ALL ■ IN ■ ONE

CCSP™: Cisco® Certified Security Professional Certification

EXAM GUIDE

Robert E. Larson
Lance Cockcroft

McGraw-Hill/Osborne
New York • Chicago • San Francisco • Lisbon
London • Madrid • Mexico City • Milan • New Delhi
San Juan • Seoul • Singapore • Sydney • Toronto

The McGraw-Hill Companies

McGraw-Hill/Osborne
2100 Powell Street, 10th Floor
Emeryville, California 94608
U.S.A.

To arrange bulk purchase discounts for sales promotions, premiums, or fund-raisers, please contact **McGraw-Hill**/Osborne at the above address. For information on translations or book distributors outside the U.S.A., please see the International Contact Information page immediately following the index of this book.

CCSP™: Cisco® Certified Security Professional Certification All-in-One Exam Guide (Exams 642-501 SECUR, 642-521 CSPFA, 642-511 CSVPN, 642-531 CSIDS, and 642-541 CSI)

Copyright © 2003 by The McGraw-Hill Companies. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

1234567890 DOC DOC 019876543

Book p/n 0-07-222692-7 and CD p/n 0-07-222693-5
parts of
ISBN 0-07-222691-9

Publisher
Brandon A. Nordin

Vice President & Associate Publisher
Scott Rogers

Acquisitions Editor
Nancy Maragioglio

Project Editor
Lisa Wolters-Broder

Acquisitions Coordinator
Jessica Wilson

Technical Editors
Joe Phago
Ole Drews Jensen

Copy Editor
Marcia Baker

Proofreaders
Brian Galloway
Linda Medoff

Indexer
Rebecca Plunkett

Compositors
Apollo Publishing Services
George Toma Charbak

Illustrators
Lyssa Wald
Melinda Moore Lytle
Michael Mueller

Series Design
Peter F. Hancik

This book was composed with Corel VENTURA™ Publisher.

Information has been obtained by **McGraw-Hill**/Osborne from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, **McGraw-Hill**/Osborne, or others, **McGraw-Hill**/Osborne does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the results obtained from the use of such information.

This book is dedicated to my parents, Lou and Elmer Larson,
who provided resources and direction when I was young,
plus freedom, inspiration, and support as I got older.
—Bob

About the Authors

Robert E. Larson lives in the Seattle, Washington area with his wife Jerri and four adult children. Bob has worked full-time as a computer trainer and course developer since 1985, including network training since 1995. Bob got involved with the Cisco Networking Academy program in 1998. He is currently the Cisco Regional Academy contact at Bates Technical College in Tacoma, plus teaches evening and weekend CCNP, Security, and CCIE prep classes at Green River Community College. Bob is currently a member of the Cisco Networking Academy Advisory Council. This is Bob's third Cisco certification book, having also written a CCNA and CCNP book. Bob taught the first Academy CCNA series in Africa in 1999 in Cape Town, South Africa. He has also taught CCNP-level courses in Birmingham, England; Dillingen, Germany; and Vienna, Austria.

Lance Cockcroft, Net+, CCA, MCSE, MCT, CCNP, CCDP, has been a Senior Engineer for many ISP and telecommunications companies, including Bellsouth, Atlanta Broadband, and Southeastern Networks. Lance is currently the Cisco Product Manager for Self Test Software, Cisco's only authorized test prep vendor. Lance writes and oversees the production of all Cisco practice tests for Self Test Software. Lance attended and continues to teach for Kennesaw State University and Southern Polytechnic University located in his hometown of Marietta, Georgia.

About the Technical Reviewers

Ole Drews Jensen began working with computers 21 years ago, and five years later made it his profession. He started out as a programmer in a wide variety of languages, but soon got involved with administering servers and networks. Today Ole is the Systems Network Manager for an enterprise company with several subsidiaries in the recruiting industry, where one of the largest is Carlton Staffing. Ole holds the following certifications: CCNP, MCSE, and MCP+I, and is currently pursuing the new CCSP.

Setotolwane Johannes "Joe" Phago, CCIE # 7105, CCNP, Cisco Firewall Specialist, Cisco VPN Specialist, B.Sc. Computer Science (University of the North, S.A.). He was the first Black South African CCIE and is a graduate of the first Cisco Networking Academy in Africa. Joe is currently Senior Network Analyst at Standard Bank of South Africa, a leading banking and financial services company in S.A. and Africa with a presence on virtually all continents.

CONTENTS

Introduction xxi

Part I Introduction to Network Security 1

Chapter 1 Understanding Network Security Threats 3

Identify the Need for Network Security 4
Identify the Causes of Network Security Problems 5
Technology Weakness 6
Policy Weakness 7
Configuration Weakness 8
The Four Primary Types of Network Threats 8
Unstructured Threats 8
Structured Threats 9
Internal Threats 10
External Threats 10
The Four Primary Types of Network Attack 11
Reconnaissance Attacks 11
Access Attacks 14
Denial of Service (DoS) Attacks 16
Data Manipulation Attacks 20
Cisco AVVID and SAFE Strategies 22
AVVID 22
SAFE 23
Cisco Security Wheel 23
Network Security Policy 25
Why Create a Network Security Policy 25
The Balancing Act 26
A Security Policy Is to Be Shared 28
Who Should Help Create the Security Policy? 29
Assets and Threats 30
Evaluating a Network Security Policy 32
Example of a Network Security Policy 35
Securing the Network 35
Wireless Communication Policy 36
Monitoring Network Security 37
Improving Network Security 38
Chapter Review 39
Questions 40
Answers 44

Chapter 2 Securing the Network 47

Secure Network Design Example 48
Inside Network 49
Outside Network 49
Demilitarized Zone (DMZ) 49
Securing Network Devices 50
Physically Secure the Devices 50
Securing Administrative Access 50
Using Access Control Lists to Secure the Network 57
Standard ACLs 57
Extended Access Lists 64
Named Access Lists 66
Time-Based Access Lists 66
Chapter Review 71
Questions 71
Answers 74

Part II Securing the Network Perimeter 75

Chapter 3 Cisco AAA Security Technology 77

The Cisco AAA Model 78
NAS Servers 78
Why Authenticate? 79
AAA Benefits 82
TACACS+, RADIUS, and Kerberos Support 83
AAA System Components 88
AAA as Facilitator 88
Authentication 92
Authorization 96
Accounting 99
Testing AAA Configuration 103
The show Commands 103
The debug Commands 103
Chapter Review 104
Questions 105
Answers 107

Chapter 4 Cisco Secure ACS and TACACS+/RADIUS Technologies 109

Describe Cisco Secure ACS 110
CiscoSecure ACS for Windows and UNIX 110
Features and Architecture of Cisco Secure ACS for Windows 111
Features and Benefits 111
Cisco Secure ACS Benefits 112
Cisco Secure ACS for Windows Internal Architecture 113
System Performance 117
Features of CiscoSecure ACS for UNIX 118
Features and Benefits 118
Preparing to Install UNIX ACS 119

Installing Cisco Secure ACS 3.0 for Windows . . . 119
Hardware Requirements . . . 120
Operating System Requirements . . . 120
Third-Party Software Requirements . . . 120
NAS Minimum IOS Requirements . . . 121
Network Requirements . . . 121
Back Up Server Data . . . 121
Gathering Information Required During Installation . . . 122
Administering and Troubleshooting Cisco Secure ACS for Windows . . . 122
Navigation Bar . . . 123
Configuration Area . . . 125
Display Area . . . 125
Accessing the HTML Interface . . . 125
Suggested Configuration Sequence . . . 128
TACACS+ Overview . . . 132
Configuring Cisco Secure ACS and TACACS+ . . . 133
Configure NAS to TACACS+ Server Communication . . . 134
Verifying TACACS+ . . . 136
The show Commands . . . 136
The debug Commands . . . 136
Configure NAS to RADIUS Server Communication . . . 137
Chapter Review . . . 138
Questions . . . 139
Answers . . . 141

Chapter 5 Securing Cisco Perimeter Routers . . . 143

Perimeter Router Terms and Concepts . . . 143
Simple Secure Network Design . . . 144
Eavesdropping . . . 147
Router Solutions . . . 147
Hub and Switch Issues . . . 149
Limit Unneeded TCP/IP and Other Services . . . 150
TCP and UDP "Small Services" . . . 150
Finger . . . 150
NTP . . . 150
CDP . . . 150
Denial of Service Attacks . . . 150
Controlling Directed Broadcasts . . . 151
Flood Management . . . 151
Antispoofing with RPF Checks . . . 152
Unauthorized Access . . . 152
Address Filtering . . . 152
Dynamic (Lock-and-Key) Access Lists . . . 152
Reflexive Access Lists . . . 157
Lack of Legal IP Addresses . . . 161
NAT Technology and Terminology . . . 162
Static NAT . . . 163
Dynamic NAT . . . 165
Dynamic NAT with Overloading (PAT) . . . 167
Rerouting Attacks . . . 169

Event Logging on Perimeter Routers . . . 170
Access List Violation Logs . . . 171
Chapter Review . . . 171
Questions . . . 172
Answers . . . 174

Chapter 6 IOS Firewall Feature Set—CBAC . . . 175

Introduction to Cisco IOS Firewall . . . 175
Router-Based Firewall Functionality . . . 176
Integration with Cisco IOS Software . . . 176
Feature Summary . . . 178
Context-Based Access Control (CBAC) . . . 179
Quick Access List Review . . . 179
CBAC Advantages . . . 179
CBAC Limitations . . . 181
CBAC Process . . . 181
Configuring CBAC . . . 182
IOS Firewall Management . . . 198
Command Line Interface . . . 198
ConfigMaker . . . 199
Chapter Review . . . 200
Questions . . . 201
Answers . . . 203

Chapter 7 IOS Firewall—Intrusion Detection System . . . 205

Intrusion Detection System (IDS) . . . 205
IOS Firewall Intrusion Detection System . . . 206
Devices Supporting the IOS Firewall IDS Features . . . 206
Cisco IDS Attack Signatures . . . 208
Cisco Secure IDS Director Support . . . 209
Performance Implications . . . 210
IOS IDS vs. Cisco Secure IDS . . . 210
Cisco IOS Firewall IDS Configuration Task List . . . 211
Initializing the IOS Firewall IDS . . . 212
The ip audit smtp spam Command . . . 212
The ip audit po max-events Command . . . 212
Initializing the Post Office . . . 212
The ip audit notify Command . . . 213
The ip audit po local Command . . . 214
The ip audit po remote Command . . . 215
Creating and Applying Audit Rules . . . 216
Creating an Audit Rule . . . 217
Apply the Audit Rule to the Interface(s) . . . 220
Verifying the IDS Configuration . . . 222
The show ip audit statistics Command . . . 222
The show ip audit configuration Command . . . 223
The show ip audit interface Command . . . 223
The show ip audit all Command . . . 224
Chapter Review . . . 224
Questions . . . 225
Answers . . . 227

Chapter 8 IOS Firewall—Authentication Proxy 229

Cisco IOS Firewall Authentication Proxy 229
How the Authentication Proxy Works 230
Applying the Authentication Proxy 232
Comparison with the Lock-and-Key Feature 233
Compatibility with Other Features 233
Security Vulnerability Issues 236
Before Configuring Authentication Proxy 236
Authentication Proxy Configuration Task List 238
AAA Server Configuration 238
AAA Router Configuration 244
Enable AAA 244
Define the Security Server 244
Define Login Authentication Methods List 249
Enable Authorization Proxy (auth-proxy) for AAA 250
Activate Authentication Proxy Accounting 251
ACL Entry for Return Traffic from the AAA Server 252
Configuring the HTTP Server 253
Authentication Proxy Configuration on the Router 254
The ip auth-proxy auth-cache-time Command 254
The ip auth-proxy auth-proxy-banner Command 255
The ip auth-proxy name Command 255
The auth-proxy Interface Configuration 257
Verify Authentication Proxy Configuration 257
The auth-proxy Cache 258
The debug Commands 259
CBAC Configuration 259
Chapter Review 260
Questions 260
Answers 263

Part III Virtual Private Networks (VPNs) 265

Chapter 9 Cisco IOS IPSec Introduction 267

Virtual Private Networks 268
Remote–Access 269
Site-to-Site 270
Layer 2 VPNs 271
Layer 3 VPNs 272
Other VPN Implementations 273
Why Use VPNs? 274
VPN Analogy 274
Tunneling Protocols 275
Layer Two Forwarding (L2F) Protocol 276
Layer 2 Tunneling Protocol (L2TP) 276
Generic Routing Encapsulation (GRE) 276
How IPSec Works 276
Cisco IOS IPSec Technologies 277
IPSec Security Overview 278

Transport and Tunnel Mode . . . 281
IPSec Transforms and Transform Sets . . . 286
Cisco IOS Cryptosystem Components . . . 288
How Encryption Works . . . 288
Cryptography Types . . . 290
Encryption Alternatives . . . 290
Hashing . . . 292
Diffie-Hellman Key Agreement (DH) . . . 293
Security Association (SA) . . . 294
IKE SAs versus IPSec SAs . . . 295
Five Steps of IPSec Revisited . . . 296
Step 1—Determine Interesting Traffic . . . 296
Step 2—IKE Phase One . . . 297
Step 3—IKE Phase Two . . . 300
Step 4—IPSec Data Transfer . . . 301
Step 5—Session Termination . . . 301
IPSec Support in Cisco Systems Products . . . 301
Chapter Review . . . 302
Questions . . . 303
Answers . . . 305

Chapter 10 Cisco IOS IPSec for Preshared Keys . . . 307

Configure IPSec Encryption Tasks . . . 307
Task 1 Prepare for IKE and IPSec . . . 309
Task 2 Configure IKE . . . 317
Task 3 Configure IPSec . . . 321
Task 4 Test and Verify IPSec . . . 329
Configuring IPSec Manually . . . 333
Configuring IPSec Manually Is Not Recommended . . . 334
Chapter Review . . . 335
Questions . . . 336
Answers . . . 339

Chapter 11 Cisco IOS IPSec Certificate Authority Support . . . 341

CA Support Overview . . . 341
Digital Certificates . . . 342
Certificate Distribution . . . 343
IPSec with CAs . . . 344
How CA Certs Are Used by IPSec Peers . . . 344
Cisco IOS CA Standards . . . 345
Simple Certificate Enrollment Protocol (SCEP) . . . 345
CA Servers Interoperable with Cisco Routers . . . 346
Enroll a Device with a CA . . . 348
Configure CA Support Tasks . . . 348
Task 1—Prepare for IKE and IPSec . . . 349
Task 2—Configure CA Support . . . 351
Task 3—Configure IKE . . . 369
Task 4—Configure IPSec . . . 371
Task 5—Test and Verify IPSec . . . 372

RSA Encrypted Nonces Overview 372
Task 2—Configure RSA Keys 373
Chapter Review 374
Questions 377
Answers 379

Chapter 12 Cisco IOS Remote Access Using Cisco Easy VPN 381

Introduction to Cisco Easy VPN 381
Cisco Easy VPN Server 382
Client Connection Process 382
Cisco Easy VPN Remote 383
Split Tunneling 384
Cisco VPN 3.6 Client 385
How the VPN Client Works 385
Connection Technologies 385
Easy VPN Server Configuration Tasks 386
Preconfiguring the Cisco VPN 3.6 Client 386
Creating a New Connection Entry 387
Trying Out the New Connection 389
Customizing the Connection 390
Management Center for VPN Routers 392
Features and Benefits 393
Router MC Server Requirements 394
Router MC Client Requirements 394
Router MC User Permissions 395
Easy VPN Remote Phase Two 396
Supported VPN Servers 396
Phase Two Features 396
Cisco VPN Firewall Feature for VPN Client 402
Overview of Software Client Firewall Feature 402
Defining a Client Firewall Policy 403
The Are You There Feature 403
The Central Policy Protection Feature 404
Client/Server Feature 406
Client Firewall Statistics 407
Chapter Review 408
Questions 409
Answers 411

Chapter 13 Cisco VPN Hardware Overview 413

Cisco Products Enable a Secure VPN 413
What's New? 414
Cisco VPN 3002 Client Devices 414
Cisco VPN 3002 Client Models 415
Client and Network Extension Modes 416
Standards Supported 417
Cisco VPN 3002 Hardware Client Features 417
Cisco VPN 3000 Concentrator Devices 419
Cisco VPN 3000 Concentrator Models 419

Standards Supported 423
Cisco VPN 3000 Concentrator Features 424
VPN 3000 Concentrator Client Support 426
Chapter Review 429
Questions 430
Answers 432

Chapter 14 Cisco VPN 3000 Remote Access Networks 435

VPN Concentrator User Interfaces and Startup 436
Quick Configuration 437
Command-Line Interface (CLI) Basics 439
Concentrator Manager (Web Interface) 443
VPN Concentrators in IPSec VPN Implementations 450
Remote Access Networks 451
LAN-to-LAN Networks 451
Remote Access VPNs with Preshared Keys 452
Preshared Keys 453
Initial Configuration 454
Setting the Public Interface 455
Defining the Default Gateway (Optional) 456
Adding the Static Routes 458
General System Information 459
Define Inside Address Assignment Method 459
Define Inside Address Pool for Remote Users 461
Configuring Groups and Users 461
Other Configuration Options 473
Digital Certificates 477
Certificate Types 477
VPN Concentrator and Certificates 477
Enrolling and Installing Certificates 478
Using SCEP to Manage Certificates 479
Using the Certificates 484
Configure Cisco VPN Client Support 486
VPN Client Autoinitiation Feature 487
The vpnclient.ini File 487
Preparation 488
Configuration 488
VPN 3000 Configuration 489
Administer and Monitor Remote Access Networks 489
Administration 489
Monitoring 494
Chapter Review 495
Questions 496
Answers 499

Chapter 15 Configuring Cisco VPN 3002 Remote Clients 501

The VPN 3002 in the Network 502
VPN Modes 503
IPSec VPNs 504
Configuring the 3002 Device 506

