, and a search bar. Below the navigation, there are sections for 'Cisco Press is the only authorized publisher...', 'CERTIFICATION INFO', 'STORE | NEWSLETTERS | SERIES', 'CISCO NETWORKING ACADEMY', 'On INFORMAT', 'OnCERTIFICATION Video Podcasts', 'Network World's Cisco Subnet', and 'Safari'. A sidebar on the right features 'Quick Links', 'Become a Member', 'Most Popular' books, 'Just Released' books, and 'Coming Soon' books. A vertical banner on the right side of the page reads 'Get Prepared. Get ahead. Cisco Learning Network'.

Subscribe to the monthly Cisco Press newsletter to be the first to learn about new releases and special promotions.

Visit ciscopress.com/newsletters.

While you are visiting, check out the offerings available at your finger tips.

–Free Podcasts from experts:

- OnNetworking
- OnCertification
- OnSecurity



View them at ciscopress.com/podcasts.

–Read the latest author articles and sample chapters at ciscopress.com/articles.

–Bookmark the Certification Reference Guide available through our partner site at informit.com/certguide.

Connect with Cisco Press authors and editors via Facebook and Twitter, visit informit.com/socialconnect.





Cisco Learning Network

Free Test Prep and Beyond.

- ✓ Access review questions
- ✓ Watch Quick Learning Modules (QLMS)
- ✓ Search for jobs and network with others
- ✓ Take self-assessments
- ✓ Participate in study groups
- ✓ Play online learning games

Register for a free membership
and get started now.
www.cisco.com/go/learningnetwork

Cisco Learning Network
A social learning site brought to you by Learning@Cisco

PEARSON IT CERTIFICATION

Browse by Exams ▾

Browse by Technology ▾

Browse by Format

Explore ▾

I'm New Here - Help!

Store

Forums

Safari Books Online

Pearson IT Certification

THE LEADER IN IT CERTIFICATION LEARNING TOOLS

Visit pearsonITcertification.com today to find:

- IT CERTIFICATION EXAM information and guidance for



CompTIA

Microsoft

vmware

Pearson is the official publisher of Cisco Press, IBM Press, VMware Press and is a Platinum CompTIA Publishing Partner—CompTIA's highest partnership accreditation

- EXAM TIPS AND TRICKS from Pearson IT Certification's expert authors and industry experts, such as

- *Mark Edward Soper* – CompTIA
- *David Prowse* – CompTIA
- *Wendell Odom* – Cisco
- *Kevin Wallace* – Cisco and CompTIA
- *Shon Harris* – Security
- *Thomas Erl* – SOACP



- SPECIAL OFFERS – pearsonITcertification.com/promotions
- REGISTER your Pearson IT Certification products to access additional online material and receive a coupon to be used on your next purchase