Command-Line Interface (CLI) . . . 506
The Hardware Client Manager (Web Interface) . . . 511
Common Configuration Tasks . . . 515
Upgrading the Software . . . 515
Quick Configuration . . . 517
System Status . . . 519
PPPoE Support . . . 519
Basic Configuration for the VPN 3002 . . . 521
Set the System Time, Date, and Time Zone . . . 522
Optional—Upload an Existing Configuration File . . . 523
Configure the Private Interface . . . 523
Configure the Public Interface . . . 526
Configure the IPSec . . . 527
Choose Client (PAT) Mode or Network Extension Mode . . . 528
Configure DNS . . . 529
Configure Static Routes . . . 529
Change the Admin Password . . . 530
Modifying Options . . . 531
Other VPN 3002 Software Features . . . 532
Interactive Hardware Client Authentication . . . 532
Individual User Authentication . . . 533
LEAP Bypass . . . 535
IPSec Backup Servers . . . 536
IPSec Server Load Balancing . . . 537
H.323 Support in PAT Mode . . . 540
Simple Certificate Enrollment Protocol (SCEP) . . . 541
XML Management . . . 542
Reverse Route Injection (RRI) . . . 542
AES Support and Diffie-Hellman Group 5 . . . 543
Push Banner to VPN 3002 . . . 544
Delete with Reason . . . 544
Auto-Update Feature . . . 546
VPN 3002 Hardware Clients . . . 546
Cisco VPN Software Clients . . . 546
Configuring Auto-Update . . . 546
Chapter Review . . . 547
Questions . . . 549
Answers . . . 551

Chapter 16 Cisco VPN 3000 LAN-to-LAN Networks . . . 553

The VPN Concentrators in LAN-to-LAN VPNs . . . 553
Chapter Scenario . . . 555
LAN-to-LAN Networks with Preshared Keys . . . 555
Configure Network Lists . . . 556
Define the IKE Proposals (Optional) . . . 560
Create the Tunnel . . . 561
LAN-to-LAN Networks with Digital Certificates . . . 566
NAT Issues . . . 567
NAT Transparency . . . 568
IPSec over TCP . . . 569
IPSec over NAT-T . . . 570

IPSec over UDP . . . 571
LAN-to-LAN VPN with Overlapping Network Addresses . . . 572
LAN-to-LAN Routing . . . 575
Default Gateways . . . 576
Reverse Route Injection . . . 577
Virtual Router Redundancy Protocol . . . 578
Chapter Review . . . 581
Questions . . . 582
Answers . . . 584

Part IV PIX Firewalls . . . 585

Chapter 17 CiscoSecure PIX Firewalls . . . 587

Firewall and Firewall Security Systems . . . 587
Packet Filter . . . 588
Proxy Filter . . . 589
Stateful Packet Filter . . . 589
CiscoSecure PIX Firewall Technology . . . 589
PIX Adaptive Security Algorithm . . . 591
The PIX Firewall Family . . . 592
Tested and Certified . . . 595
VPN Support . . . 595
PIX Management Options . . . 596
Cisco Mobile Office Support . . . 596
Cisco Catalyst 6500 Implementation . . . 596
Basic PIX Firewall Configuration . . . 597
PIC Command-Line Interface . . . 597
The nameif Command . . . 599
The interface Command . . . 599
The ip address Command . . . 601
The nat Command . . . 601
The global Command . . . 602
The route Command . . . 604
Chapter Review . . . 604
Questions . . . 605
Answers . . . 607

Chapter 18 Getting Started with the Cisco PIX Firewall . . . 609

Basic PIX Firewall Configuration . . . 609
Verifying Configuration and Traffic . . . 612
ICMP Traffic to the Firewall . . . 612
The show icmp Command . . . 614
The debug icmp trace Command . . . 614
Time Setting and NTP Support . . . 614
How NTP Works . . . 614
NTP and PIX Firewalls . . . 615
Syslog Configuration . . . 617
The logging Commands . . . 618
FTP and URL Logging . . . 620
Verifying and Monitoring Logging . . . 621

DHCP Server Configuration 625
Configuring the DHCP Server Feature 626
DHCP Client 631
Using NAT/PAT with DHCP Client 632
Firewalls as a DHCP Client and Server 632
Chapter Review 633
Questions 634
Answers 637

Chapter 19 Access Through the PIX Firewall **639**

Adaptive Security Algorithm 639
Security Levels 640
Stateful System 642
Translations 643
Connections 643
Translations and Connections 644
Transport Protocols 646
Static Translations 649
Network Address Translation 654
Port Address Translations (PAT) 658
Using NAT and PAT Together 659
Names and Name Commands 659
Configuring DNS Support 660
Access Control Lists (ACLs) 661
Using Access Lists 661
Access-Group Statement 662
Basic ACL Statements 662
ICMP ACL Statements 663
TurboACL 664
Downloadable ACLs 666
Content Filtering 668
ActiveX Blocking 669
Java Blocking 669
Websense Filtering 670
Object Grouping 673
Overview of Object Grouping 673
Getting Started with Group Objects 674
Configuring Object Groups with ACLs 675
Nested Object Groups 676
Conduit Statements 676
Configuring Conduits 677
PIX Routing Configuration 678
The Route Command 678
Routing Options 680
Multicast Traffic 682
Chapter Review 682
Questions 683
Answers 685

Chapter 20 Advanced PIX Firewall Features **687**

Remote Access 687
Telnet Access 688

HTTP Access . . . 689
Secure Shell (SSH) Access . . . 690
AAA Support for Telnet, HTTP, and SSH Sessions . . . 691
AAA on the PIX Firewall . . . 691
Defining the AAA Server . . . 691
Local User Database . . . 693
Configuring AAA Features . . . 695
Access Lists with AAA . . . 699
Command-Level Authorization . . . 700
Firewall Privilege Levels . . . 701
Configuring Cisco Secure ACS for Windows . . . 702
Advanced Protocol Handling . . . 702
Application Inspection . . . 702
The fixup protocol Command . . . 703
Supported Applications and Protocols . . . 704
Fixup Protocol Examples . . . 706
Other Supported Protocols and Applications . . . 709
Attack Guards . . . 710
DNS Control . . . 711
Flood Defender . . . 711
FragGuard and Virtual Reassembly . . . 712
TCP Intercept . . . 714
Unicast Reverse Path Forwarding . . . 714
ActiveX Blocking, Java Filtering, and URL Filtering . . . 715
Intrusion Detection . . . 715
Define Default Audit Actions . . . 716
Disabling Individual Signatures . . . 716
Create Named Audit Rules . . . 717
Apply the Audit Rule to the Interface(s) . . . 717
PIX Firewall IDS Syslog Messages . . . 718
Shunning . . . 718
Managing SNMP Services . . . 719
PIX Firewall SNMP Support . . . 719
SNMP Contact and Location . . . 720
SNMP Management Station . . . 721
SNMP Community Key . . . 721
Enabling SNMP Traps . . . 722
Verify SNMP Configuration . . . 722
Logging to the SNMP Management Station . . . 722
Chapter Review . . . 723
Questions . . . 724
Answers . . . 726

Chapter 21 Firewalls and VPN Features . . . **729**

Pix Firewall Enables a Secure VPN . . . 729
IPSec VPN Establishment . . . 731
Five Steps of IPSec . . . 731
IPSec Configuration Tasks . . . 732
Task 1: Prepare to Configure VPN Support . . . 732
Task 2: Configure IKE Parameters . . . 733

Task 3: Configure IPSec Parameters . . . 740
Task 4: Test and Verify VPN Configuration . . . 747
Cisco VPN Client . . . 748
Client Mode . . . 748
Network Extension Mode . . . 748
Establishing Preliminary Connectivity . . . 749
Easy VPN Remote Configuration . . . 749
Scale PIX Firewall VPNs . . . 750
Network Management Options . . . 750
PPPoE and the PIX Firewall . . . 752
Chapter Review . . . 754
Configuring IPSec . . . 754
Configuring IPSec for RSA Encrypted Nonces . . . 757
Configuring CA Support Tasks . . . 757
Questions . . . 760
Answers . . . 763

Chapter 22 Managing and Maintaining the PIX Firewall . . . 765

PDM Overview . . . 765
Versions and Device Support . . . 767
PDM Operating Requirements . . . 767
PIX Firewall Requirements . . . 767
Workstation Requirements . . . 768
Cisco Secure Policy Manager Considerations . . . 769
Web Browser Considerations . . . 769
Prepare for PDM . . . 771
Installing PDM on a PIX Firewall . . . 771
Minimum PIX Configuration . . . 772
Starting PDM . . . 772
Using the PDM Startup Wizard . . . 774
Using PDM to Configure the PIX Firewall . . . 775
Using PDM to Create a Site-to-Site VPN . . . 776
Using PDM to Create a Remote Access VPN . . . 780
CiscoWorks Management Center for PIX Firewalls (PIX MC) . . . 783
System Requirements . . . 783
PIX Failover Feature . . . 784
Understanding Failover . . . 785
Failover Configuration with Failover Cable . . . 789
LAN-Based Failover Configuration . . . 792
Verifying Failover Configuration . . . 793
Password Recovery . . . 794
Before Getting Started . . . 794
PIX Devices with a Floppy Drive . . . 795
PIX Devices Without a Floppy Drive . . . 796
Upgrading the PIX OS . . . 797
Older Upgrade Methods . . . 798
Chapter Review . . . 800
Questions . . . 801
Answers . . . 803

Part V Intrusion Detection Systems (IDS) 805

Chapter 23 Intrusion Detection System Overview 807

Security Threats 807
Internal Threats 808
External Threats 808
Unstructured Threats 809
Structured Threats 809
The Attack Types and Phases 809
Attack Types 810
Attack Phases 811
Intrusion Detection Systems Overview 816
Host- and Network-Based IDSs 817
IDS Triggers 821
Summary 827
Questions 829
Answers 832

Chapter 24 Cisco Secure Intrusion Detection System 835

CIDS Operations and Functionality 836
Monitoring 836
Analyzing 841
Communications 841
Centralized Alarm Display and Management 845
Sensor Response 848
CIDS Architecture 850
CIDS Software Architecture 851
CIDS Commands 860
CIDS Directory Structure 861
CIDS Log Files 863
Chapter Review 866
Questions 867
Answers 871

Chapter 25 Sensor Installation and Configuration 873

Sensor Deployment Considerations 873
Network Entry Points 874
Network Size and Complexity 877
The Amount and Type of Traffic 877
Sensor Installation 878
Connecting to Your Network Sensor Appliance 878
Sensor Bootstrap 880
IDS Device Manager 885
Connecting to the IDS Device Manager 886
IDS Device Manager GUI Interface 887
Device Area Configuration 890
Configuration Area 894
Monitoring Area 911
Administration Area 912
Chapter Review 917

Questions 918
Answers 919

Chapter 26 Signature and Alarm Management 921

CIDS Signatures 922
Signature Series 922
Signature Implementations 924
Signature Structure 925
Signature Classes 926
Signature Types 927
Signature Severity 929
Event Viewer 930
Managing Alarms 931
Event Viewer Customization 936
Preference Settings 938
Chapter Review 940
Review Questions 941
Answers 943

Part VI Cisco SAFE Implementation 945

Chapter 27 Cisco SAFE Implementation 947

Preparation Documents 947
Exam Topics 948
Security Fundamentals 948
Architectural Overview 948
Cisco Security Portfolio 948
SAFE Small Network Design 949
SAFE Medium Network Design 949
SAFE Remote-User Network Implementation 949
Skills Required for the Exam 950
Chapter Review 950
Questions 951
Answers 954

Appendix A Access Control Lists 955

Access List Basics 955
Two-Step Process 956
Numbered ACL Common Characteristics 957
The Numbers Matter 957
Standard Access Lists 958
Building a Standard ACL 958
Verifying ACLs 963
Show Run Command 963
Show Access-Lists Command 964
Show IP Interfaces Command 964
Extended Access Lists 965
Creating an Extended Access List 965
Named Access Lists 971

Appendix B About the CD 975

System Requirements 975
LearnKey Online Training 975
Installing and Running MasterExam 976
MasterExam 976
Electronic Book 976
Lab Exercises 976
Help 976
Removing Installation(s) 977
Technical Support 977
LearnKey Technical Support 977

Index 979

INTRODUCTION

Before You Get Started

Welcome to the *CCSP™: Cisco® Certified Security Professional Certification All-in-One Exam Guide*. This book is here to help you prepare to take–and pass–the following Cisco security certification exams. Even more importantly, it is here to share a pool of knowledge that should help you become more employable in the field. If you strive for knowledge and experience, the certification will come. The CCSP exams are:

- Securing Cisco IOS Networks
- Cisco Secure PIX Firewall Exam
- Cisco Secure Virtual Private Networks
- Cisco Secure Intrusion Detection Systems Exam
- Cisco SAFE Implementation Exam

In this section, we discuss skill building and exam preparation alternatives, the certification exam situation itself, the Cisco certification programs in general, and how this book can help you prepare for Cisco certification exams. We will look at the following:

- Things to do to prepare
- CCNA exam insights
- Cisco Certification Information

CCSP Certification Program

The Cisco Certified Security Professional is a brand-new CCNP-level certification track being driven by the rapidly changing and growing world concern about security. For that reason there have been and will continue to be a great number of changes and additions to the program. There have been three major changes in the program in its first year. At the same time, some of the security products have gone through major upgrades, adding many new and useful features.

What this means to you is that it is very important to keep on top of the current exam numbers and exam objectives. Use the Cisco web site at www.cisco.com and the Learning and Events link to get to the latest certification information. The direct link is: http://www.cisco.com/en/US/learning/le3/learning_career_certifications_and_learning_paths_home.html.

In developing this book, we tried to include the information that is required to pass the various certification exams while at the same time anticipating any new topics that might become exam objectives in the near future.

Because the book covers all five exams, much of the security overview information that appears at the beginning of every book has been consolidated into Chapter 1. Other exam sections may use topics covered in the SECUR exam as foundation. The following table shows the relationships between the exams and chapters. The X indicates the material should be included, while an R is recommended.

CCSP All-In-One	SECUR	CSVPN	CSPFA	CSIDS	CSI
Chapter					
Introduction to Network Security					
1. Understanding Network Security Threats	X				X
2. Securing the Network	X				X
Securing the Network Perimeter					
3. Cisco AAA Security Technology	X				X
4. CiscoSecure ACS and TACACS+ Technologies	X				X
5. Securing Cisco Perimeter Routers	X				X
6. IOS Firewall Feature Set - CBAC	X		R		X
7. IOS Firewall Feature Set - Intrusion Detection System	X			R	X
8. IOS Firewall Feature Set - Proxy Authentication	X				X
Virtual Private Networks (VPNs)					
9. Cisco IOS IPSec Introduction	X	X	R		X
10. Cisco IOS IPSec for Pre-Shared Keys	X	R	R		X
11. Cisco IOS IPSec Certificate Authority Support	X	R	R		X
12. Cisco IOS Remote Access Using Cisco Easy VPN	X	X			X
13. Cisco VPN Hardware Overview		X			X
14. Cisco VPN 3000 Remote Access Networks		X			X
15. Configuring Cisco VPN 3002 Remote Clients		X			X
16. Cisco VPN 3000 LAN-to-LAN Networks		X			X
PIX Firewalls					
17. Cisco PIX Firewall Technology and Features			X		X

CCSP All-In-One	SECUR	CSVPN	CSPFA	CSIDS	CSI
PIX Firewalls					
18. Getting Started with the Cisco PIX Firewall			X		X
19. Access Through the PIX Firewall			X		X
20. Advanced PIX Firewall Features			X		X
21. Firewalls and VPN Features			X		X
22. Managing and Maintaining the PIX Firewall			X		X
Intrusion Detection Systems (IDS)					
23. IDS Overview and CSIDS Installation				X	X
24. Alarms and Signatures				X	X
25. CIDS Installation and Configuration				X	X
26. Signature and Alarm Management				X	X
Cisco SAFE Strategy					
27. Cisco SAFE Strategy					X
Appendix A - Access Control Lists	R		R		

How to Protect Yourself Against Exam Changes

Become very familiar with the Cisco web site and how to perform searches for documents. Use the site to stay current on any exam changes. Be sure to look at both the exam description and the *Recommended Training* descriptions. Both will have objectives and topics covered usually as bulleted lists. Consider printing these out and using them as check-off guides to monitor your learning progress. It will also help you to spot new technologies or features introduced in later descriptions.

Release Notes

As you are preparing for a particular topic, perform searches for release notes on that topic, for example *VPN 3000 Concentrator release notes*. Look over the results looking for the latest version; they are not always sorted with the latest at the top. Look particularly at the System Requirements, Upgrading, and New Features sections. Pay particular attention to and feature that was recently added to either the exam or course description on the certifications pages.

Technical Documentation

On the Cisco site, go to the products section for the technology that you are studying and use the links on the left side to find *Technical Documentation* section where you will often find User Guides, Command Reference, Configuration Guides, etc. Each of these documents is available in HTML format and many are available as PDFs.

Find the User Guide or Configuration Guide for the technology (PIX, VPN Concentrator, etc.) and look up the features that are new to you. This is also an excellent way to

get a different perspective than the one presented in this or any other book. If you do not have access to some of the technologies (some are very expensive to acquire just for study purposes) look for the Getting Started Guide. Spend some time studying the parts of these documents that are new or unclear for you.

Finally, search for any configuration examples. These documents are often listed under the *Technical Documentation* heading of the product information, or use the search feature. These are typically very specific and usually include diagrams, instructions, configuration output, and useful links. For technologies with web-based interfaces, many include step-by-step instructions with web captures of the entire process.

NOTE Many documents do not require a CCO account, but if asked to login you will be given an opportunity to apply for a CCO account. The process will only require answering some questions. Even the most limited level may make additional documents available to you.

Remember Your Goal

You are, after all, attempting to become recognized as an expert in these technologies. Don't sell yourself short. Look over the most recent (latest version) documents so that you are not surprised by look-and-feel changes or the addition of a key feature on a menu or screen.

Things to Do to Prepare

I cannot emphasize enough how important it is to get some hands-on experience with Cisco devices whenever possible. The exams ask many questions involving the command syntax or web interface page feature options. Experience configuring devices is the best way to become comfortable with any Cisco technology. I have tried to include enough screen captures to assist you if hands-on experience is not possible. The last section covered using Cisco documentation to checkout new features, but it is equally as valuable for building familiarity with devices you do not have access to. In this section we will look at some other options.

Unlike some other certification, memorizing a long list of facts is not necessarily the best approach for Cisco exams. You must be able to apply the information and see it from other perspectives. The following list of resources that can help you study and prepare:

This Book and Related Materials

Preparing for any Cisco certification exam (including the CCSP) requires you to obtain and study materials designed to provide comprehensive information about the subject matter that will appear on your specific exam. This book contains the framework to prepare to pass the exam. The task now is to apply and absorb that information and become

comfortable with it. This will present different levels of challenge based on your experience with networking. Obviously, someone who has been working in the field for a period of time will and possibly has another advanced certification, such as CCNP, will have a solid base of knowledge and skills that they can build on. I think this book can be a good tool for that person.

The other type of CCSP student I find is the recent CCNA who is interested in getting into the IT field but has little or no real networking experience. I have tried to write this book for that person, as well. The latter student may need some background material, and may need to look at things from two or more perspectives; the Cisco web site and online articles can help with this.

Labs and Exercises

On the CD-ROM you will find labs and exercises for most of the technologies covered. Even if you do not have access to the required equipment, look over the labs. They have a methodology that will be useful as well as many screen captures or sample output to augment the materials in the related chapter.

SAFE and AVVID Documents

The fifth and final exam for CCSP is the *Cisco SAFE Implementation Exam (CSI 642-541 CSI)*. While based on the series of SAFE documents, such as the *SAFE Blueprint for Small, Midsize, and Remote-User Networks*, every technology, topic, or configuration process covered on the other four exams is fair game. Do yourself a favor and start by downloading the SAFE documents in PDF form. Read them at least the *SAFE Blueprint for Small, Midsize, and Remote-User Networks* before getting too far into the book. Then as you learn about each technology review how it fits into the SAFE strategy. Make sure that you can configure the main connections, such as router VPN to PIX VPN. The SAFE documents have additional configuration examples that should help broaden your knowledge.

Classroom Training

Whether you use this book or not, classroom training for many people is the preferred way to learn complex technologies. In this field that classroom training should be combined with hands-on experience with real routers and switches. There are several possible courses to follow:

Cisco Networking Academies

I believe in this program for the average person. Since 1987, Cisco Systems has set up Networking Academies in more than 10,500 locations around the world. Many are in high schools and the rest are in community colleges, technical colleges, trade schools, universities, and at some service organizations. This highly developed multimedia curriculum, combined with abundant hands-on experience offered part-time, can create a

solid foundation. The academies offer CCNA, CCNP, and Fundamentals of Security (SECUR and CSPFA) training and are now branching out to include non-Cisco technologies like UNIX and web design. To learn more about the Academy Program or to locate one in your area, check the following web site: http://www.cisco.com/warp/public/779/edu/academy/.

Cisco Training Partners

In larger cities, for the working administrator with solid foundation skills who truly meets the course prerequisites, these short, often five-day courses can be a quick way to fill in the gaps, gain limited hands-on experience, and move on to certification. I really like these programs for working professionals with a lot of experience. For them, this type of training can be an excellent value. On the other hand, if a person really doesn't fit the target audience and can't keep up with the class, this can be a very expensive reality check. For more information, go to http://www.cisco.com/ and click on the Learning link.

Buying Equipment

Many students do purchase equipment, particularly if their long-term goal is CCNP, CCSP, or CCIE. Cisco vendors like Blackbox and www.cdw.com offer catalogs and knowledgeable support people. I have always had very good luck with eBay (www.ebay.com). Do a search on "Cisco" at the eBay site and there will be thousands of items. The key is that you can't be in a hurry. Watch for the deal that you want, and be ready to walk away. If you are worried about fraud, deal only with sellers who have made many transactions (a number after their ID) and have an easily viewable performance record.

There are two ways to use eBay. First, look at the people offering items. Many have web sites linked to their auctions. See what kind of businesses they are and what other "deals" they have going. Second, if I'm buying a bigger item, I only buy from an auction that will take a credit card. I then use a card that guarantees my purchases. I've bought hundreds of items and I don't feel that I've ever been hurt. I've never had an item fail to be delivered pretty much as advertised.

Virtual Labs and Simulators

While I think simulators do not replace hands-on experience, they are significantly better than nothing at all. It is my understanding that www.boson.com is working on a simulator for these exams.

Practice Exams

I really hate the thought of a person taking a test repeatedly until they know enough of the questions to pass. This leads to what the industry refers to as "paper certifications" or worse "vapor certifications." It's bad for the industry and can't be all that great for the individual. What value is the certification if you get fired from the job because you can't do the work?

Having said that; I do believe in taking practice tests once you have trained and prepared yourself. This serves two purposes. First, it may point out gaps or weaknesses in your training plan. Second, and more importantly, it helps to prepare you for the exam itself. If you have taken the CCNA or CCNP exams, you already know that Cisco exams are like none you've taken before. While they are fair and valid, they are not designed to pass a lot of students. They are designed to see if you know the exam material forwards and backwards. My students have found that the exams at www.boson.com are both challenging and helpful.

Cram Sessions and Brain Dumps

There are web sites called brain dumps, where test-takers try to list as many test questions as they can remember. First, my personal opinion is that these are a waste of time and energy. Second, they violate the non-disclosure agreement that every test-taker agrees to when they take the exam. In the end, you compromise your integrity for a bit of short-term-memory fodder.

What time I've spent at the sites that I'm aware of, I've found a mix of good and bad questions, questions from old exams, questions from the wrong exams, and a small amount of mischief. There are better ways.

One site I like is http://studyguides.cramsession.com/. They have a series of study guides, usually 12-20 pages, for many exams that I recommend to all of my students. While they do not give you questions, they give you lists of things to know. I really do not believe they replace studying. The practice that I follow, and recommend to my students, is that each night for the week before a scheduled exam, read the Cramsession just before bed. Typically, it will lead me to question some points, and after researching I put the results on the margins of the study guide. Their study guide is the only thing that I ever take to a test site. I try to review it once before going into the test site.

Do you need all of the things covered in this section? Probably not. But I've tried to offer a mix to helpful tools and suggestions.

CCSP Exam Insights

Once you have prepared for your exam, you need to register with a testing center. Each computer-based CCSP exam costs $125 (North America), and if you don't pass, you may retest for an additional $125 for each try. In the United States and Canada, tests are administered by Prometric Testing Centers.

You can sign up for a test through Prometric's web site at http://www.2test.com, or you can register by phone at 800-204-EXAM (within the United States or Canada). The web site will not allow you to schedule exams within 48 hours, so use the phone registration for shorter scheduling intervals. It is possible in some markets to take tests on the same day. Be prepared to wait through voice messages.

To sign up for a test, you will need a valid credit card.

To schedule an exam, call the toll-free number or visit the web page at least one day in advance. Before booking the exam make sure that you understand the cancellation process

and deadlines, currently before 7 P.M. Central Standard Time the day before the scheduled test time (or you will be charged, even if you don't appear to take the test).

When you want to schedule a test, have the following information ready:

- Exam number and title
- Your name–Exactly the way that you want it to appear on your certificate.
- Your social security, social insurance, or Prometric number (SP)
- A method of payment–Credit card
- Contact telephone numbers–In case of a problem so they can reach you.
- Mailing address–Where you want your certificate mailed.
- Email address–For contact purposes. You will get a confirmation via e-mail.

Once you sign up for a test, you will be informed as to when and where the test is scheduled. Try to arrive at least 15 minutes early–personally, due to traffic congestion, I tell students to show up an hour early. You can always relax and review your notes. I've sat in exams next to students who have showed up late for whatever reason. They seem miserable and I suspect the stress and tension will be reflected in their score.

Photo ID

You will need to bring two forms of identification to the testing site. One form *must* be a photo ID such as a driver's license or a valid passport. The other must have a signature. The test cannot be taken without the proper identification.

Gum, Candy, and Cough Drops

Do yourself a favor and bring something with you. It can always just sit there ignored. But the last thing you want is a dry throat or coughing to disrupt your testing and the silence for your peers.

The Exam Process

When you show up at the testing center, you will need to sign in with an exam coordinator. He or she will ask you to show the two forms of signature identification. After you have signed in and your time slot arrives, you will be asked to deposit any items with you such as books, bags, pagers, or calculators. Make sure that you know where the restrooms and drinking fountain are located. You don't want to plan to need them, but even worse is to have to search for them. You will be escorted into a closed room.

All exams are closed book. You will be furnished with one or two blank sheets of paper and a pen or, in some cases, an erasable plastic sheet and an erasable pen. Before the exam–take a few minutes and write out any important material on the blank sheet. This is particularly important for any formulas or detailed data that you might forget under the

stress of the exam. You can refer to this piece of paper any time you like during the test, but you will have to turn it in when you leave.

You will have some time to compose yourself, to record this information, and to take a sample orientation exam before you begin the real test. You will also be required to complete a computer-based survey to track demographics of the test candidates. Typically, if an exam has a 75-minute time limit, you will have 90 minutes to take the sample exam, complete the survey, and take the actual exam. Once you start the actual exam you now have only the exam time limit.

Typically, the room will have up to a dozen computers. Each workstation will be separated from the others by dividers designed to keep you from seeing your neighbor's computer. Keep in mind that the people next to you could be taking a certification exam from an industry totally unrelated to yours, so don't be concerned if someone starts after you or finishes before you. Most test rooms use closed circuit cameras. This permits the exam coordinator to monitor the room.

The exam coordinator will have preloaded the appropriate Cisco certification exam. If there is a problem with the exam, such as version number, screen doesn't display all data, the screen or desk area is dirty, etc., let the coordinator know right away. Do not put yourself at a disadvantage. You can start as soon as you are seated in front of the computer. I suggest that you sit back for a minute and relax. Take a deep breath. If the chair is adjustable, adjust it. Move your arms and legs to release any tension. You are going to be sitting there almost 90 minutes.

All Cisco certification exams allow a certain maximum amount of time in which to complete the work (this time is indicated on the exam by an on-screen counter/clock, so you can check the time remaining whenever you like). All Cisco certification exams are computer generated and most use a multiple-choice format, often with six to eight choices. It is possible, if not likely, that several questions will refer to an exhibit containing dozens of commands from which you will be expected to select one as the answer to a specific question.

Most Cisco exams use some form of simulator in a few questions to test your configuration skills. Typically these are fundamental activities not obscure activities, so make sure that you know how to configure the basics.

While this may sound quite simple, the questions not only are constructed to check your mastery of basic facts and skills about the subject material, but they also require you to evaluate one or more sets of circumstances or requirements. Often, you are asked to give more than one answer to a question, although you will always be told how many to choose. *You get only one pass through the questions*—you cannot mark a question and return to it later.

When you complete a Cisco certification exam, the exam will tell you whether you have passed or failed. All test objectives are broken into several topic areas and each area is scored on a basis of 100 percent. Particularly if you do not pass the exam, select the option on the screen that asks if you want to print the report. The test administrator will print it for you. You can use this report to help you prepare for a second effort, if needed. Once you see your score, you have the option of printing additional copies of the score report. It is a good idea to print it twice.

Remember, if you need to retake an exam, you will have to schedule a new test with Prometric and pay another $125.

Exam Design

All Cisco tests use one of following basic question types:

- Multiple-choice with a single answer
- Multiple-choice with two or more answers (the question will indicate how many answers)
- Multipart with one or more answers (the question will indicate how many answers)
- CLI-based questions (many times, an exhibit will present a sample IOS configuration in which you are asked to choose the correct command or interpret the configuration's output, per the question's directions)
- Drag and drop where steps need to be arranged in order, technologies need to be labeled, or you need to fill in the blanks. Expect a couple of these.
- Simulations to test configuration skills. This will typically be a step in an overall device configuration, such as configuring an interface. Expect no more than a couple of these.

Take the time to read a question at least twice before selecting an answer, and pay special attention to words such as "not" that can radically change the question. If a question seems very simple, great—but read it over once more to make sure that you aren't missing something.

Always look for an Exhibit button as you examine each question. The Exhibit button brings up graphics used to help explain a question, provide additional data, or illustrate network design or program behavior. My perception is that there are fewer exhibits than in the past, with drawings and images included on the screen with the question.

Cisco exams do not allow you to return to questions, so you must make sure to answer the question as best you can before proceeding to the next one. The exam will clearly state before you start whether you can mark answers and return.

Cisco's Testing Format

All Cisco exams are fixed-length with a fixed number of questions. Each candidate will get the same number of questions; the order of the questions can vary, as can the specific questions. If you retake an exam assume there will be different questions. From time to time, questions are replaced and others may not be scored.

Cisco provides a counter in the upper-right corner (near the remaining time) showing the number of questions completed and the number outstanding. Monitor your time to make sure that you have completed at least one-quarter of the questions one-quarter of the way through the exam period and three-quarters of the questions three-quarters of the way through. Have the calculations done in advance, such as 16 questions by 18 minutes.

If you are not finished with 10 minutes remaining, try to pick up the pace. At five minutes remaining, use the remaining time to guess your way through any remaining questions. Guessing is better than not answering because blank answers are always wrong, but a guess may turn out to be right. The important thing is to answer every question.

Some Basic Question-Handling Strategies

For those questions that take only a single answer, usually two or three of the answers will be obviously incorrect, and a couple of the answers will be plausible. Of course, only one can be correct. Unless the answer leaps out at you, begin the process of eliminating those answers that are most obviously wrong.

Many questions assume that the default behavior of a particular command or option is in effect. If you know the defaults and understand what they mean, this will help you with your choice.

Cisco exams are generally pretty straightforward and not intended to beat you out of your certification, but then again they are not designed to be easy. Pay attention, particularly with syntax. Knowing the difference between *access-list 1 deny any* and *access list 1 deny any* should be assumed (note the hyphen).

If the answer seems immediately obvious, reread the question to look for a trap; sometimes those are the ones you are most likely to get wrong.

Typically, at least one answer out of the possible choices for a question can be eliminated immediately because the answer does not apply to the situation or the answer describes a nonexistent issue or option.

If faced with guessing among two or more potentially correct answers, reread the question. Try to picture how each of the possible remaining answers would alter the situation. Be especially sensitive to terminology; sometimes the choice of words (e.g., "remove" instead of "disable") can make the difference between a right answer and a wrong one.

Cisco Certification Program

The Cisco Certification Program currently includes the following separate certificates with various specialty tracks. You should become familiar with and visit regularly Cisco's website at www.cisco.com/go/certifications/.

Cisco reserves the right to change the number of questions and time limits for the exams as it sees fit. Cisco tries to keep this information confidential, although you can check either figure when you register for an exam. The http://studyguides.cramsession.com/ site usually has pretty reliable information about number of questions and time limits.

Receiving Your Certificate

After passing the necessary certification exam(s) and agreeing to Cisco's nondisclosure terms, you will be certified. Official certification normally takes from four to six weeks. The package includes a welcome kit that contains a number of elements:

- Official certificate (suitable for framing)
- A laminated wallet card

- A graduation letter
- A license to use the Cisco certification logo, in advertisements, promotions, documents, resumes, letterhead, business cards, and so on.
- Access to the online Tracking System

Tracking Cisco Certification Status

As soon as you pass any Cisco exam, you must complete a certification agreement. To do this, go to Cisco's Web site www.cisco.com/go/certifications/ and select the Tracking System link. You can also mail a hard copy of the agreement to Cisco's certification authority. You will not be certified until you complete a certification agreement and Cisco receives it in one of these forms.

The Certification Tracking Web site also allows you to view your certification information. Cisco will contact you via email and explain your certification and its use.

Recertification

Cisco requires three-year recertification for the non-CCIE programs. The best place to keep tabs on the Cisco Career Certifications program and its related requirements is on the Web. The URL for the program is www.cisco.com/go/certifications/.

PART I

Introduction to Network Security

- **Chapter 1** Understanding Network Security Threats
- **Chapter 2** Securing the Network

CHAPTER

Understanding Network Security Threats

In this chapter, you will learn to:

- Identify the need for network security
- Recognize the causes of network security problems
- Distinguish the four primary types of threats
- Know the four primary types of network attack
- Discover Cisco AVVID and SAFE, and how they relate to network security
- Learn about the Cisco Security Wheel
- Understand network security policy
- Improve network security

To understand, in part, why we are where we are today, you only have to remember that PC is the acronym for personal computer. The PC was born and, for many years, evolved as the tool of the individual. In fact, much of the early interest and growth came as a rebellion to what appeared as exclusionary attitudes and many restrictions of early data-processing departments. Admittedly, many PCs were tethered to company networks, but even then there was often considerable flexibility in software selection, settings preferences, and even sharing of resources such as folders and printers.

As a result, a huge industry of producers developed and sold devices, software, and services targeted at meeting user interests and needs, often with little or no thought about security. Prior to the Internet, a person could keep their computer resources safe simply by being careful about shared floppy disks.

Today, even the PCs of most individuals routinely connect to the largest network in the world (the Internet) to expand the user's reach and abilities. As the computing world grew, and skills and technology proliferated, people with less than honorable intentions discovered new and more powerful ways to apply their craft. Just as a gun makes a robber a greater threat, computers give the scam artist, terrorist, thief, or pervert the opportunity to reach out and hurt others in greater numbers and from longer distances.

This book provides a variety of techniques and technologies to protect computing resources from unauthorized access and loss. This chapter lays the foundation by looking at the need for network security. What are the threats? Who are these people who

Author's Bias

One of the reasons hacking and other forms of network intrusions occur so often is because too many people inside and outside the industry think something is special about computer crime. A mystique surrounds some activities. Some even go so far as to create colorful terms, such as "ethical hackers" or "white hat hackers." The bottom line is this: the person who gains *unauthorized access to another's computer* is no less a criminal than the burglar who gains access to your home. Web site hackers are no more honorable or deserving of special treatment than any other vandal, regardless of their cause or motivation. Once you own, or work for, a company that's had to waste the equivalent of many annual salaries to defend against attacks, fight off an attack, or restore damaged resources, the "victimless" rational of computer crime goes up in smoke.

threaten the data, and what are some of the methods they use? In addition, you'll find many references to outside resources for additional information.

While this book addresses the requirements of the various certification exams, recognizing that the diversity of security threats is far too large for any single book is important. Furthermore, the nature and source of many threats changes on a daily basis, making it important to start building a set of resources, such as web sites, news groups, trade associations, vendor distribution lists, and so forth that can help you try to stay abreast of the changes. Each technology, such as wireless, voice, web pages, and e-mail systems, has its own set of threats that a person must remain aware of.