Articles & Chapters



Blogs



Books



Cert Flash Cards Online



eBooks



Mobile Apps



Newsletters



Podcasts



Question of the Day



Rough Cuts



Short Cuts



Software Downloads



Videos



CONNECT WITH PEARSON IT CERTIFICATION

Be sure to create an account on pearsonITcertification.com and receive members-only offers and benefits



```
R1(config)# line console 0
R1(config-line)# login
R1(config-line)# password M3rcury$12
R1(config-line)# exit
R1(config)# line vty 0 4
R1(config-line)# login
R1(config-line)# password V3nus$2012
```

```
R1(config)# security passwords min-length 10
```

```
R1(config)# username SecAdmin secret 0 Curium2012
```

```
R1(config)# username SecAdmin secret 5 $1$uypB$vAWRWP.qFQqb65.KxVxKg1
```

```
R1(config)# no service password-recovery
```

```
WARNING:
```

```
Executing this command will disable password recovery mechanism. Do not execute this  
command without another plan for password recovery.
```

```
Are you sure you want to continue? [yes/no]: yes
```

```
R1(config)#
```

```
R1# show secure bootset
```

```
IOS resilience router id FHK085031MD
```

```
IOS image resilience version 12.3 activated at 05:00:59 UTC Fri Feb 10 2006
```

```
Secure archive flash:c1841-advsecurityk9-mz.123-14.T1.bin type is image (elf) []  
file size is 17533860 bytes, run size is 17699528 bytes
```

```
Runnable image, entry point 0x8000F000, run from ram
```

```
IOS configuration resilience version 12.3 activated at 05:01:02 UTC Fri Feb 10 2  
006
```

```
Secure archive flash:.runcfg-20060210-050102.ar type is config
```

```
configuration archive size 4014 bytes
```

```
R1(config)# privilege exec level 2 ping
```

```
R1(config)# enable secret level 2 Cariboo2012
```



```
R1> enable 2
```

```
Password: Cariboo2012
```

```
R1# show privilege
```

```
Current privilege level is 2
```

```
R1(config-view)# commands parser-mode {include | include-exclusive | exclude} [all]  
[interface interface-name | command]
```

```
R1> enable view
```

```
Password:
```

```
R1# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R1(config)# parser view NetOps
```

```
R1(config-view)# secret 0 hardtocrackpw
```

```
R1(config-view)# commands exec include ping
```

```
R1(config-view)# commands exec include all show
```

```
R1(config-view)# commands exec include telnet
```

```
R1(config-view)# commands exec include traceroute
```

```
R1(config-view)# commands exec include write
```

```
R1(config-view)# commands exec include configure
```

```
R1(config-view)# commands configure include access-list
```

```
R1(config-view)# commands configure include all interface
```

```
R1(config-view)# commands configure include all ip
```

R1# **enable view NetOps**

Password: hardtocrackpw

R1#

Jan 3 13:45:03.887: %PARSER-6-VIEW_SWITCH: successfully set to view 'NetOps'.

R1#?

Exec commands:

configure Enter configuration mode

enable Turn on privileged commands

exit Exit from the EXEC

ping Send echo messages

show Show running system information

telnet Open a telnet connection

traceroute Trace route to destination

write Write running configuration to memory, network, or terminal

R1# **configure terminal**

R1(config)#?

Configure commands:

access-list Add an access list entry

do To run exec commands in config mode

exit Exit from configure mode

interface Select an interface to configure

ip Global IP configuration subcommands

```
R1(config)# aaa new-model
R1(config)# aaa local authentication attempts max-fail 10
R1(config)# aaa authentication login default local
R1(config)# aaa authentication login MGT-ACCESS local enable
R1(config)# enable secret SnowyDay2012
R1(config)# username admin privilege 15 view root secret sanfran2012
R1(config)# username FWadmin privilege 10 view CCP_Firewall secret 1StopUn0w
R1(config)# line con 0
R1(config-line)# login authentication MGT-ACCESS
R1(config-line)# end
R1# debug aaa authentication
```

User Access Verification

Username: **wrongusername**

Password:

Feb 11 11:06:47.971: AAA/BIND(0000001B): Bind i/f

Feb 11 11:06:47.971: AAA/AUTHEN/LOGIN (0000001B): Pick method list 'MGT-ACCESS'

Feb 11 11:06:48.223: AAA/AUTHEN/ENABLE(0000001B): Processing request action LOGIN

Feb 11 11:06:48.223: AAA/AUTHEN/ENABLE(0000001B): Done status GET_PASSWORD

Feb 11 11:06:49.231: AAA/AUTHEN/ENABLE(0000001B): Processing request action LOGIN

Feb 11 11:06:49.235: AAA/AUTHEN/ENABLE(0000001B): Done status FAIL - bad password

% Authentication failed

Username: **admin**

Feb 11 11:06:51.239: AAA/AUTHEN/LOGIN (0000001B): Pick method list 'MGT-ACCESS'

Password:

R1>

```
aaa new-model
!
aaa authentication login TACACS_SERVER group tacacs+ local
aaa authorization exec default group tacacs+
aaa authorization network default group tacacs+
aaa accounting exec default start-stop tacacs+
aaa accounting network default start-stop tacacs+
aaa accounting commands 15 default stop-only group tacacs+
!
!
tacacs-server host 10.0.1.11
tacacs-server key ciscosecure
!
line vty 0 4
  login authentication TACACS_SERVER
```

Router# **debug tacacs**

```
13:53:35: TAC+: Opening TCP/IP connection to 192.168.60.15 using source 192.48.0.79
13:53:35: TAC+: Sending TCP/IP packet number 416942312-1 to 192.168.60.15 (AUTHEN/
START)
13:53:35: TAC+: Receiving TCP/IP packet number 416942312-2 from 192.168.60.15
13:53:35: TAC+ (416942312): received authen response status = GETUSER
13:53:37: TAC+: send AUTHEN/CONT packet
13:53:37: TAC+: Sending TCP/IP packet number 416942312-3 to 192.168.60.15 (AUTHEN/
CONT)
13:53:37: TAC+: Receiving TCP/IP packet number 416942312-4 from 192.168.60.15
13:53:37: TAC+ (416942312): received authen response status = GETPASS
13:53:38: TAC+: send AUTHEN/CONT packet
13:53:38: TAC+: Sending TCP/IP packet number 416942312-5 to 192.168.60.15 (AUTHEN/
CONT)
13:53:38: TAC+: Receiving TCP/IP packet number 416942312-6 from 192.168.60.15
13:53:38: TAC+ (416942312): received authen response status = FAIL
13:53:40: TAC+: Closing TCP/IP connection to 192.168.60.15
```


Router# debug tacacs events

%LINK-3-UPDOWN: Interface Async2, changed state to up

00:03:16: TAC+: Opening TCP/IP to 192.168.58.104/1049 timeout=15

00:03:16: TAC+: Opened TCP/IP handle 0x48A87C to 192.168.58.104/1049

00:03:16: TAC+: periodic timer started

00:03:16: TAC+: 192.168.58.104 req=3BD868 id=-1242409656 ver=193 handle=0x48A87C

(ESTAB)

expire=14 AUTHEN/START/SENDAUTH/CHAP queued

00:03:17: TAC+: 192.168.58.104 ESTAB 3BD868 wrote 46 of 46 bytes

00:03:22: TAC+: 192.168.58.104 CLOSEWAIT read=12 wanted=12 alloc=12 got=12

00:03:22: TAC+: 192.168.58.104 CLOSEWAIT read=61 wanted=61 alloc=61 got=49

00:03:22: TAC+: 192.168.58.104 received 61 byte reply for 3BD868

00:03:22: TAC+: req=3BD868 id=-1242409656 ver=193 handle=0x48A87C (CLOSEWAIT)
expire=9

AUTHEN/START/SENDAUTH/CHAP processed

00:03:22: TAC+: periodic timer stopped (queue empty)

00:03:22: TAC+: Closing TCP/IP 0x48A87C connection to 192.168.58.104/1049

00:03:22: TAC+: Opening TCP/IP to 192.168.58.104/1049 timeout=15

00:03:22: TAC+: Opened TCP/IP handle 0x489F08 to 192.168.58.104/1049

00:03:22: TAC+: periodic timer started

00:03:22: TAC+: 192.168.58.104 req=3BD868 id=299214410 ver=192 handle=0x489F08
(ESTAB)

expire=14 AUTHEN/START/SENDPASS/CHAP queued

00:03:23: TAC+: 192.168.58.104 ESTAB 3BD868 wrote 41 of 41 bytes

00:03:23: TAC+: 192.168.58.104 CLOSEWAIT read=12 wanted=12 alloc=12 got=12

00:03:23: TAC+: 192.168.58.104 CLOSEWAIT read=21 wanted=21 alloc=21 got=9

00:03:23: TAC+: 192.168.58.104 received 21 byte reply for 3BD868

00:03:23: TAC+: req=3BD868 id=299214410 ver=192 handle=0x489F08 (CLOSEWAIT)
expire=13