Identify the Need for Network Security

Pointing at the Internet and indicating that as the point in time when security had to become a part of everyone's computing strategy is easy. Business and individuals alike were faced with protecting their computing resources from the many possible dangers that lurked in the Net. The Internet opened a large door onto a busy street filled with seemingly unlimited commercial and intellectual opportunities. Unfortunately, within that busy street reside the same opportunists we fear in our noncyber lives.

Another way the Internet impacts security is its worldwide reach as a reference library for security experts and, unfortunately, the hacker community as well. In a few minutes, a search for hack, crack, phreak, or spam yields many sites, some with many links to other links.

But blaming the Internet is somewhat unfair. The Internet simply happened to be the first attractive new service with strong mass appeal that brought with it significant security risks. Others that followed include wireless communications and connectivity, instant messaging, and enhanced e-mail services, and undoubtedly more will follow. Increased security awareness and implementation is, by necessity, one of the prices that must be paid for new services that connect people.

Unfortunately, all organizations aren't alike and, therefore, a one-plan-fits-all approach to security won't work. Many factors—from internal company policies to topologies and services supported—impact the decisions about the proper security strategy. Even within an organization, the security requirements can require many different solutions. A single LAN branch location has different security issues than a WAN link or a campus VLAN environment.

Even after the organization assesses its security risks and starts to develop a plan, problems often exist in knowing whether various multivendor tools will work together and be supportable in the long term. One common problem with any multivendor environment (not only networking) is the inevitable finger-pointing when things go wrong. So often, a decision about single vendor or multivendor solutions must be made. Cisco is a big believer in single-vendor, end-to-end solutions—the company was built through acquisitions and R&D to that end, but it's also a solid supporter of standards-based technologies. Standards-based solutions can at least reduce some of the interoperability issues involved in a multivendor solution.

Cisco network and security products are developed under Cisco's AVVID and SAFE strategies to ensure solid standards-based implementations. Both strategies are covered later in this chapter in the "Cisco AVVID and SAFE Strategies" section.

NOTE Multivendor implementations require more than just knowing that the technologies will work together. There can also be a significant support commitment and cost in maintaining resident experts on multiple vendor products. In addition to having to know how to install and provide production support, someone must be a security expert on each vendor line to keep on top of security announcements, vulnerabilities, patches, upgrades, and so forth. The future can change the balance completely. While products from two vendors might "play well together" initially, what happens in the future when a new technology develops and one vendor chooses a standards-based approach while the other chooses a proprietary solution, or maybe not to play at all?

Identify the Causes of Network Security Problems

While many causes exist for security problems, at least three types of fundamental weaknesses open the door to security problems.

- Technology weakness
- Policy weakness
- Configuration weakness

Obviously, we could probably add human weakness and some others, but our purpose is to concentrate on those issues that, once recognized, can be managed, monitored, and improved within a security strategy.

Technology Weakness

Every technology has some known or unknown inherent weaknesses, or vulnerabilities that can be exploited by a sufficiently motivated troublemaker. Some weaknesses are publicized widely in the media because they're associated with a well-known product. Don't fall into the faulty logic that because you don't hear about the other products, they must be secure. Just because no one cares enough to hack a product, doesn't mean it's necessarily secure.

TCP/IP Wasn't Designed for Security

Starting right at the top, TCP/IP wasn't designed with security as a high priority. One of the drawbacks to being the first at anything is the inability to see how others might manipulate and transform a technology into something else. The designers were looking for a reliable vehicle to allow research organizations to share information. The many early protocols and tools that make up the TCP/IP suite were developed in an environment of trust and openness.

Today, various Request for Comments (RFCs), security best practices, security services, and an array of products from many vendors work together to reduce the risks inherent in the environment.

Computer and Network Operating Systems

Regardless of the manufacturer or whether it's an open standard or proprietary product, every operating system (OS) has vulnerabilities that need to be addressed through patches, upgrades, and best practices. Every time a major upgrade comes out, the possibility for new or even revived vulnerabilities can, and does, appear.

While a company tries to produce and deliver a secure final product, the addition of new features, implementation of new standards, and even hardware changes can lead to potential problems that don't get caught in prerelease testing.

Given the number of lines of code in most modern OSs, it isn't wholly unreasonable that some problems will slip through. While our focus is security, the OS developers and product testers are looking at usability, accessibility, features, performance, stability, backward compatibility, and many other characteristics, plus security. Right or wrong, it's also important to remember that security hasn't always been the highest priority of developers, product managers, customers, product reviewers, financial analysts, writers, and so forth.

Network Device Weaknesses

Whether IOS based or embedded in the circuitry, such as application-specific integrated circuit (ASIC), network devices can have vulnerabilities, often called "holes," that can be exploited. Some might lay dormant for years until someone stumbles across one, and either exploits it or documents it. Often the process of documenting and notifying the user base of a problem lays out a roadmap to troublemakers.

When possible patches, IOS upgrades, and best practices should be applied to eliminate or mitigate known problems. In some cases, it might be determined that the device

should be abandoned or moved to a part of the network that would be impacted less by the problem.

To find security advisories and related information without a CCO ID, go to http://www.cisco.com and do a search on security.

Policy Weakness

Policy weakness is a catchall phrase for company policies, or a lack of policies, that inadvertently lead to security threats to the network system. Chapter 2 covers in detail the importance and implementation of a written security policy, which is the essential foundation of a good security implementation.

The following examples are some of the policy issues that can negatively impact a businesses computer system:

- **No written security policy** Lack of a documented and adopted plan means the security efforts evolve and are enforced, if at all, in a best-effort manner.
- **Lack of disaster recover plan** Without a plan, the efforts to fight a network attack—or even a physical emergency such as fire, flood, or earthquake—are left to the judgment and knowledge of the staff on hand. Even the best-trained and most experienced staff can make foolish decisions when faced with an unexpected catastrophic event.
- **No policy for software and hardware additions or changes** Whether motivated by increasing productivity or recreation, any addition or upgrade to software or hardware can introduce unexpected security vulnerabilities. Adding an unauthorized wireless access point to a network can throw open a virtual garage door to the network and the company resources. Similarly, an unauthorized screensaver might also be harvesting passwords, user IDs, and other information for someone else.
- **Lack of security monitoring** Even if a secure network is developed, failure to monitor logs and processes or weak auditing allows new vulnerabilities and unauthorized use to evolve and proliferate. The worst case would be not recognizing that a serious loss had occurred or was continuing.
- **Employment policies** Frequent staff turnover, lower than typical compensation, and lack of training opportunities can all impact network security by bringing new untested and underskilled employees into positions of authority and responsibility.
- **Internal policies** Lax business attitudes and practices often create temptations and a relatively safe environment for the opportunist within to ply their craft. This is the "we are all like family here" syndrome. Unfortunately, even some of the best families have a thief in their midst. Similarly, infighting, backbiting, power struggles, or turf struggles can lead to security issues or divert attention, allowing problems to go undetected.

Configuration Weakness

Many network devices have default settings that emphasize performance or ease of installation without regard for security issues. Installation without adequate attention to correcting these settings could create serious potential problems. Some common configuration issues include the following:

- Ineffective access control lists failing to block intended traffic
- Default, missing, or old passwords
- Unneeded ports or services left active
- User IDs and passwords exchanged in clear text
- Weak or unprotected remote access through the Internet or dial-up services

Monitoring vendor announcements and advisories, combined with industry news services, can identify the most common, best-known vulnerabilities and often include the appropriate mitigation solution.

STUDY TIP Know the three causes of security problems.

The Four Primary Types of Network Threats

In an attempt to categorize threats both to understand them better and to help in planning ways to resist them, the following four categories are typically used.

- Unstructured threats
- Structured threats
- Internal threats
- External threats

Unstructured Threats

Unstructured threats often involve unfocused assaults on one or more network systems, often by individuals with limited or developing skills. The systems being attacked and infected are probably unknown to the perpetrator. These attacks are often the result of people with limited integrity and too much time on their hands. Malicious intent might or might not exist, but there is always indifference to the resulting damage caused to others.

The Internet has many sites where the curious can select program codes, such as a virus, worm, or Trojan horse, often with instructions that can be modified or redistributed as is. In all cases, these items are small programs written by a human being. They aren't alive and they can't evolve spontaneously from nothing. Some common terms to be aware of include the following:

Virus	A program capable of replicating with little or no user intervention, and the replicated programs also replicate.
Worm	A form of virus that spreads by creating duplicates of itself on other drives, systems, or networks. A worm working with an e-mail system can mail copies of itself to every address in the e-mail system address book. Code Red and Nimda are examples of high-profile worms that have caused significant damage in recent years.
Trojan horse	An apparently useful or amusing program, possibly a game or screensaver, but in the background it could be performing other tasks, such as deleting or changing data, or capturing passwords or keystrokes. A true Trojan horse isn't technically a virus because it doesn't replicate itself.

The person launching an unstructured attack is often referred to as a *script kiddy* because that person often lacks the skills to develop the threat themselves, but can pass it on anonymously (they think) and gain some perverse sense of satisfaction from the result. E-mail delivery methods have replaced "shared" game disks as the vehicle of choice for distributing this type of attack.

NOTE The term "script kiddy" is a common derogatory term and should be used with caution, if at all. Script kiddy is included here so you know what it means. Remember, the difference between an unstructured attack and a series of all-out denial-of-service attacks might be that the latter attacker is offended or angry.

Unstructured attacks involving code that reproduces itself and mails a copy to everyone in the person's e-mail address book can easily circle the globe in a few hours, causing problems for networks and individuals all over the world. While the original intent might have been more thoughtless than malicious, the result can be a loss of user access while systems are being protected, a loss of reputation if the news that a company's site has been attacked, or a loss of user freedoms as more-restrictive policies and practices are implemented to defend against additional attacks.

In some organizations, if the network is down, entire groups of people can't do their jobs, so they're either sent home or they sit and wait without pay because their income is tied to sales. So even if the hacker "thought" no one would be hurt, the result is often that they just beat some single parent or new hire out of a day's pay.

Each of these results can be quantified in currency and often result in large numbers if and when the perpetrator is prosecuted.

Structured Threats

Structured threats are more focused by one or more individuals with higher-level skills actively working to compromise a system. The targeted system could have been detected through some random search process, or it might have been selected specifically. The attackers are typically knowledgeable about network designs, security, access procedures, and hacking tools, and they have the ability to create scripts or applications to further their objectives.

Structured attacks are more likely to be motivated by something other than curiosity or showing off to one's peers. Greed, politics, racism (or any intolerance), or law enforcement (ironic) could all be motives behind the efforts. Crimes of all types where the payoff isn't directly tied to the attack, such as identity theft or credit card information theft, are also motivations.

International terrorism and government-sponsored attacks on another country's computer infrastructure are becoming well documented. Systems of interest might include utilities, public safety, transportation systems, financial systems, or defense systems, which are all managed by large data systems, each with vulnerabilities.

Internal Threats

Internal threats originate from individuals who have or have had authorized access to the network. This could be a disgruntled employee, an opportunistic employee, or an unhappy past employee whose access is still active. In the case of a past network employee, even if their account is gone, they could be using a compromised account or one they set up before leaving for just this purpose.

Many surveys and studies show that internal attacks can be significant in both the number and the size of any losses. If dishonest employees steal inventory or petty cash, or set up elaborate paper-invoicing schemes, why wouldn't they learn to use the computer systems to further their ambitions? With access to the right systems, a trusted employee can devastate an unsuspecting organization.

All too often, employers fail to prosecute this type of activity. The reasons range from fear of the activity becoming public knowledge to knowing that, quite often, record-keeping systems haven't been developed either to provide adequate evidence or to prove that the transactions, no matter how ludicrous, weren't authorized.

NOTE I was helping a dentist reconstruct a substantial loss by an office manager when we were all served with papers threatening all sorts of repercussions if we spoke to anyone, including the police, about the matter. Because less than three days had passed since the loss was inadvertently exposed, I was shocked at the coolness and speed of the reaction. After a little research, I found this was at least the third dentist in seven years who had been scammed by the same person. The bottom line is that the bonding company and the dentist came to terms, and I never heard another word about it.

External Threats

External threats are threats from individuals outside the organization, often using the Internet or dial-up access. These attackers don't have authorized access to the systems.

In trying to categorize a specific threat, the result could possibly be a combination of two or more threats. The attack might be structured from an external source, but a serious crime might have one or more compromised employees on the inside actively furthering the endeavor.

STUDY TIP Be sure to know the four primary types of threats. They could appear on all four exams.

The Four Primary Types of Network Attack

While there are many variations and often different names, the four most common types of network attacks are

- Reconnaissance attacks
- Access attacks
- Denial-of-service attacks
- Data manipulation attacks

STUDY TIP Some texts and certification exams might consider only the first three as specific types of network attacks, with data manipulation being a variety of access attack.

Reconnaissance Attacks

A *reconnaissance attack,* as the name implies, is the efforts of an unauthorized user to gain as much information about the network as possible before launching other more serious types of attacks. Quite often, the reconnaissance attack is implemented by using readily available information.

Public Information

Employee names and e-mail addresses provide a good start in guessing the user name for an employee's account. Common practice is to use an employee's first initial and last name as the user name for their network computer account. E-mail addresses are also a common user name for computer accounts. Large companies usually have their phone numbers assigned in blocks from the local telephone company, and many large corporations have their own dialing prefix. By using this information, the intruder can begin war dialing all the company phone numbers looking for a dial-up server. Once a dial-up server is found, the intruder can begin guessing account user names based on an employee's first initial and last name or their e-mail addresses. Brute force password crackers are freely available on the Internet. Once a user name has been guessed, it's only a matter of time before a weak password can be cracked.

A *war dialer* is a program used to dial blocks of phone numbers until it finds a computer on the other end of the line. Once a computer is found, the war dialer application records the number dialed for later use by the intruder.

To use a user account on a server or a network, you must first have the user name and password. Discovering the user names is a fairly straightforward process described in the preceding paragraph. Attackers use *password crackers* to crack the passwords to user

accounts. Some password crackers find the encrypted password files on the server and decrypt them. When a hacker is unable to retrieve the password files, then brute force password crackers are used. *Brute force password crackers* attempt to log in to a computer account over and over, using multiple password combinations. Some cracking software uses dictionary files, while others attempt every combination of each key on the keyboard—a time-consuming ordeal.

The following are commonly used password crackers:

Microsoft Windows	UNIX
L0phtCrack 4	Qcrack by the Crypt Keeper
PWLVIEW	CrackerJack by Jackal
Pwlhack 4.10	John the Ripper by Solar Designer
PWL-Key	Crack by Alec Muffet
ntPassword	

Internet Protocol (IP) address information is publicly available via the ARIN and many other Internet registering authorities. From www.arin.net, anyone can begin a search using a single known IP address. The search will yield the complete block of IP addresses belonging to the company. Domain Naming Systems (DNS) is another publicly available system that can provide a wealth of information regarding the IP addressing and naming strategies of virtually any company connected to the Internet.

For a company to host its own e-mail, web, ftp, or any other service on the Internet, it must first have each of these servers listed within the DNS infrastructure. These DNS servers list the names of the servers, along with the IP addresses that can be used to access these services. To mitigate these risks, security conscious companies could choose to host these servers and services outside their private networks with a hosting company. This added security is usually rendered obsolete, however, by adding backend connections from the hosting facilities back to their private networks.

Electronic Reconnaissance

The attacker must perform electronic reconnaissance to find what systems and resources are on the network. Unless the attacker has prior knowledge of the target network, he or she must find where the company resources are logically located. Once the company IP addresses are known (see the preceding section, "Public Information"), the attacker can begin to probe and scan the network. The intruder can scan the network looking for vulnerable hosts, applications, or infrastructure equipment.

Scanning the network is typically done using a *ping sweep* utility that pings a range of IP addresses. The purpose of this scanning is to find what hosts are currently live on the network. The ping sweep identifies viable targets on the network. Once the IP address of viable hosts is known, the attacker can then begin to probe those hosts to gather additional information, such as the OS or applications running on those hosts.

Probing is attempting to discover information about the hosts that are on the network. Probing is accomplished by looking for open ports on the available host computers. *Ports* are like virtual doorways to the computer. For a computer to offer or use services on the network, it must first have an open port. Web servers typically use port 80, while FTP servers use port 21. An attacker can find out what services are running on a computer by discovering what ports that computer has opened.

TCP/IP uses port addresses to locate services running on host computers. The port numbers used by an application are that application's address on that host. The address for a web application located on host 10.0.0.1 would be 10.0.0.1:80. This address specifies the host address 10.0.0.1 and the application address of 80. Most common applications use well-defined port numbers. A list of well-known port numbers managed by the Internet Assigned Number Authority (IANA) can be viewed at http://www.iana.org/assignments/port-numbers.

The more ports that are open, the more potential for someone to exploit the services running on the host computer. Once the attacker knows which ports are open, he/or she can use this information further to discover the OS and the application servicing the port.

The purpose of this scanning and probing is to find weaknesses on the network. Intruders know the vulnerabilities of certain OSs and the applications they run. The intruder increases his or her chance of succeeding by finding the weakest point on the network and later attacking that vulnerability. The attacker continues to discover information about the network until they have a complete map of the hosts, servers, and weaknesses to exploit in the future.

Reconnaissance Tools

The most common and widely used hacking tools are *reconnaissance tools*. Many of these tools have been developed by hackers to aid them in their illicit activities. Other tools used by hackers are the same tools commonly used by network engineers to view problems on the network.

As security and intrusion detection have gotten more sophisticated, so has the software used by hackers. Intrusion-detection software looks for people looking at the network. Hackers know that scanning and probing a network is likely to create suspicion and might generate alarms. Because of this, hackers have begun to develop new software that attempts to hide the true purpose of its activity. Reconnaissance tools in common use today include the following:

NMAP	WHOIS
SATAN	Ping
Portscanner	Nslookup
Strobe	Trace

Access Attacks

Access attack is a catch-all phrase to encompass a variety of forms of unauthorized access of computer resources. An access attack could be an outside individual, or a group that uses various methods to gain entry to a network and, from there, steals confidential information or engages in destruction of resources. An access attack could also be an inside (trusted) user getting into areas they aren't authorized to use. Their intentions could be curiosity or the same as the outside hackers.

Gaining Initial Access

In many cases, the first objective is to gain initial access, so additional reconnaissance can be conducted. This reconnaissance could include scouting out resources, IP addresses, and possibly running a network discovery (mapping) program or even a sniffer-type packet-capturing utility, hoping to capture administrative-level passwords.