AUTHEN/START/SENDPASS/CHAP processed

00:03:23: TAC+: periodic timer stopped (queue empty)

```
Switch(config)# interface fa0/1  
Switch(config-if)# switchport mode trunk
```

```
SwitchX# show interfaces fa0/11 switchport
```

```
Name: Fa0/11
```

```
Switchport: Enabled
```

```
Administrative Mode: trunk
```

```
Operational Mode: down
```

```
Administrative Trunking Encapsulation: dot1q
```

```
Negotiation of Trunking: On
```

```
Access Mode VLAN: 1 (default)
```

```
Trunking Native Mode VLAN: 1 (default)
```

```
SwitchX# show interfaces fa0/11 trunk
```

```
Port      Mode      Encapsulation Status      Native vlan
```

```
Fa0/11    desirable 802.1Q      trunking    1
```

```
Port      Vlans allowed on trunk
```

```
Fa0/11    1-4094
```

```
Port      Vlans allowed and active in management domain
```

```
Fa0/11    1-13Trunking Native Mode VLAN: 1 (default)
```

```
Switch(config)# vlan 2
```

```
Switch(config-vlan)# name Marketing
```

SwitchX# show vlan id 2

VLAN Name	Status	Ports
2	active	Fa0/2, Fa0/12

VLAN	Type	SAID	MTU	Parent	RingNo	BridgeNo	Stp	BrdgMode	Trans1	Trans2
2	enet	100002	1500	-	-	-	-	-	0	0

. . .
SwitchX#

```
SwitchX(config)# interface range fastethernet 0/2 - 4
```

```
SwitchX(config-if)# switchport access vlan 2
```

```
SwitchX# show vlan
```

VLAN	Name	Status	Ports
1	default	active	Fa0/1
2	switchlab99	active	Fa0/2, Fa0/3, Fa0/4

SwitchX# show vlan brief

VLAN Name	Status	Ports	
1	default	active	Fa0/1
2	switchlab99	active	Fa0/2, Fa0/3, Fa0/4
3	vlan3	active	
4	vlan4	active	
1002	fddi-default	act/unsup	
1003	token-ring-default	act/unsup	

VLAN Name	Status	Ports	
1004	fddinet-default	act/unsup	
1005	trnet-default	act/unsup/4	

Switch# show spanning-tree vlan 21

VLAN0021

Spanning tree enabled protocol rstp

Root ID Priority 32789
 Address 88f0.77c5.0f80
 Cost 19
 Port 1 (FastEthernet0/1)
Hello Time 2 sec Max Age 20 sec Forward Delay
15 sec

Bridge ID Priority 32789 (priority 32768 sys-id-ext 21)
 Address
Hello Time d0c2.82c5.6b00
 2 sec Max Age 20 sec Forward Delay
15 sec

 Aging Time 300 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/1	Root	FWD	19	128.1	P2p
Fa0/2	Desg	FWD	19	128.2	P2p
Fa0/8	Desg	FWD	19	128.8	P2p

```
Switch# show spanning-tree summary
```

```
Root bridge for: Bridge group 1, VLAN0001, VLAN0004-VLAN1005
```

```
VLAN1013-VLAN1499, VLAN2001-VLAN4094
```

```
EtherChannel misconfiguration guard is enabled
```

```
Extended system ID is enabled
```

```
Portfast is enabled by default
```

```
PortFast BPDU Guard is enabled
```

```
Portfast BPDU Filter is disabled by default
```

```
Loopguard is disabled by default
```

```
UplinkFast is disabled
```

```
BackboneFast is disabled
```

```
Pathcost method used is long
```

```
<output omitted>
```

```
Switch#
```

```
switch# show interfaces gigabitethernet 4/1 status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Gi4/1	err-disabled	100	full	1000	1000BaseSX	

```
switch# show errdisable recovery
```

```
ErrDisable Reason      Timer Status
```

```
-----
```

```
Udld                    Disabled
```

```
Bpduguard              Enabled
```

```
security-violatio     Disabled
```

```
channel-misconfig     Disabled
```

```
<output omitted>
```

```
Timer interval: 300 seconds
```

```
Interfaces that will be enabled at the next timeout:
```

```
Interface      Errdisable reason      Time left(sec)
```

```
-----
```

```
Gi4/1         bpduguard              290
```

```
Switch(config-if)# switchport port-security [mac-address mac-address  
[vlan {vlan-id | {access | voice} } ] ] | [mac-address sticky  
[mac-address| vlan {vlan-id | {access | voice} } ]] [maximum value  
[vlan {vlan-list | {access | voice} } ] ]
```

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

```
Switch(config-if)# switchport mode access  
Switch(config-if)# switchport port-security  
Switch(config-if)# switchport port-security maximum 2  
Switch(config-if)# switchport port-security violation shutdown  
Switch(config-if)# switchport port-security mac-address sticky  
Switch(config-if)# switchport port-security aging time 120
```

```
sw-class# show port-security
```

```
Secure Port    MaxSecureAddr  CurrentAddr    SecurityViolation  Security Action  
                (Count)        (Count)        (Count)
```

```
-----  
Fa0/12         2              0              0                 Shutdown  
-----
```

```
Total Addresses in System (excluding one mac per port) : 0
```

```
Max Addresses limit in System (excluding one mac per port) : 1024
```



```
sw-class# show port-security interface fa0/12
Port Security          : Enabled
Port status           : Secure-down
Violation mode        : Shutdown
Maximum MAC Addresses : 1
Total MAC Addresses   : 2
Configured MAC Addresses : 0
Aging time            : 120 mins
Aging type            : Absolute
SecureStatic address aging : Disabled
Security Violation Count : 1
```

```
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
```

```
R1(config)# interface fa 0/0
```

```
R1(config-if)# ipv6 address 2001:0DB8:2222:7272::72/64
```

```
R1(config)# interface fa 0/0
```

```
R1(config-if)# ipv6 address 2001:0DB8:0:1::/64 eui-64
```

```
R1(config)# ipv6 unicast-routing
R1(config)# interface fa0/0
R1(config-if)# ipv6 address 2001:db8:c18:1::/64 eui-64
R1# show ipv6 interface fa0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::218:B9FF:FE21:9278
  Global unicast address(es):
  2001:DB8:c18:1:218:B9FF:FE21:9278, subnet is 2000:1:2:3::/64
  Joined group address(es):
  FF02::1:FF21:9278
  FF02::1
  FF02::2
  MTU is 1500 bytes
<output omitted>
```

```
r1(config)# access-list 1 deny 172.16.4.0 0.0.0.255
r1(config)# access-list 1 permit any
( implicit deny all = access-list 1 deny 0.0.0.0 255.255.255.255 )
r1(config)# interface ethernet 0
r1(config-if)# ip access-group 1 out
```

```
r1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.255 172.16.3.0 0.0.0.255 eq 21
r1(config)# access-list 101 deny tcp 172.16.4.0 0.0.0.255 172.16.3.0 0.0.0.255 eq 20
r1(config)# access-list 101 permit ip any any
(implicit deny all)
(access-list 101 deny ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255)
r1(config)# interface ethernet 0
r1(config-if)# ip access-group 101 out
```

```
Router(config)# access-list 102 permit tcp any host 200.1.1.2 established
Router(config)# access-list 102 permit tcp any host 200.1.1.2 eq smtp
Router(config)# interface serial 0
Router(config-if)# ip access-group 102 in
```



```
Router# show access-lists
```

```
Extended IP access list OUTBOUND
```

```
10 permit tcp 10.0.0.0 0.255.255.255 any eq www (67 matches)
```

```
20 permit tcp 10.0.0.0 0.255.255.255 any eq 443 (1190 matches)
```

```
30 permit udp 10.0.0.0 0.255.255.255 any eq domain
```

```
40 deny ip any any (41 matches)
```

```
r1# show ip interfaces FastEthernet 0/0
```

```
FastEthernet 0/0 is up, line protocol is up
```

```
Internet address is 10.0.2.1/24
```

```
Broadcast address is 255.255.255.255
```

```
Address determined by non-volatile memory
```

```
MTU is 1500 bytes
```

```
Helper address is not set
```

```
Directed broadcast forwarding is disabled
```

```
Multicast reserved groups joined: 224.0.0.5 224.0.0.6
```

```
Outgoing access list is not set
```

```
Inbound access list is OUTBOUND
```

```
Proxy ARP is enabled
```

```
<text omitted>
```

```
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.12 eq www
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.12 eq ftp
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.13 eq www
access-list 100 deny tcp host 10.6.252.65 host 171.8.2.13 eq ftp
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.12 eq www
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.12 eq ftp
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.13 eq www
access-list 100 deny tcp host 10.6.252.66 host 171.8.2.13 eq ftp
```