War dialers can be used to dial a large number of phone numbers looking for modems. A new variation involves sitting in a parking lot or in a building across the street with a laptop and a wireless NIC, looking for unsecured or poorly secured access points.

Again, don't overlook the person on the inside who already has access through an authorized user name and password. Whether connecting from outside or from an inside host, they have the first hurdles resolved.

Social Engineering

The term *social engineering* relative to security came from early hacking efforts on telephone systems and long-distance services. Social engineering is based on the concept of why risk breaking into a system by brute force or tools when you can get some friendly employee to help you do it? Social engineering is generally a hacker's clever manipulation of an employee's natural human tendencies to trust and want to be helpful.

More than one company with elaborate authentication processes, firewalls, virtual private networks (VPNs), and network monitoring software has been left wide open to an attack by an employee unwittingly giving away key information in an e-mail or by answering questions over the phone with someone they don't know. This is one area where the would-be hacker can benefit from a friendly demeanor, a good smile, and knowledge of looking and acting like they belong.

Don't make the mistake of thinking only lower-level employees are prone to this. The fear of appearing not to cooperate with an obviously important activity has led to the comprise of many a manager.

Password-Based Attacks

To use a user account on a server or network, you must first have the user name and password. Discovering the user names is a fairly straightforward process described in the preceding section. Attackers use password crackers to crack the passwords to user accounts. Some password crackers find the encrypted password files on the server and decrypt them. When a hacker is unable to retrieve the password files, then brute force password

crackers are used. Brute force password crackers attempt to log in to a computer account over and over using multiple password combinations. Some cracking software uses dictionary files, while others attempt every combination of each key on the keyboard, a time-consuming ordeal.

Commonly used password crackers include the following:

Microsoft Windows	UNIX
L0phtCrack 4	Qcrack by the Crypt Keeper
PWLVIEW	CrackerJack by Jackal
Pwlhack 4.10	John the Ripper by Solar Designer
PWL-Key	Crack by Alec Muffet
ntPassword	

A good password system locks the account after a limited number of tries to thwart this type of attack. The successful hacker has the same access to resources as the users whose accounts they compromised to gain access to those resources.

General password security lapses can put a password in the hands of an intruder. This can be something as simple as passwords written on a desk pad, an appointment calendar, or an address book, to gaining access to a person's home or laptop computer where the logon password is being remembered by the OS. More than one company's security has been compromised by a child accessing the system from home or a friend's house using a password appropriated from a parent.

One-time passwords (OTP) systems and/or cryptographic authentication can almost eliminate the threat of password attacks. OTPs involve using "something you have," such as password-token generator software on your computer, plus something "you know," such as a PIN number. The token software uses the PIN to generate what appears as a unique password. Once the token is used, it won't work again, thwarting the intruder with a sniffer product.

If standard passwords must be used, strong passwords—those that would be difficult to guess—can help. Strong passwords should be at least eight characters long and contain both uppercase and lowercase letters, numbers, and special characters (such as 23!!pandA). While randomly generated passwords might be the best, they're hard to remember and often lead users to write them down.

Gaining Trusted or Privileged Access

Once initial access has been accomplished, the hacker will attempt to exploit any privileges associated with that access, including the ability to get into shared resources. If the initial account has limited access permissions, the hacker will try to gain administrator privileges (root inUNIX). With the higher privileges, the hacker can expand their influence by creating additional accounts they have access to, clean up any logs or history of their activities, and install additional tools for reconnaissance.

Denial of Service (DoS) Attacks

Denial of service (DoS) attacks in their many forms are by far the most infamous, and possibly the most threatening to organizations who conduct any business over the Internet. The primary purpose of any DoS attack is to deny access to a device—or better, an entire network—by bombarding it with useless traffic. This attack has two ways to bury the target. First, the packets themselves can consume 100 percent of a device's resources, thereby preventing it from doing its regular work. Because a firewall or intrusion detection system could often easily defeat this type of attack, the second threat is far greater. The second threat is that the organization's connection(s) to the Internet is filled to capacity with this useless traffic, thereby preventing in or out communications. For this reason, a DoS attack typically can only be defeated by the efforts of the organization's ISP.

Because the ISP's upstream connection, called a *fat pipe,* is typically many times larger than the connection to each customer, the ISP might be completely oblivious to the attack. If the ISP's staff and service policies are less than optimal, the organization under attack might seem doomed. Figure 1-1 shows the relative capacity of the ISP's link to the Internet versus the much smaller links to their customers.

The true DoS attack launched by a single host generally isn't used, except by the least-experienced hackers. Figure 1-2 shows a traditional DoS attack. The two most devastating variations are the distributed denial of service (DDoS) and the distributed deflection denial of service (DRDoS). Both of these attacks enlist the assistance of others, often hundreds, of unsuspecting hosts to assist in the attack, thereby significantly increasing the size of the attack, further shielding the source, and making it harder to defend against.

DDos

DDoS attacks start by the attacker(s) placing Zombie (technically, "bot," short for "robot") programs in a series of compromised computers hooked by relatively high-bandwidth connections to the Internet. These Zombies are programmed to monitor specific Internet Relay Chat (IRC) chat rooms to receive further instructions. The

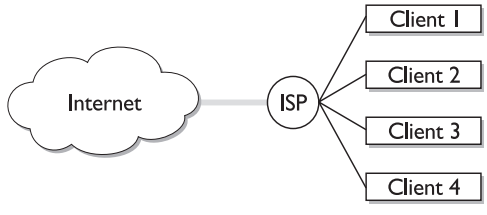

Figure 1-1 Bandwidth comparison for ISP to client links vs. ISP upstream links

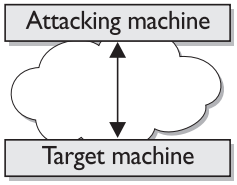

Figure 1-2 DoS attack with a single attacker and a single target

Zombie attack is directed and coordinated by a Zombie Master, who sends instructions to the individual Zombie, who then begins generating a flood of malicious traffic aimed at the target. Figure 1-3 shows a DDoS attack.

Early DoS attacks on some famous web sites involved many computers on university campuses and even some from security agencies. These computers had unprotected security holes, were online around the clock, and provided large connections to the Internet. Today, DSL and cable modem connections make many home and small business computers more attractive as Zombie sites because they often lack the security features and staff to defend against the intrusion.

Some Zombies, once in place, download and install additional applications that can map the local network, capture passwords or keystrokes, and report findings to the instigators of the attacks.

DRDoS

The latest variation on the DoS, the *DRDoS*, involves one or more hosts sending a series of TCP SYN requests or ICMP ping requests to many unsuspecting, even thoroughly secure, hosts using the "spoofed" source address of the target. When these hosts respond to what appears to be a legitimate, nonthreatening request, they collectively create an unsupportable flood of packets aimed at the target. Figure 1-4 shows a DRDoS attack. Again, even if the target device(s) can determine what's happening, only a cooperative ISP can block the traffic before it buries the target's Internet connection.

If the originating source continues to vary the type of packets sent to the reflectors, the filters at the ISP have only temporary or limited usefulness before they need to be changed.

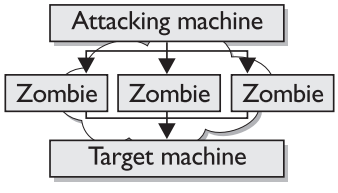

Figure 1-3 DDoS attack involving Zombie remote hosts

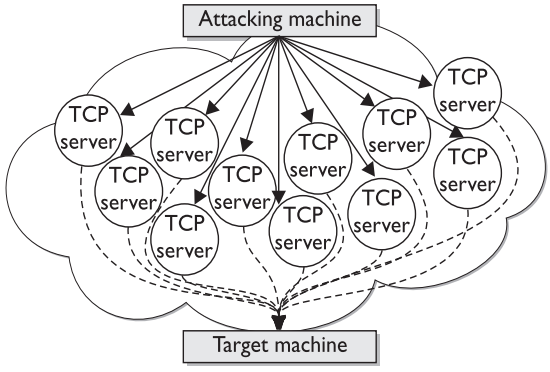

Figure 1-4 DRDoS attack showing the interim hosts

Well-Known DoS Attacks

Knowing about common, well-known attacks can be useful and interesting, and when someone indicates an attack is a variation of the Ping of Death, you will know what that means. Well-known attacks include the following:

- **TCP SYN Flood** Uses the TCP establishment handshake to conduct attacks by creating TCP "half-open" connections, tricking the target or reflector into thinking a session is being established.
- **Ping of Death** Sends one or more oversized ping packets to crash or disable servers and other computer systems. Sending illegal IP datagrams (larger than 65,536 bytes) is possible because of packet fragmentation during transmission. When the fragments are reassembled at the target, it can overflow the buffer and cause a reboot, crash, or hang.
- **Trinoo** A distributed tool (bot) used to launch coordinated UDP flood DoS attacks from many sources. A *Trinoo* network consists of a small number of masters and a large number of bots.
- **Tribe Flood Network (TFN) and Tribe Flood Network 2000 (TFN2K)** Like Trinoo, variations of TFN use a distributed tool to launch coordinated DoS attacks from many sources against the target(s), often using spoofed source IP addresses. TFN bots can generate UDP flood attacks, TCP SYN flood, ICMP echo request flood, and ICMP directed broadcast (for example, smurf) DoS attacks.
- **Stacheldraht** (German for "barbed wire") Combines features of the Trinoo DDoS tool with those of the original TFN, and adds encrypted communications between the attacker and stacheldraht masters and automated agent updates.
- **Trinity** Preys on Linux servers and uses IRC channels to unleash IP packet floods on targeted host machines

Terrorism, Act of War, and Legal Implications

Variations of the DoS attack are likely to be a major component of global terrorism and even a part of government-sponsored acts of aggression against its perceived enemies. The possible devastating that effect a massive distributed attack could have on a country's command and control systems, financial systems, utility grids (power, telephone, transportation, and soon), and other services is something to remember.

While we hope most of us will never be in a position to directly defend against such an attack, it's critical that the resources under our control do not become unwitting hosts to any kind of DDoS attack.

Many lawyers became quite computer savvy and versed in the areas of financial responsibility in preparation for the feeding frenzy they expected from the Year 2000 "bugs." While that threat never materialized, many lawyers are now advising victims of DDoS attacks that the unwitting hosts of the attack bots might have financial liability because of not detecting and eliminating the devices. Security practices that allow these unauthorized residents to do their dirty deeds could carry a hefty price tag.

While the original hacker is careful to conceal their identity and address, they have a whole lot less interest in protecting the bot hosts. If, in fact, lawsuits against remote sites become common, it's not inconceivable that the bot site might be the ultimate target of an attack.

Motivation and Good Sense

While many reasons or rationalizations exist that an individual or group of individuals might choose to launch a form of DoS attack on a network, one thing common to many attacks is anger. Real or imagined, the attacker blames the site, the owners of the site, or the users of the site for some slight, injustice, or wrong doing. Add to this the apparent anonymity the attacker enjoys, and it's generally a no-win situation to provoke or even incur the interest of these individuals needlessly. The size and scope of the Internet means your site can literally fall prey to a "sniper" 12,000 miles way.

Attackers typically have the time, and the cost to them is close to zero. The target is in the opposite position: once the attacks begin, time is virtually nonexistent. The costs, direct and in lost business or reputation, start to soar. Be well aware that no Internet Police Department or anyone else is going to handle this for you.

Don't make yourself a target. Practice good security measures and involve law enforcement in all criminal acts, but be forewarned that personal attacks and even belittling statements like script kiddies might precipitate a career of fighting these attacks. Sometimes, even protective security measures as a result of an attack within your network might escalate the attack.

Time isn't as universal as many of us think. When a network is under attack and the administrator has brought in all the high-priced talent, added new technologies, and possibly even lined up law enforcement, it's common to want the attack to continue long enough to identify and catch the attacker. Remember, other than anger or adrenalin, the hacker has nothing invested and could recognize they can even cause greater losses by being unpredictable. In most cases, the worst that can happen is that the hacker gets locked out.

Techniques to Counteract DoS Attacks

While the threat of DoS attacks can't be eliminated, it can be reduced through the following three methods:

- **Anti-DoS features** Proper implementation and configuration of anti-DoS features available on routers and firewalls can help limit the effectiveness of an attack. These features could include limiting the number of half-open connections allowed at any given time or limiting the number of certain types that can originate from a source address.
- **Antispoofing features** Proper implementation and configuration of antispoofing features on routers and firewalls can help limit a hacker's ability to mask their identity. RFC 2827 filtering should be configured at a minimum (see the upcoming section "IP Spoofing").
- **ISP traffic rate limiting** The ISP agrees to filtering limits on the amount of nonessential traffic that can cross link(s) to the company at one time. The filtering might limit the volume of ICMP traffic, a common source of distributed denial of service (DDoS) attacks, into a network because it's used only for diagnostic purposes.

Data Manipulation Attacks

Data manipulation, or impersonation, is made possible by vulnerabilities in IP protocols and related applications. *Data manipulation attacks* are often called "man-in-the-middle" attacks because the attacks typically involve an individual located between TCP/IP-exploited IP vulnerabilities. Common forms of these attacks include IP spoofing, session replay, session hijacking, rerouting, repudiation, and vandalizing web pages.

IP Spoofing

An *IP spoofing attack* involves an external or internal hacker who pretends to be using a trusted computer by using the address of that computer. The hacker either uses an IP address within the range of trusted internal addresses for the network or an authorized external address that's both trusted and allowed access specified network resources. IP spoofing is often a tool used as part of other attacks, such as any variation of DoS attack, to hide the hacker's identity.

IP spoofing is often limited to the introduction of malicious data or commands into an existing data stream in a peer-to-peer network session. Spoofing a source address might enable data to be sent through a router interface with filtering based on the source address.

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

- **RFC 2827 filtering** Basically, *RFC 2827 filtering* means filtering out any IP addresses from coming into a network segment that should already be on that segment. If the entire 195.17.1.0 network is attached to a router interface, then no legitimate packets with source addresses in that network should be coming in through the router. This should be applied to edge routers for sure, but it can also be used on internal routers to prevent spoofing within the network. Similarly, limiting any outbound packets leaving the network to ones that have source addresses assigned to that network can prevent a network's hosts from spoofing other networks. This could be the result of an attacker on the inside or a DoS bot on a local host participating in an attack on an outside host. If the company can get its ISP to perform RFC 2827 filtering on packets coming into the network, it would preserve the bandwidth of the link and kill some hack attempts.

NOTE Spoofing could be virtually eliminated if all ISPs filtered client traffic to allow only source addresses assigned to that client. If hackers can't spoof it, this makes going undetected harder.

- **RFC 1918 filtering** *RFC 1918 filtering* means filtering out RFC-defined "private" addresses from entering or exiting the network segment. Because they have no business on the Internet, they shouldn't be there. If private addresses are used in the network, RFC 2827 filtering will include them.
- **Non-IP address authentication** IP spoofing is worthwhile when devices use IP address–based authentication. If you use additional authentication methods, IP spoofing attacks lose much of their value. *Cryptographic authentication* is the strongest form of additional authentication, but if this isn't possible, use strong, two-factor authentication, such as OTP.

Session Replay and Hijacking

Session replay is a form of a man-in-the-middle attack, where the intruder captures a packet sequence and modifies part of the data before forwarding it on normally. This type of attack relies on an inherent weakness in data traffic authentication.

Session hijacking is a form of a man-in-the-middle attack where the attacker takes over an IP session that's underway by spoofing source and/or destination addressing and altering TCP sequence numbering. Typically, a packet sniffer is used to set up the hijacking by allowing the user to see the existing traffic.

Rerouting

Rerouting involves either gaining access to a router to change the route table entries, or spoofing the identity of routers or hosts so traffic is directed to a compromised device. Spoofing ARP replies is even possible. It causes a host to forward packets intended for a specific host or the default gateway to be forwarded instead to another local host. The new destination host can perform its assigned task and then forward the packet on to the correct destination.

Repudiation

Repudiation is the denial of having been a part of a data exchange. This repudiation might be to avoid responsibility for an action. *Nonrepudiation* is a security feature that helps ensure that data has been sent and received by the parties claiming to have sent and received it. Nonrepudiation guarantees that the sender of a message can't later deny (repudiate) having sent the message. Similarly, the recipient can't deny having received the message.

Methods for implementing nonrepudiation include the following:

- **Digital signatures** Unique identifier for an individual, much like a written signature
- **Confirmation services** The message transfer agent creates digital receipts indicating messages were sent and/or received
- **Timestamps** The date and time a document was composed, proving a document existed at a certain time

Cisco AVVID and SAFE Strategies

Cisco Architecture for Voice, Video, and Integrated Data (AVVID) and SAFE are Cisco-comprehensive strategies that help organizations successfully and securely develop and implement end-to-end network designs.

AVVID

The approaching convergence of the telephone services, videoconferencing, IP data networks, software-based services, and a never-ending supply of new technologies introduces many opportunities for both disaster and success. AVVID provides a standards-based network architecture and a comprehensive set of best practices, which allows businesses to develop business and technology strategies that scale to meet the changing demands of e-business. AVVID provides end-to-end networking solutions that help organizations plan rapid deployment of emerging technologies and new Internet-based solutions.

The AVVID end-to-end networking solution includes network client devices, network infrastructure from network platforms to intelligent network services, Internet middleware tools, systems integrator responsibilities, and Internet business solutions. The AVVID strategy addresses the three primary concerns of network deployment: performance, scalability, and availability.

These documents and related links are available free, without a CCO ID from the Cisco web site. To find the current list of documents, either go to the following web site or go to http://www.cisco.com and do a search on AVVID: http://www.cisco.com/warp/public/779/largeent/avvid/cisco_avvid.html.

TIP For the exams and for your own development in the industry, go to the site and at least download the white paper on the AVVID architecture and become familiar with it. This white paper provides a detailed overview that you can use to supplement all the chapters in this book.

SAFE

Cisco's strategy for secure networks (SAFE) started with the original "SAFE: A Security Blueprint for Enterprise Networks," a 66-page plan to provide best-practice information to those involved in designing and implementing secure networks. SAFE represents a defense-in-depth approach to network security design, focusing on the expected threats and the best ways to mitigate those threats, rather than a single set of rules to follow. The result is a "layered approach" to security design and implementation, intended to prevent a failure of one security system from compromising the organization network resources.

Since its introduction, the SAFE program has expanded to include many other blueprints, including "SAFE Extending the Security Blueprint to Small, Midsize, and Remote-User Networks," and a growing number of SAFE documents on topics such as wireless, IP telephony, IPSec VPN, Nimda attack mitigation, and Code-Red attack mitigation. These white papers are available free, without a CCO ID from the Cisco web site. To find the current list of documents, either go to following web site or go to http://www.cisco.com and do a search on SAFE: http://www.cisco.com/warp/public/779/largeent/issues/security/safe.html.

TIP For the exams and for your own development in the industry, go to the site and download at least the original "SAFE: A Security Blueprint for Enterprise Networks" and become familiar with it. This provides a detailed overview you can use to supplement all the chapters in this book. At a minimum, review the foundation material in Appendix B.

Cisco Security Wheel

Cisco developed a process they call the Security Posture Assessment (SPA) to describe a company's network security efforts as a living, evolving entity. The SPA is represented graphically in Figure 1-5. As the graphic shows, developing a network security program is an iterative process that must be continually managed to reduce the risk of loss, while efficiently using company resources.

STUDY TIP You should assume this graphic and the process it represents can be part of all four security exams. As you learn about a new technology, make sure you know where that technology fits in the Security Wheel. For example, intrusion detection systems (IDS) would be a part of the monitoring process.

Even if a company had the capital resources and attempted to develop the "perfect" network security solution, it would still be only the beginning of an on-going process. Like a perfect wave for a surfer or a perfect breeze for a sailor, the perfect security system is at best a moment in time, if not an illusion. The factors that led the company to put in the security system have been busily evolving and changing at the same time. The nearly constant changes occurring in technologies used in the network, types and sources of threats, even changes in data flows within the organization continually introduce new

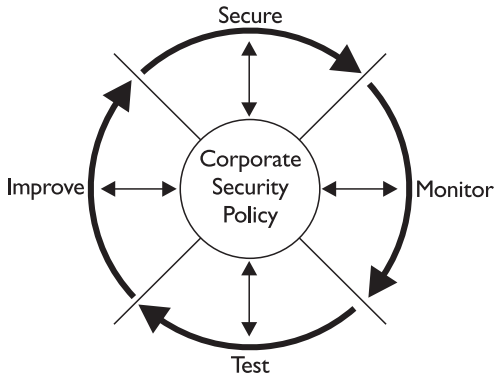

Figure 1-5 Cisco Security Wheel depicting the evolution of a security system

risks that must be anticipated and mitigated. The wheel identifies the four stages of developing a secure system.

- **Secure** After carefully studying the security policy, it's time to secure the network by implementing the processes and technologies required to protect the organization's data and intellectual resources. These could include technologies like VPNs for telecommuters and branch locations, or the addition of firewall devices in the network.
- **Monitor** The security processes and technologies need to be monitored to make sure they provide the security expected. This could involve a variety of activities, ranging from scanning log files to using network management software to detect intrusions, failed attempts, and internal misuse of resources.
- **Test** The test stage can include testing new processes to make sure they meet expectations, testing established processes to see if internal or external changes might have made them less than secure, and periodic audits to see that all processes and policies are being implemented as designed and whether security problems are being dealt with properly.
- **Improve** The improve stage involves developing new plans to adjust the security program for changes in both the internal and the external environment. From implementing "staged" improvements that were built into the original plan to reacting to the latest security threat that could be stalking the network, improvements in both technology and processes are a must.

At the center of the wheel is the network security policy, sometimes referred to as the *corporate* or *enterprise network security policy*. This component, if properly implemented, is the blueprint for the four evolutionary processes of the wheel to follow. The next section looks at the security policy in greater detail.

Network Security Policy

The organization security policy, like the vacation policy or the family leave policy, is an official company document that lays out the expectations of the organization, the processes to be implemented, and the sanctions for those that fail to comply. Without a well-defined and accepted security policy, security becomes an ad hoc process governed by the person in charge at the moment and can (at best) lack any effective usefulness to the organization or (at worse) could lead to significant losses of resources and/or opportunities.

Security policies are covered in detail in the "Site Security Handbook" (RFC 2196). The handbook document defines a security policy as "A formal statement of the rules by which people who are given access to an organization's technology and information assets must abide." It goes on to say that "the main purpose of a security policy is to inform users, staff, and managers of their obligatory requirements for protecting technology and information assets. The policy should specify the mechanisms through which these requirements can be met. Another purpose is to provide a baseline from which to acquire, configure, and audit computer systems and networks for compliance with the policy. Therefore, an attempt to use a set of security tools in the absence of at least an implied security policy is meaningless."

Why Create a Network Security Policy

Any set of policies requires time and effort to produce a useful and effective tool that can be implemented and administered. Some of the benefits and purposes for taking the time to develop a well-structured policy include the following:

- Provides a blueprint for security purchases and implementations
- Defines technologies that can and cannot be added to the network
- Provides the procedures to follow in case of a security incident
- Defines which practices and behaviors are acceptable and which are not
- Defines the sanctions for violations of the policies
- Provides a process and targets to audit existing security
- Defines responsibility throughout all levels of the organization for implementation, monitoring/auditing, sanctioning, funding, and supporting the policies
- Creates a basis for an enforceable legal action
- Provides the baseline for the next step in the evolution of the network security

The Balancing Act

An organization security policy determines how secure or insecure the network and intellectual resources are, what features or functionality is included in the network, and how easy the network resources are to access. A good security policy can only evolve after the organization clearly understands and defines its security goals. Only at that time can intelligent decisions be made about which tools to use, which technologies to allow and/or support, and what restrictions need to be defined and communicated.

Recognizing that the goals of other organizations might not be the same as yours is important. Even competitors within the same industry might have different security needs based on their perceptions, organizational structure, and even whether they are industry leaders or followers. Furthermore, the goals of an organization's vendors—or would-be vendors—might not necessarily be the best to follow. Many network devices include default settings that allow "wide open" operation to maximize throughput and to facilitate adding new devices with little thought for overall network security. For example, notice how many vendors use the top throughput ratings for wireless systems, knowing full well that when security is incorporated, the numbers drop substantially.

A security policy is always the result of compromises and balancing between the following key tradeoffs:

- Security versus ease-of-use
- Security versus services provided
- Security cost versus risk of loss

In the next sections, you will see each of these compromises and the impact each one has on the resulting security.

Security vs. Ease of Use

On the one extreme, a system with no passwords, unlimited access from anywhere, and few restrictions on user behaviors provides the easiest environment for users to work and create within. This also creates an environment where company resources and intellectual property could be easily damaged, lost, or stolen. On the other extreme, frequently changing passwords, restrictive "need to know" access, and draconian penalties for any mistakes can secure the resources at the expense of users being unable to or unwilling to do their jobs to the fullest.

A natural conflict will always exist between the users and the security requirements of an organization. Users often see any restrictions placed on them as interfering with the company's capability to compete and be efficient. Security personnel often see users as security risks, instead of the tools of production that ultimately pay the bills. Figure 1-6 represents the balance between ease of use and network security. The same representation could also be applied to balancing services and cost versus security.

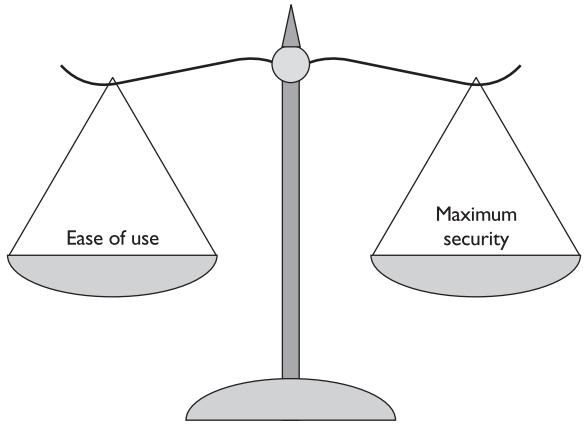

Figure 1-6 Balancing security needs with user ease of use

Security vs. Services Provided

Every service and technology provided to users carries some level of security risks. Some services, such as wireless technologies, involve risks that could be somewhat obvious. Others might involve hidden risks that go unexploited for an extended period of time. In some cases, the decision might be made that the risks outweigh the potential benefits and the service should be eliminated rather than efforts made to secure it.

Security Cost vs. Risk of Loss

This might be the most difficult comparison to make. Implementing a form of security seldom involves just the cost of acquisition, such as the purchase price of a firewall or an authentication server like TACACS+. There can be network performance implications if the new technology increases latency through a device. The new technology or policy could reduce user access or increase user effort in using the technology.

Just as identifying costs associated with implementing some security changes might be difficult, many types and levels of risk can also accrue if the service or policy isn't implemented. The loss could include the following:

- *Loss of company data or intellectual property*, such as the accidental or intentional corruption or deletion of files
- *Loss of service*, such as the loss of a server from a DoS attack, the loss of data storage space because of the replication of nonbusiness files, or the loss of a web site defaced by a hacker
- *Loss of privacy*, such as the copying and/or viewing of company or personnel information by unauthorized users
- *Loss of reputation*, such as the company embarrassment and possible loss of business associated with the disclosure that customer files or credit card information has been accessed by outsiders

NOTE I once had a client who got a call from a police department several states away informing him that his web site was hosting a porn exchange site. The officer provided the exact address and told the client to "take care of it." The officer went on to explain that the call was a courtesy because he was convinced this was an unauthorized use of the server, but that they could just as easily have showed up with warrants and turned the whole place upside down. The site was removed. The company developed a security program, but the client was fearful for a long time that the word might somehow get out and damage his reputation within his client community.

Ultimately, each organization must evaluate its network and intellectual resources to determine the appropriate level of security needed to safeguard those resources, while still enabling the organization to meet its primary mission. This evaluation process often takes time to reach equilibrium and can be altered significantly by events such as a network intrusion or loss of a valuable resource. The lucky groups are those that can learn vicariously from the experiences of others and possibly avoid significant losses of their own.

Network Security vs. Network Operations

Many organizations separate their network security and network operations departments because their missions and pressures are somewhat different. While operations strives to make network resources available in as quick and efficient a method as possible, they might not have the time, resources, or training to analyze the security implications of all situations. Security personnel should have the resources and training to evaluate the organization's needs, without the direct pressures of maintaining a production network. Note, I'm not saying that security personnel can ignore the pressures of running an efficient and reliable production network, only that their main objective is different.

In a perfect world, balance would be struck by having a security staff with previous production experience to better understand their peers in operations, and the operations personnel would have sufficient training and management support for security to perform their jobs effectively, while complying with the security policy.

A Security Policy Is to Be Shared

A *security policy* must be a formal written statement of company policy that has the full support of management and owners. A security policy must be easily disseminated to and scrutinized by users at all levels, operations staff, and managers as a set of security rules that covers all types of information technology, as well as the information stored and manipulated by that technology. To be effective, a security policy must be communicated to users in a clearly understandable form and acknowledged by them through a signed statement, indicating they have read, understand, and will abide by the policy.

Acceptable Use Policy (AUP)

An acceptable use policy (AUP) might logically be included as part of the final security policy. Because of the convergence of technologies, it's common for AUPs to include

telephone, copier, personal digital assistant (PDA), and pager, as well as fax activities. The AUP should spell out specifically what users can and cannot do on the various components that make up the network, including the type of activities and traffic allowed on the networks. The AUP should be as explicit as possible to avoid ambiguity or misunderstanding, particularly if sanctions are imposed for failure to comply. For example, an AUP might list prohibited activities like "browsing and engaging in transactions on web auction sites." The AUP could be explained at all employee orientation sessions and signed by each user. This agreement and training should be updated periodically as a refresher and definitely any, time a significant change is made.

Unfortunately, it's not uncommon for new employees to learn about new limitations on access to resources, such as Internet access for personal use only, to find the same information hasn't been distributed to the existing employees. It's particularly dysfunctional when the sanctions involved with a policy are implemented against employees who had no way of knowing they were violating a policy. Handled poorly, this can lead to mistrust, lack of support, and even refusal to use certain resources.

Network Operations and Network Security Training

Networking employees at all levels will require additional training and sign-offs on those portions of the security policy that impact their jobs, but that aren't covered in the AUP. Again, this should be handled proactively with enough detail that each employee understands their responsibilities and the limits of their authority. No one wants to learn that the configuration file they sent to Cisco TAC for help with a problem might cost them their job as a breach of the company security policy.

Who Should Help Create the Security Policy?

For a security policy to be effective, it must have the acceptance and support of all levels of users within the organization. Especially important is that corporate management and ownership (board of directors) fully support the security policy process; otherwise, little chance exists that it will be successful. Also critical is that the resulting policy will eventually fit within the organization and its culture. In particular, a first security policy or a radical change in policy might require some transition time for people to learn and assimilate the new rules. The following people are representative of those who should typically be involved in the creation and review of security policy for a larger organization:

- Company security administrator
- Security incident response team representatives
- IT technical staff representatives (network operations)
- Administrators of organization business units
- Representatives of the user groups
- Responsible upper management
- Corporate legal counsel (in some countries)

The wide variety and sizes of businesses make it impossible to define a single list. The nature of the business and the level of and types of employee contracts and bargaining units might dictate some other attendees. Just because a security policy is necessary and reasonable doesn't set aside a company's requirements to negotiate changes in work rules. More than one organization has been required to rehire with back pay an employee terminated under a security policy rule because it conflicted with a bargaining agreement.

Another group that should be represented is any internal auditors required by industry standards or governmental regulations. Because some policies dictate production of logs, backups, and documentation, it's critical that those policies comply with any relevant laws, regulations, industry standards, or court orders.

If the resulting policy statements are to reach the broadest possible acceptance, the group must be an appropriate mix of involved representatives (stakeholders) that can formulate a set of rules that balance the security requirements with the technical expertise available or obtainable. These policies must have an acceptable impact on the company business model, particularly in any areas perceived to create a competitive advantage. Finally, the budget and policy authority must be present to make sure these policies are supported throughout the organization and funded adequately during both good times and bad. If done properly, the policy should yield the highest level of appropriate security in the most cost-effective manner.

Assets and Threats

Developing a security policy, as in any risk analysis, involves determining what needs to be protected, what it needs to be protected from, and how best to protect it. So the first things to do in the process are

- Identify the assets
- Identify the threats

The process then involves examining all possible risks, ranking those risks by level of severity, and, finally, making cost-effective decisions on what you want to protect and to what extent. It makes no sense to spend more to protect something than its actual worth.

Identifying the Assets

When identifying the assets that need to be protected, some might be obvious, like valuable proprietary information such as product blueprints or designs, intellectual property, and the many hardware components that make up the network. Others might not be so obvious, though, and are often overlooked, such as the people using the systems. While the company doesn't own the people, it could have invested in their skills and development over the years. Similarly, the company might rely heavily on those skills to meet its business objectives. Some users might have no readily identifiable replacements within the current workforce.

The point is to list everything that could be impacted in any way by a security problem:

- **Hardware** Servers, workstations, laptops, printers, scanners, FAX units, routers, switches, firewalls, intrusion detection devices, wireless access points, IP telephones, palm-sized devices, pagers, projection systems, electronic white boards, and communication lines. Don't forget devices that might be at telecommuters' homes, such as DSL routers, printers, and so forth. The move to combine resources like printers and copiers should be acknowledged, even if not yet implemented.
- **Software** User software licenses, custom and off-the-shelf enterprise applications, virus protection software, network and workstation OSs, network device OS, network management applications, utilities of all types, diagnostic programs, and communication/FAX programs.
- **Data** Financial records, business plans and strategies, customer and employee information, sales records (including credit card information), product designs and parts lists, inventories, production schedules, and customer and vendor contracts. Many of these could be parts of one or more databases, while others might be many individual documents in the system. Each type must be identified by its location during execution, where they're stored online, where they're archived offline, any backups, audit logs, and whether they're ever transmitted over communication links. It isn't uncommon to discover entire classes of strategic documents stored only on local hard drives.
- **People** Users, administrators of all types, help desk people, and hardware maintenance.
- **Documentation and licenses** For OSs, applications, hardware, systems, and administrative procedures. Don't forget service agreements and warranties.
- **Supplies** Paper, toner and ink cartridges, and batteries.
- **WAN and Internet services** Contracts and service agreements for communications links, web hosting services, and related contracted services of any kind. Because these services could be in negotiation for some time, be sure to include any works in progress.

While not technically a network component and not appropriate for all companies, as previously mentioned, any company doing business over the Internet ought to consider its reputation and the trust relationships it's developed as an asset. Any attack that damages this reputation could have serious implications for the future well being of that company and its stakeholders.

Identifying the Threats

Once the assets to be protected are identified, it's necessary to identify and assess the threats to those assets. As you saw earlier in this chapter, different names and levels of

detail can be applied to these security threats, but, generally, they break down to the following:

- Reconnaissance
- Unauthorized access
- Data manipulation
- Denial of service

The basic goal of security for each asset is to ensure availability, confidentiality, and integrity. Each threat should be examined to assess how it could impact the company assets. Each company can have different perceptions and assessments of the threats, which should be identified and addressed.

STUDY TIP Data manipulation is often included in unauthorized access threats, so keep that in mind if you are asked to identify only three threats.

Evaluating a Network Security Policy

Just as the preceding section showed that different organizations can have different perceptions of threats and place different values on assets, the security policy reflects those differences. Most security studies agree that some of the key characteristics of a security policy should include the following:

- It must be implementable through network administration technologies, by publishing rules and acceptable use policies, or other appropriate methods.
- It must be enforceable with security tools, where appropriate, and with sanctions, where actual prevention isn't technically or financially feasible.
- It must clearly define the areas of responsibility for the users, administrators, and management. Maybe as important, it should clearly identify the limits of authority for each group under predictable circumstances.

Any policy that has serious deficiencies in any of these characteristics stands a better-than-good chance of failing to meet the company objectives.