```
object-group network SOURCES
```

```
host 10.6.252.65 host 10.6.252.66
```

```
object-group network DESTINATIONS
```

```
host 171.8.2.12 host 171.8.2.13
```

```
object-group service APPLICATIONS
```

```
tcp www tcp ftp
```

```
access-list 100 deny object-group APPLICATIONS object-group SOURCES object-group  
DESTINATIONS
```

```
Router# config t
Router(config)# object-group network INTERNAL-NETS
Router(config-network-group)# description Subnets inside the Firewall
Router(config-network-group)# 10.10.0.0 255.255.255.0
Router(config-network-group)# 10.10.1.0 255.255.255.0
Router(config-network-group)# 10.10.2.0 255.255.255.0
Router(config-network-group)# 10.10.10.0 255.255.255.0
```

```
permit icmp any any nd-na  
permit icmp any any nd-na  
deny ipv6 any any
```

```
permit protocol {source-ipv6-prefix/prefix-length | any | host
source-ipv6-address | auth} [operator [port-number]] {destination-
ipv6-prefix/prefix-length | any | host destination-ipv6- address |
auth} [operator [port-number]] [dest-option-type [doh-number | doh-
type]] [dscp value] [flow-label value] [fragments] [log] [log-
input] [mobility] [mobility-type [mh-number | mh-type]] [reflect
name [timeout value]] [routing] [routing-type routing-number]
[sequence value] [time-range name]
```

```
permit protocol {source-ipv6-prefix/prefix-length | any | host
source-ipv6-address | auth} [operator [port-number]] {destination-
ipv6-prefix/prefix-length | any | host destination-ipv6- address |
auth} [operator [port-number]] [dest-option-type [doh-number | doh-
type]] [dscp value] [flow-label value] [fragments] [log] [log-
input] [mobility] [mobility-type [mh-number | mh-type]][routing]
[routing-type routing-number] [sequence value] [time-range
name] [undetermined-transport]
```



```
ipv6 access-list RFC4890
```

```
permit icmp any any echo-reply permit icmp any any echo-request permit icmp any any  
1 3
```

```
permit icmp any any 1 4
```

```
permit icmp any any packet-too-big permit icmp any any time-exceeded permit icmp any  
any parameter-problem permit icmp any any mld-query
```

```
permit icmp any any mld-reduction permit icmp any any mld-report permit icmp any any  
nd-na
```

```
permit icmp any any nd-ns
```

```
permit icmp any any router-solicitation
```

```
class-map type inspect match-any OUTBOUND-PROTOCOLS
  match protocol http
  match protocol smtp
  match protocol ftp
!
policy-map type inspect ACCESS-POLICY
  class type inspect OUTBOUND-PROTOCOLS
  inspect
!
zone security PRIVATE
zone security INTERNET
!
interface fastethernet 0/0
  zone-member security PRIVATE
!
interface serial 0/0/0
  zone-member security INTERNET
!
zone-pair security PRIV-TO-INTERNET source PRIVATE destination INTERNET
  service-policy type inspect ACCESS-POLICY
!
```

```
interface FastEthernet0/0
 ip address 200.200.1.2 255.255.255.0
 ip nat outside
!
interface FastEthernet0/1.4
 ip address 10.10.0.1 255.255.255.0
 ip nat inside
!
ip nat inside source list 1 interface FastEthernet0/0 overload
!
access-list 1 permit 10.10.0.0 0.0.0.255
!
```

Router# **show ip nat translations**

Pro	Inside global	Inside local	Outside local	Outside global
TCP	200.200.1.51:1050	10.10.10.20:1050	75.75.75.750:23	172.16.100.10:23
TCP	200.200.1.52:1776	10.10.10.10:1776	150.150.1.40:25	150.150.1.40:25