What Belongs in a Network Security Policy

Each organization will develop a policy based on a variety of factors. Even after exhaustive study and development, the policies need to be updated as new technologies develop or become financially feasible. Some common components of a good security policy include the following:

- **Statement of authority and scope** Who sponsors and authorizes the policy, as well as who it impacts.
- **Access policy** Defines all access privileges and responsibilities to network assets by specifying acceptable use rules for users, network operations staff, and management. These policies should cover adding devices to a network, adding software to systems, modifying software and OS settings, and how and by whom external connections to the network are to be made. The policies should specify any required connect warning messages about "authorized usage and monitoring." Parts of the access policy might become part of acceptable use policy (next).
- **Acceptable use policy** Which user practices and behaviors are acceptable and which are not. This policy often includes technologies, including telephones, cell phones, pagers, copiers, fax machines, computers, access to the Internet, and so forth.
- **Privacy policy** Defines what are reasonable privacy expectations regarding monitoring of e-mail and access to users' files.
- **Remote access policy** Defines how remote users and telecommuters will access the organization networks. This policy might be further broken down to cover specific technologies, such as ISDN access policy, DSL access policy, and cable modem access policy.
- **Wireless access policy** Defines if and under what circumstances a wireless devices or devices can be used with the company network.
- **Antivirus policy** Defines which tools will be used and how they'll be implemented.
- **Password policy** Defines what passwords will look like and how often they must be changed, and it should authorize audits of password files to ensure compliance.
- **Authentication policy** A more comprehensive form of the password policy that defines a local access password policy and establishes guidelines for remote authentication processes, which might include OTPs and the devices that generate them.
- **Router and switch security policy** Defines minimal security configuration for all routers and switches connecting to a production network.
- **Availability statement** Defines what users can expect for resource availability. Known risks, redundancy, and recovery issues should be stated. Hours for routine maintenance downtime should be specified, as well as any notification process used before the system is taken down. Contact information for reporting system and network failures should be included.

- **Accountability policy** Defines the responsibilities of users, network operations staff, and management. This should cover guidelines for routine monitoring, scheduled and unscheduled audits, and guidelines for incident handling (what to do, what not to do, and who to contact in case an intrusion is suspected).
- **IT system and network maintenance policy** Defines how internal and external maintenance people are allowed to handle and access company technology. This should address if, and under what conditions, remote maintenance is allowed and how such access is controlled. If outsourcing is allowed, how it should be managed and the processes that need to be followed should be defined.
- **Violations reporting policy** Defines the types of violations that must be reported and to whom the reports are made. Remember, a low-key atmosphere and even anonymous reporting can result in a greater likelihood that a suspected violation will be reported.The policy should include specific contact information for users, staff, and management for each type of policy violation. Guidelines should include how to handle outside queries about a security incident, detailing who should be responding, and defining the procedures depending on where the contact is from. There may be a different policy or contact person when working with an interested third party, law enforcement, or the media. Defines the procedures if the security incident involves information that might be considered confidential or proprietary. If appropriate, any cross-references to security procedures company policies and/or applicable laws and regulations.

Here's an example. Recently in our area, a policeman was ordered reinstated with about a year's full pay and benefits, plus overtime that would have been earned. While the court agreed that viewing pornography on a department laptop in a public place was a violation of the AUP in place, it disagreed with the penalty for a first violation, determining that termination was too severe. The $86,000 is probably less damaging than the harm done to the department's reputation and credibility.

Make Time for Training and Signing Off

Once the security policy is established, reviewed as needed, and approved by the highest levels in the company, it should be clearly communicated to all users, network staff, and management. Remember, this training of the employees is the only opportunity to express the importance of the effort, the seriousness of the company commitment, and the need for their active participation. Each person should be able to retain for future reference any appropriate sections, including at least the AUP.

Having all personnel sign a statement that indicates they've read, understood, and agreed to abide by the policy is common practice and logical. Note, this signing is of questionable value in protecting the company resources if the policies aren't explained or treated with respect by management. The last security policy I signed was a modified distribution list attached to a stack of papers passed around a conference table. I remember

thinking that because virtually nobody read the attached document, it probably wasn't going to modify much existing behavior.

Keep It Flexible

For a security policy to be viable over the long term, it must be as flexible as possible, allowing for growth and changes within the architecture. As much as possible, the policy should be independent from specific hardware and software because these items can be moved or removed on short notice.

The processes for updating the policies should be defined in detail to include the steps to do, the people to include, and who must sign-off on the resulting changes.

Finally, the policies must be reviewed on a regular basis to make sure they're successfully meeting the company security needs. The policies should also be reviewed if any new major technology is added to the network to make sure the technology is covered from day one.

Example of a Network Security Policy

A security policy can take many forms and styles, but one that's easy to get started with is a series of templates produced by the SANS Institute called "The SANS Security Policy Project," available from their web site at http://www.sans.org/newlook/resources/policies/policies.htm.

The System Administration, Networking, and Security (SANS) Institute was established in 1989 as a cooperative research and education organization to provide a forum for security professionals, auditors, system administrators, and network administrators to share the lessons they have learned and find solutions to the challenges they face. Many SANS resources, such as news digests, research summaries, security alerts, and white papers, are free. For more information, their web site is http://www.sans.org.

One of many useful sites SANS hosts is the "SANS/FBI Top Twenty List" (http://www.sans.org/top20.htm), which summarizes the "Twenty Most Critical Internet Security Vulnerabilities." This web site also includes step-by-step instructions and pointers to additional information useful for correcting the flaws.

The writers at SANS have produced about two dozen templates for each of the major policies to be included in a security policy. Each policy is a separate document that allows for easy addition and editing, and each follows the same style shown in the example in the upcoming sidebar.

Securing the Network

Once the security policy is approved, it's time to make sure that any security measures specified or appropriate to the policy are properly installed and configured. The rest of this book focuses on many of those technologies. In the next chapter, beginning in the section on "Improving Network Security," you start looking at those basic configurations that should be a part of every device configuration.

Wireless Communication Policy

1.0 Purpose

This policy (provided here with permission from SANS) prohibits access to <Company Name> networks via unsecured wireless communication mechanisms. Only wireless systems that meet the criteria of this policy or have been granted an exclusive waiver by InfoSec are approved for connectivity to <Company Name>'s networks.

2.0 Scope

This policy covers all wireless data communication devices (that is, personal computers, cellular phones, PDAs, and so forth) connected to any of <Company Name>'s internal networks. This includes any form of wireless communication device capable of transmitting packet data. Wireless devices and/or networks without any connectivity to <Company Name>'s networks don't fall under the purview of this policy.

3.0 Policy

To comply with this policy, wireless implementations must do the following: maintain point-to-point hardware encryption of at least 56 bits; maintain a hardware address that can be registered and tracked, that is, a MAC address; and support strong user authentication that checks against an external database, such as TACACS+, RADIUS, or something similar.

EXCEPTION A limited-duration waiver to this policy for Aironet products has been approved, if specific implementation instructions are followed for corporate and home installations.

4.0 Enforcement

Any employee found to have violated this policy could be subject to disciplinary action, up to and including, termination of employment.

5.0 Definitions

- **Terms** (add terms here)
- **Definitions** (add definitions here)
- **User Authentication** A method by which the user of a wireless system can be verified as a legitimate user independent of the computer or OS being used.

6.0 Revision History

The references to <Company Name> would be replaced by the company or client name and references to **InfoSec** would be replaced by the name of the company department responsible for network security. The reference to Aironet refers to a pioneer wireless manufacturer, since acquired by Cisco. The design with standardized sections and titles makes it easy to add a policy unique to the organization that would fit right in.

Monitoring Network Security

The security policy should specify the methods to be implemented in the routine monitoring of the network. The purpose of security monitoring is not only to observe a network attack, but also to point out potential weaknesses that could be exploited. The one thing monitoring should verify is whether the security policy is being followed.

Monitoring could be as simple as an orderly collection and review of the various log files that network devices generate as a normal part of operation. Simply viewing failed login authentications for a server can indicate attempts to break into the system or maybe just some individuals that need additional training. At the other end of the spectrum are sophisticated devices like IDS that can monitor traffic looking for patterns or signatures that would indicate something is amiss. If a potential problem is discovered, the IDS sensor can notify the IDS director management console, which can then start a process to block (shun) the attack. It could involve creating an access control list in a router or firewall specifically to block further contact from that source. IDS technology is covered in Chapter 7, and then again in Chapters 23 to 26.

Auditing Network Security

Whereas monitoring is a routine general process of watching the network for potential problems with the network security, auditing is a more specific test or series of tests on the network to determine the effectiveness of the security measures, as well as compliance with the policy and any relative regulations. Auditing, like monitoring, should be specifically defined, and allowed within the formal security policy.

The random nature of audits as compared to routine monitoring makes them useful in detecting unauthorized internal activity within the system. For this reason, audits shouldn't become routine or scheduled. Much of their usefulness is diminished if your troublemakers can predict when an audit will occur.

Shareware tools are available on the Internet and even many operator-triggered system tests can be automated by the use of scripts. Checklists of features to audit for many technologies are available on the various security web sites. Some auditing is low tech. It can be as simple as periodically reviewing the tape backups to confirm the backup was

run when scheduled, that it covered everything specified, and that the tape has a usable file. Other reasons for audits might include the following:

- Any addition to or change to systems in the network. Has the new system or the change (upgrade) possibly brought with it some unknown problem? This shouldn't be an optional audit.
- Spot checks of policy compliance.
- Regularly checking password files for compliance with the password policy.
- Router and switch configuration compliance to make sure all appropriate security requirements are included.

Commercial auditing tools or even network intrusion tools like Cisco Secure Scanner can perform ongoing auditing, including looking for potential security holes. Some of these tools come from a variety of vendors, but might only run on certain management stations. This is one area where UNIX seems to have a clear edge.

Improving Network Security

Just as the image of the security wheel implies, network security is a constantly evolving and growing process. This process is driven by the changing and growing nature of the business on one side leading to more and more resources and possibly links to more outside sources. Pressing from another side is the bad guys outside of your network who are constantly gathering better tools, often the same ones you'll be learning about. They also don't have a security policy preventing them from trying the latest and greatest hack posted on the Internet. They could also have "cracked" copies of licensed tools and software the company can't afford. If potential attacks from two growing fronts weren't bad enough, internal users are becoming savvier about the workings of the network. Economic turmoil can often bring out a side of people that even they might not have known existed under other circumstances.

So what is the company to do? Start by recognizing that what you have today in the form of network security becomes the baseline from which the future is built. There's no going back to last year, so today becomes the beginning of time. Planning and development must always look at the next level of safety the network security can be moved to. Meanwhile, caution and good practices would suggest the following:

- Monitor the security alerts from all network device vendors and install the recommended patches and upgrades.
- Stay current on the latest threats, vulnerabilities, and tools by monitoring security web sites and newsgroups, such as www.sans.org, www.cert.org and www.cisecurity.org.
- Implement and follow the existing network security policies, including incident investigation and reporting. Lax implementation and enforcement leads to potential vulnerabilities and can undermine commitment to security.

- Update the security policy on a regularly scheduled basis, plus any time a new technology is added to the network or an existing technology is removed.
- Ongoing security training and awareness should be a priority at all levels within the company.
- Encourage a sense of trust and friendliness to encourage employees to "ask first" when in doubt and to encourage reporting of potential security incidents.

Chapter Review

Technologies like the Internet, wireless connectivity, instant messaging, and others have introduced new levels of concern for network security. In addition to providing additional access points to the network, the Internet is also a reference resource for hackers and security experts. Unfortunately, no rules or limitations exist on the information that can be posted, so it's easy for harmful information to be widely distributed.

The secure network design example is to lay a basic design foundation for discussions throughout the book. Basically, three types of networks are separated by a firewall device, which can be a router running firewall features, a server running firewall features, or a dedicated firewall device. The three types of networks are the following:

- **Inside** Those trusted hosts and networks that make up the area to be protected.
- **Outside** Those hosts and networks that pose a threat to the inside network. In many cases, the outside includes a perimeter router, the ISP, the Internet, and any networks attached to it.
- **DMZ** A network of shared servers, or bastion hosts, that provide resources to both the outside network and the inside network users.

The firewall configuration allows relatively free access from the outside to the DMZ; access from the DMZ to the inside is limited to sessions originating inside; and access from the outside to the inside is generally blocked, except in rare well-protected exceptions.

The three most common weaknesses or vulnerabilities that cause network security problems include technology weaknesses, such as IP or OS holes; policy weaknesses, such as missing or weak security policies; and configuration weaknesses, such as insecure default settings.

The four primary threat designations are unstructured, structured, inside, and outside. These designations can be used in combinations like outside—unstructured to better define the type of attack.

The four most common attack types include reconnaissance attacks, access attacks, denial of service attacks, and data manipulation attacks. Many references, including the PIX Firewall exam, consider data manipulation attacks to be variations of access attacks.

Cisco Architecture for Voice, Video, and Integrated Data (AVVID) and SAFE are Cisco-comprehensive strategies that help organizations to successfully and securely develop and implement end-to-end network designs.

Questions

1 Which of the following is *not* a common cause of network security problems?

A. Technology weakness

B. Configuration weakness

C. Planning weakness

D. Policy weakness

2. Which four of the following are primary types of network threats?

A. Planned threats

B. Unstructured threats

C. External threats

D. Structured threats

E. Internal threats

3. A virus is an example of which of the types of network threats?

A. Planned threats

B. Unstructured threats

C. External threats

D. Structured threats

E. Internal threats

4. Which of the following is *not* one of the four primary types of network attack?

A. Access attacks

B. Data manipulation attacks

C. Reconnaissance attacks

D. Programmed attacks

E. Denial of service attacks

5. In a reconnaissance attack, which one of the following is *not* a part of target discovery?

A. Ping sweeps

B. Ping of Death

C. Port scans

D. DNS queries

6. Which one of the following is an example of social engineering relative to security?

A. Guessing a password based on a person's vanity license plate

B. All users in a department using the same user name and password

C. Giving your password to the home office technician, so they can test your account

D. Using a password of eight characters, uppercase and lowercase, plus numbers

7. A standard DoS attack typically includes which one of the following?

A. An attacker, a series of Zombies, and the target

B. An attacker, a war dialer, and the target

C. An attacker, a series of Zombies, and the target

D. An attacker and the target

8. Which of the following is *not* a well-known DoS attack?

A. Ping of Death

B. Tribe Flood Network

C. Trinoo

D. TCP SYN Flood

E. Script kiddy

9. IP spoofing is an example of which of the following?

A. DoS attack

B. Reconnaissance attacks

C. Data manipulation attacks

D. Brute force attack

10. RFC 2827 filtering refers to which of the following?

A. Filtering private IP addresses

B. Filtering based on trusted external addresses

C. Filtering based on source addresses that belong on a network segment

D. Filtering based on destination addresses that belong on a network segment

11. Which of the following involves denying responsibility for a transaction?

A. Session replay

B. Rerouting

C. Repudiation

D. Session hijacking

12. Which of the following provides a standards-based network architecture and comprehensive set of best practices that allow businesses to develop business and technology strategies, which scale to meet the changing demands of e-business?

A. SAFE

B. DrDos

C. AVVID

D. Digital signatures

13. RFC 1918 filtering refers to which of the following?

A. Filtering private IP addresses

B. Filtering based on trusted external addresses

C. Filtering based on source addresses that belong on a network segment

D. Filtering based on destination addresses that belong on a network segment

14. Which of the following best describes a good security plan for an organization?

A. Set it and forget it

B. Ad hoc

C. Evolutionary

D. Rigid

15. Which one of the following is *not* a part of the Cisco Security Wheel?

A. Monitor

B. Administer

C. Test

D. Secure

16. What is at the center of the Cisco Security Wheel?

A. Strong network support staff

B. Improvement

C. Network security policy

D. Flexibility

17. According to the "Site Security Handbook" (RFC 2196) a security policy is

A. An informal set of suggestions by which people who are given access to an organization's technology and information assets should abide.

B. A formal statement of the rules by which people who are given access to an organization's technology and information assets must abide.

C. A formal statement of the rules by which network administrators can control access to an organization's technology and information assets.

D. RFC 2196 doesn't deal with this issue.

18. Which of the following is *not* a reason to create a network security policy?

A. It provides a blueprint for security purchases and implementations.

B. It defines technologies that can and cannot be added to the network.

C. It provides the procedures to follow in case of a security incident.

D. It frees up network security personnel to work on other projects.

E. It provides a process and targets to audit existing security.

19. Developing a security policy is often described as a balancing act. Which three of the following are common compromises that need to be made?

A. Security cost versus risk of loss

B. Security cost versus local or federal regulations

C. Security versus ease of use

D. Security versus services provided

E. Local or federal regulations versus ease of use

20. When evaluating risk associated with a security breach, which three of the following might result from someone accessing the companies sales order entry system?

A. Loss of reputation

B. Loss of company data or intellectual property

C. Loss of temper

D. Loss of service

E. Loss of privacy

21. Which of the following might be included in an acceptable use policy (AUP)?

A. Acceptable and unacceptable Internet activities

B. Penalties or sanctions for violating the policy

C. Acceptable and unacceptable e-mail activities

D. Acceptable and unacceptable telephone use

E. Acceptable and unacceptable copier or fax use

F. All of the above

22. Which of the following is likely to result in a security policy that will meet the needs of an organization?

A. The network security staff develops the policy.

B. The network security staff and network operations staff develop the policy.

C. Representatives of all major groups that use and manage the networks, plus representatives of management develop the policy.

D. The network security staff and company management develop the policy.

23. In developing a security policy, as in any risk analysis, the first two things you must identify are the

A. Opportunities

B. Assets

C. Hardships

D. Threats

24. According to the RFC, which of the following is *not* a characteristic required of a good security policy?

A. It must be implementable through network administration technologies, by publishing rules and acceptable use policies, or other appropriate methods.

B. It must clearly define the areas of responsibility for the users, administrators, and management. Maybe as important, it should clearly identify the limits of authority for each group under predictable circumstances.

C. It must be state of the art, implementing all the latest, most advanced technologies and procedures to protect the company resources.

D. It must be enforceable with security tools, where appropriate, and with sanctions, where actual prevention isn't technically or financially feasible.

25. When a group of sales reps plug a wireless access point into a network jack so they can use wireless NICs on their laptop, they

A. Greatly increase their mobility and freedom

B. Increase the number of usable connections to the network

C. Probably violate the wireless access policy of the security policy

D. Add a hub that will reduce the bandwidth to each user

26. Which of the following statements best describes monitoring and auditing?

A. They are two terms for the same process.

B. Auditing is routine scheduled reviewing of security, while monitoring is random and unpredictable reviewing of security.

C. Monitoring catches bad guys outside the network, while auditing catches bad guys inside the network.

D. Monitoring is routine scheduled reviewing of security, while auditing is a random and unpredictable reviewing of security.