```
asa5505(config)# interface vlan21
```

```
asa5505(config-if)# nameif inside
```

```
!  
asa5505(config)# interface vlan21  
asa5505(config-if)# nameif inside  
INFO: Security level for "inside" set to 100 by default. ciscoasa(config)# setup  
Pre-configure Firewall now through interactive prompts [yes]? <Enter> Firewall Mode  
[Routed]: <Enter>  
Enable password [<use current password>]: cisco123  
Allow password recovery [yes]? <Enter>  
Clock (UTC):  
  Year [2012]: <Enter>  
  Month [Aug]: <Enter>  
  Day [26]: <Enter>  
  Time [10:21:49]: 15:34:00  
Inside IP address [0.0.0.0]: 10.0.2.1  
Inside network mask [255.255.255.255]: 255.255.255.0  
Host name [ciscoasa]: ASA Domain name: cisco.com  
IP address of host running Device Manager: 10.0.2.11  
Use this configuration and write to flash? Y  
!
```

!

```
asa5505(config)# asdm image disk0:/asdm-641.bin
```

!

```
Router(config)# ip ips name sdm_ips_rule
Router(config)# ip ips config location flash:/ips/ retries 1
Router(config)# ip ips notify SDEE
!
Router(config)# ip ips signature-category
Router(config-ips-category)# category all
Router(config-ips-category-action)# retired true
Router(config-ips-category-action)# exit
Router(config-ips-category)# category ios_ips basic
Router(config-ips-category-action)# retired false
!
Router(config)# interface Ethernet0/0
Router(config-if)# ip ips sdm_ips_rule in
Router(config-if)# ip virtual-reassembly
```


Router# **show ip ips configuration**

IPS Signature File Configuration Status

Configured Config Locations: flash:/ips/

Last signature default load time: 04:39:33 UTC Oct 19 2011

Last signature delta load time: 04:41:43 UTC Oct 19 2011

Last event action (SEAP) load time: -none-

General SEAP Config:

Global Deny Timeout: 3600 seconds

Global Overrides Status: Enabled

Global Filters Status: Enabled

IPS Auto Update is not currently configured

IPS Syslog and SDEE Notification Status

Event notification through syslog is enabled

Event notification through SDEE is enabled

IPS Signature Status

Total Active Signatures: 295

Total Inactive Signatures: 4008

IPS Packet Scanning and Interface Status

IPS Rule Configuration

IPS name sdm_ips_rule

IPS fail closed is disabled

IPS deny-action ips-interface is false

Obsolete tuning is disabled

Regex compile threshold (MB) 14

Interface Configuration

Interface FastEthernet0/0

Inbound IPS rule is sdm_ips_rule

Outgoing IPS rule is not set

Interface FastEthernet0/1

Inbound IPS rule is sdm_ips_rule

Outgoing IPS rule is not set

IPS Category CLI Configuration:

Category all:

Retire: True

Category ios_ips basic:

Retire: False

```
Router# show ip ips interfaces
```

```
Interface Configuration
```

```
Interface FastEthernet0/0
```

```
Inbound IPS rule is sdm_ips_rule
```

```
Outgoing IPS rule is not set
```

```
Interface FastEthernet0/1
```

```
Inbound IPS rule is sdm_ips_rule
```

```
Outgoing IPS rule is not set
```

```
Router# show ip ips signature count
```

```
Cisco SDF release version S594.0
```

```
Trend SDF release version V0.0
```

```
...
```

```
<output omitted>
```

```
...
```

```
Total Signatures: 4303
```

```
Total Enabled Signatures: 1218
```

```
Total Retired Signatures: 4008
```

```
Total Compiled Signatures: 295
```

```
Total Obsoleted Signatures: 13
```

```
Total Disallowed Signatures: 3
```

```
.
```

```
.
```

```
Router# show ip ips statistics
```

```
Signature statistics [process switch:fast switch]
```

```
signature 3041:0: packets checked [0:4] alarmed [0:4] dropped [0:0]
```

```
signature 3040:0: packets checked [0:1] alarmed [0:1] dropped [0:0]
```

```
signature 6062:1: packets checked [0:1] alarmed [0:1] dropped [0:0]
```

```
signature 6054:0: packets checked [0:3] alarmed [0:3] dropped [0:0]
```

```
Interfaces configured for ips 1
```

```
Session creations since subsystem startup or last reset 10101
```

```
Current session counts (estab/half-open/terminating) [15:764:0]
```

```
Maxever session counts (estab/half-open/terminating) [22:1182:0]
```

```
Last session created 00:00:08
```

```
Last statistic reset never
```

```
TCP reassembly statistics
```

```
received 2 packets out-of-order; dropped 0
```

```
peak memory usage 1 KB; current usage: 0 KB
```

```
router# show ip sdee alerts
```

```
Alert storage: 200 alerts using 96000 bytes of memory
```

```
SDEE Alerts
```

SigID	Sig Name	SrcIP:SrcPort	DstIP:DstPort	VRF
		or Summary Info		
1: 3040:0	TCP NULL Packet	172.17.44.101:36044	10.5.5.5:25	NONE
2: 3041:0	TCP SYN/FIN Packet	172.17.44.101:52623	10.5.5.5:25	NONE
3: 6054:0	DNS Version Request	172.17.44.101:4745	172.17.22.103:53	NONE
4: 6062:1	DNS Authors Request	172.17.44.101:4756	172.17.22.103:53	NONE

```
IOS-FW# show crypto isakmp policy
```

```
Global IKE policy
```

```
Protection suite of priority 1
```

```
  encryption algorithm: Three key triple DES
```

```
  hash algorithm: Secure Hash Standard
```

```
  authentication method: Pre-Shared Key
```

```
  Diffie-Hellman group: #2 (1024 bit)
```

```
lifetime: 86400 seconds, no volume limit
```

```
IOS-FW# show crypto ipsec transform-set
```

```
Transform set ESP-3DES-SHA: { esp-3des esp-sha-hmac }
```

```
will negotiate = { Tunnel, },
```

```
Transform set #${default_transform_set_1}: { esp-aes esp-sha-hmac }
```

```
will negotiate = { Transport, },
```

```
Transform set #${default_transform_set_0}: { esp-3des esp-sha-hmac }
```

```
will negotiate = { Transport, },
```

```
IOS-FW# show crypto map
```

```
Crypto Map "SDM_CMAP_1" 1 ipsec-isakmp
```

```
Description: Tunnel to200.200.20.2
```

```
Peer = 200.200.20.2
```

```
Extended IP access list 101
```

```
access-list 101 permit ip 10.10.0.0 0.0.255.255 10.20.10.0 0.0.0.255
```

```
Current peer: 200.200.20.2
```

```
Security association lifetime: 4608000 kilobytes/3600 seconds
```

```
Responder-Only (Y/N): N
```

```
PFS (Y/N): N
```

```
Transform sets={
```

```
ESP-3DES-SHA: { esp-3des esp-sha-hmac } ,
```

```
}
```

```
Interfaces using crypto map SDM_CMAP_1:
```

```
FastEthernet0/0
```



```
RouterA# show crypto ipsec sa
```

```
interface: FastEthernet0/0
```

```
  Crypto map tag: SDM_CMAP_1, local addr 200.200.1.2
```

```
protected vrf: (none)
```

```
local ident (addr/mask/prot/port): (10.10.0.0/255.255.0.0/0/0)
```

```
remote ident (addr/mask/prot/port): (10.20.10.0/255.255.255.0/0/0)
```

```
current_peer 200.200.20.2 port 500
```

```
  PERMIT, flags={origin_is_acl,}
```

```
#pkts encaps: 4030, #pkts encrypt: 4030, #pkts digest: 4030
```

```
#pkts decaps: 4033, #pkts decrypt: 4033, #pkts verify: 4033
```

```
#pkts compressed: 0, #pkts decompressed: 0
```

```
#pkts not compressed: 0, #pkts compr. failed: 0
```

```
#pkts not decompressed: 0, #pkts decompress failed: 0
```

```
#send errors 0, #recv errors 0
```