Answers

1. C. Planning weakness
2. B. Unstructured threats; C. External threats; D. Structured threats; and E. Internal threats
3. B. Unstructured threats
4. D. Programmed attacks
5. B. Ping of Death

6. C. Giving your password to the home office technician so they can test your account. There is never a reason to give someone your password, or for a legitimate tech to ask for it.
7. D. An attacker and the target
8. E. Script kiddy
9. C. Data manipulation attacks
10. C. Filtering based on source addresses that belong on a network segment
11. C. Repudiation
12. C. AVVID
13. A. Filtering private IP addresses
14. C. Evolutionary
15. B. Administer
16. C. Network security policy
17. B. A formal statement of the rules by which people who are given access to an organization's technology and information assets must abide
18. D. It frees up network security personnel to work on other projects.
19. A. Security cost versus risk of loss; C. Security versus ease of use; and D. Security versus services provided
20. A. Loss of reputation B. Loss of company data or intellectual property; and E. Loss of privacy; (While loss of temper might occur, this isn't one of the risk criteria.)
21. F. All of the above
22. C. Representatives of all major groups that use and manage the networks, plus representatives of management develop the policy.
23. B. Assets and D. Threats
24. C. It must be state of the art implementing all of the latest and most advanced technologies and procedures to protect the company resources. (While a good idea, this might be far from cost-effective for many organizations.)
25. C. Probably violate the wireless access policy of the security policy (While all the statements are true, within the context of this course, this is *the* important issue.)
26. D. Monitoring is routine scheduled reviewing of security, while auditing is a random and unpredictable reviewing of security.

CHAPTER 2

Securing the Network

In this chapter, you will learn how to:

- Secure network design example
- Improve network security
- Secure network devices
- Use access control lists (ACLs) to secure the network

As you saw in the preceding chapter, the network has many threats. Developing a comprehensive security program to combat those threats requires planning in the initial stages, and a lot of tweaking and revising as time goes on. Just as the organization attempts to move the network to a more secure state, the threats are also evolving, changing and often growing stronger. Network security will never be a "set it and forget it" world.

This book looks at many technologies and methods to secure access to and operation of the computer networks. While many variations exist, some of the most important ones include the following concepts.

Physical security	Once stated, this seems so obvious, but it's often overlooked. Network devices must be physically secured from unauthorized access or even theft. Password recovery techniques make it quite easy for anyone to access and reconfigure a device if they can have physical access to it.
Vulnerability patching	Network devices from workstations to routers have or develop vulnerabilities that can usually be mitigated by applying software patches, performing upgrades, and disabling any unnecessary services.
Encryption	If making the data path absolutely secure isn't possible, then encrypt the data. Encryption, such as IPSec, means anyone capturing the data will find useless gibberish.
Firewalls	Firewalls filter traffic based on predefined permit-and-deny rules. Ideally, a firewall devotes 100 percent of its resources to protecting the network.

Intrusion detection	Intrusion detection systems (IDS) detect certain patterns of data that match known "signatures" of improper activities. The IDS system can then notify network management or even implement measures to block the activity.
Authorization systems	Secure authorization systems, such as one-time passwords (OTP), limit the usefulness of a password that's been compromised or captured through sniffing activity.

In this chapter, you learn some other simple techniques to improve network security. Some of these techniques should be familiar from other certifications, but each one provides a small piece of the strategy necessary to secure the network.

Secure Network Design Example

To lay a foundation for discussion of secure networks, this section looks at some basic terms and concepts used throughout the book. In security terms, you have three types of networks to consider: inside, outside, and an optional network called the demilitarized zone (DMZ). A *firewall* is the device that separates or joins these areas. The firewall can be a router running a firewall feature set or a specialty server, or it can be a specialty device such as the Cisco PIX that does nothing but provide firewall services. Figure 2-1 shows a simplified view of the three areas and the firewall.

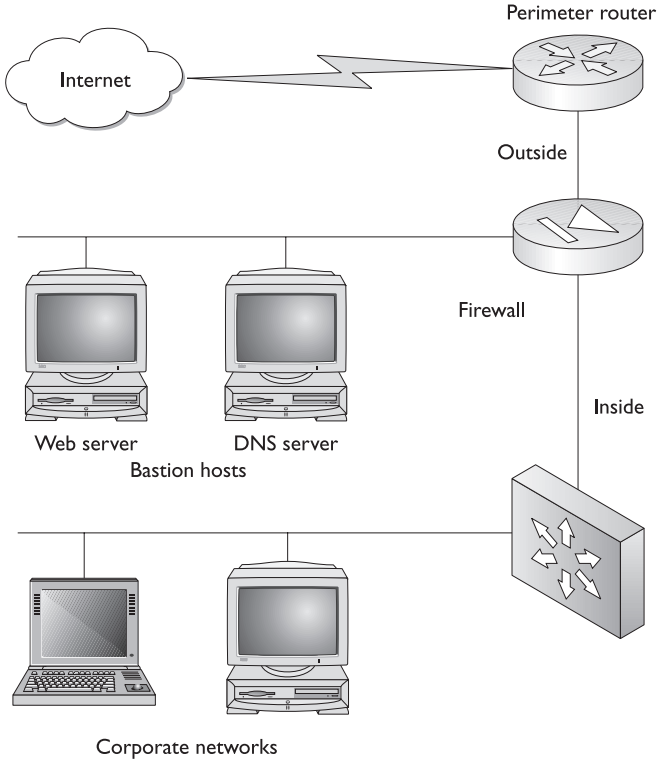

Figure 2-1 A firewall separating the three security areas

The typical firewall device has three or more LAN interfaces: one each for the inside and outside networks, and one for each DMZ network. Some early firewalls and those used in small implementations like branch locations or telecommuter residences might only have two interfaces for separating the inside network from the outside world. Today the LAN interfaces are typically Fast Ethernet or Gigabit Ethernet, but there's no reason they couldn't be Ethernet, Token Ring, or Fiber Distributed Data Interface (FDDI).

Inside Network

Inside networks are also referred to as the *internal* or *private networks*. The inside area is made up of the network(s) of the organization, including all workstations and servers not shared with the outside world. These devices are considered *trusted* and in need of protection from the outside world. The inside area is typically under one administrative authority and operates under a common security plan.

A firewall is normally used to separate the inside network from the outside world, but a firewall can also be used to separate internal departments if additional security is required. For example, a school might choose to place a firewall between the student network and the faculty network.

Outside Network

Outside networks are also referred to as the *external* or *public networks*. The outside area, or *untrusted area,* is considered to be all devices and networks beyond the direct control of the organization's administration and security policies. The outside area would typically include the perimeter router, the ISP, the Internet, and all networks attached to it.

While not necessarily the greatest actual threat to many networks, the international scope of the Internet means an organization can face threats from anywhere in the world for reasons that could seem ludicrous at home.

Demilitarized Zone (DMZ)

The demilitarized zone (DMZ) is made up of one or more isolated LAN networks that contain shared server resources, such as web, DNS, and e-mail servers. These servers are available to the outside world. These shared servers are often called *bastion hosts, bastion servers,* or even *sacrificial hosts*. Bastion hosts must be secured and receive highest priority security maintenance because of their vulnerability to the outside world and increased likelihood of attacks. A bastion server typically runs only those specific services being shared, and all other services will be stopped or turned off.

The firewall must be configured to allow quite loose, but regulated, access to the DMZ from the outside network while at the same time protecting the inside network. Inside network users need access to the server resources in the DMZ and are typically allowed limited access, possibly restricting access to only those sessions originating within the inside network.

Generally, the firewall will be configured to prevent access from the outside to the inside, possibly limiting access to those sessions originating from the inside network. Other, unsolicited, access from the outside would be blocked in most cases. One

common exception might be the e-mail server(s) if it resides in the inside network instead of the DMZ. Securing this type of connection is covered in the firewall chapters.

Securing Network Devices

The rest of the chapters in this book cover technologies and processes to secure the corporate network. Some are pretty simple, while others are more elaborate. The remainder of this chapter covers some basic techniques that should be a part of any basic router or switch configuration, regardless of its role or placement within the network. Many of these might be familiar to you from your earlier certifications. While this list of tasks is by no means exhaustive, it does represent the types of details that must be covered to make sure our networks are secure.

Physically Secure the Devices

All network devices need to be in a physically secure environment, whether in a locked data closet, a locked cabinet, or both. In most cases, if routers and switches can be physically accessed, they can be compromised. With Cisco devices, the only things required to compromise the system are a console cable kit and a terminal system, which today can include many palm-sized devices. Password recovery techniques are well known and easy to implement.

While password recovery lets a person with less than CCNA skills take control of the device, an equal fear should be the person with a screwdriver who decides to take the device(s) and worry about accessing them later. Cisco devices are typically hot items on web auction sites, often with several thousand listings on any given day. While corporate data centers typically have secure facilities, many small businesses and small branch offices might rely on Telco closets for router placement. Before agreeing to place devices in a Telco closet, consider that, in many cases, every building tenant has direct or indirect access to that closet.

Other reasons for centralizing network devices into a single room include facilitating environmental features like climate control (heating and cooling), stable power with Uninterruptible Power Supply (UPS) backups, secure access including locks and protection from over the wall or under the floor access, and possibly increased human presence to provide a deterrent. If the data room is busy with many people having access, it might make good sense to put key devices into locked cabinets.

Securing Administrative Access

One of the fundamental requirements in protecting the network is to secure the administrative access to any network device. With Cisco devices, administrative access to the device could allow someone to reconfigure features or even possibly use that device to launch attempts on other devices. Some of the basic techniques for securing administrative access would include the following:

- Setting User mode passwords
- Setting Privilege mode passwords
- Encrypting passwords in the configuration files
- Setting an MOTD banner to advise about the security restrictions
- Setting access privilege levels
- Restricting Telnet access to the device
- Restricting web browser access to the device
- Restricting SNMP access to the device

The first five techniques are covered in this chapter, while the last three are covered in Chapter 3, which deals with access control lists.

User Mode Passwords

This section reviews and expands on the techniques for assigning passwords to the three potential access points to the *User* mode, the entry level into a Cisco device. In Chapter 3, you learn how to use authentication servers, such as TACACS+ and RADIUS, with AAA authentication services for securing access to Cisco devices.

The User level on a Cisco router often has three potential access points. They include the following:

- **Console (con) port** Access for the console cable. Figure 2-2 shows a typical console port on a router.
- **Auxiliary (AUX) port** A console-like access that can be attached to an external modem for a dial-up connection.
- **Virtual terminal (vty) ports** The access points for Telnet sessions.

The default configuration for each of these interfaces, shown in the following code listing, doesn't include a password. Since the release of version 12.0 of the IOS, the virtual terminals and AUX ports require that a password is set. If none is set, the user will be rejected with the message "password required, but none set." The console port doesn't

Figure 2-2 Console port on an 800 model telecommuter router

have this requirement, so it's a good idea always to set a password to prevent anyone with a laptop and a console cable from accessing the device.

```
interface FastEthernet0/0
 ip address 192.168.0.1 255.255.255.0
!
line con 0          <-Console connection
 login
line aux 0          <-AUX connection
 login
line vty 0 4        <-Virtual terminal connections
 login
end
```

The basic password configuration for each is the same. The password is defined with the **password** command and the **login** command. Passwords can be 1 to 25 character, and can include uppercase and lowercase letters, as well as numbers, to comply with complex password requirements in the password policy. The result might look like the following listing:

```
!
line con 0
 password cisco1
 login
line aux 0
 password cisco2
 login
line vty 0 4
 password cisco3
 login
end
```

The passwords used should comply with the password policy portion of the network security policy. You could use the same password for all three, but this isn't a secure solution. Someone attempting to access the device through one of these three methods will be prompted only for a password, at which time they need to supply the appropriate one.

User Name/Password with Login Local You can require both a user name and a password, as well as have the opportunity to create different combinations for different users. The first step is to develop a local database of acceptable user name and password combinations in the Global Configuration mode. Like all passwords, these are case- ensitive, can include text and numerals, and should comply with the password policy. The user names aren't case sensitive. Two examples might include the following:

```
Rtr1(config)#username remote password acC3ss
Rtr1(config)#username scott password woLfe7
```

To finish the configuration, change the login command to **login local** for the interface(s) that you want to use this feature. In the following example, only the virtual terminal lines are being changed.

```
username remote password access
username scott password wolfe
```

```
!
line con 0
 password cisco1
 login
line aux 0
 password cisco2
 login
line vty 0 4
 login local
end
```

After making the changes, the next Telnet session login might look like the following:

```
User Access Verification

Username: remote              <-Used the remote / acC3ss combination
Password:
Rtr1>exit                     <-Successful attempt

User Access Verification

Username: scott               <-Used scott / WOLFE7 combination - note case
Password:
% Login invalid               <-Wrong case on the password)

Username: ScOtT               <-Used ScOtT / woLfe7 combination - note case)
Password:
Rtr1>                         <-Used the correct case on the password)
```

The important things to remember are that the user name isn't case sensitive, while the password is. Furthermore, if more than one entry is in the local database, then any valid combination is acceptable.

Privilege Mode Passwords

Access security for the Privilege mode involves being prompted for a password only if an **enable password** or **enable secret** password has been previously defined in Global Configuration mode. If neither is set, no security allowing any user to view and/or change the device configuration exists for the Privilege mode. Someone could even set a password and lock out other users.

The older **enable password** command followed by the desired password creates a cleartext entry in the running configuration that could be viewed by anyone seeing the configuration. The more secure **enable secret** command followed by the desired password creates an encrypted entry in the running configuration that can't be understood by anyone just seeing the configuration. If both **enable password** and **enable secret** are configured, only the **enable secret** is used. The **enable password** is ignored.

The following entries demonstrate both commands, and then use a **show run** command to display the configuration. All passwords are case sensitive and should comply with the password policy.

```
Rtr1#conf t
Rtr1(config)#enable password test
Rtr1(config)#enable secret cisco
Rtr1(config)#^z
Rtr1(config)#show run
```

```
!
enable secret 5 $1$4F6c$D5iYCm31ri1cA9WwvAU220
enable password test
```

Notice the enable secret password can't be recognized, but the enable password is easily recognized. If only the enable password had been set, anyone seeing the configuration could get the password that would let them reconfigure the router.

Password Encryption

If you look over the previous examples, you'll notice all passwords are in cleartext, except for the enable secret password. This means prying eyes might gather a password. You can secure the passwords in Global Configuration mode by typing the **service password-encryption** command. This permanently encrypts all passwords, so make sure you know what they are. The following abbreviated output demonstrates the command and the results:

```
Rtr1#conf t
Rtr1(config)#service password-encryption            <-command
Rtr1(config)#^Z
Rtr1#sho run
version 12.0
service timestamps debug uptime
service timestamps log uptime
service password-encryption                         <-command
!
hostname Rtr1
!
enable secret 5 $1$4F6c$D5iYCm31ri1cA9WwvAU220
enable password 7 15060E1F10                       <-enable password
!
username remote password 7 070E224F4B1A0A          <-local database
username scott password 7 02110B570D03             <-local database
!
line con 0
 password 7 121A0C041104                           <-line con password
 login
line aux 0
 password 7 121A0C041104                           <-line aux password
 login
line vty 0 4
 login local
!
end
```

Note, typing **no service password-encryption** won't cause the passwords to revert to cleartext. The system will no longer encrypt new passwords, but the existing ones remain encrypted. Make sure you know a password before you encrypt it.

Message of the Day Banner (MOTD)

It's possible and prudent to create a message that will appear to everyone logging in to the User mode. This message should be a polite warning of company security policies

for unauthorized access. Some courts have held that if this isn't explicitly stated—telling people to stay out, then it's an implicit invitation to come in and raise havoc.

To configure this message, use the **banner motd** command in the Global Configuration mode. The syntax is a little unusual in that you type the command, followed by a character you don't plan to include in the message. This character becomes a delimiter, in that everything you type after it until the character appears again will be part of the message. An example would be **banner motd *No Unauthorized Access*** where the asterisks indicate the beginning and the end of the message. The asterisks won't appear in the message.

You can make multiple-line messages by using SHIFT-ENTER at the end of the line and ignoring the warning message that appears the first time you try it. The following lines demonstrate this technique. Typing a new MOTD replaces any existing one.

```
Rtr1#conf t
Rtr1(config)#banner motd * Unauthorized Access Could Result In Termination
Enter TEXT message.  End with the character '*'.          <-Ignore this

If you have any trouble with this device, call:
Mark Smith in IT Tech Support
Phone: (555) 555-1111 ext 1234*
Rtr1(config)#^Z
Rtr1#
```

The next login attempt would look like the following:

```
Unauthorized Access Could Result In Termination

If you have any trouble with this device, call:
Mark Smith in IT Tech Support
Phone: (555) 555-1111 ext 1234
User Access Verification
Password:
```

Three other options can be used with the **banner {exec | incoming | login | motd}** command that allow for variations on when the banner can be used. The MOTD option has been included because it causes the warning to appear at login for all users and is the most frequently used. You can use the help feature to see the other options.

Privilege Levels

Cisco devices numbered 0 through 15 have 16 privilege levels. By default, any user who can furnish the user-level password or user name/password combination can gain User exec mode access to the device, which is privilege level 1. From there, if the user knows the enable secret password, they can access the Privilege exec mode, or privilege level 15. The three predefined privilege levels on Cisco devices include the following:

- **1** User exec mode only (prompt is router>), the default level for login
- **15** Privileged exec mode (prompt is router#), the Enable mode
- **0** Seldom used, but includes five commands: disable, enable, exit, help, and logout

To determine or confirm the current privilege level, type the **show privilege** command. It would look like this in Privilege mode:

```
Rtr1#show privilege
Current privilege level is 15
Rtr1#
```

Privilege levels 2 through 14 can be defined by the admin to provide limited features to some users by assigning specific commands to the level using the **privilege** command.

The syntax is

privilege *mode* {level *level command* | reset *command*}, where

mode	Indicates the configuration level being assigned. This includes all router configuration modes, including exec, configure, and interface.
level	Indicates the level being defined.
command	Indicates the command to be included. If you specify exec mode, then the command must be an **exec mode** command.
reset	Resets the privilege level of the command to the default privilege level.

A possible application of this feature might look like the following lines, which are creating a new Privilege mode for a part-time administrator.

```
Rtr1(config)#privilege exec level 7 ping
Rtr1(config)#privilege exec level 7 show startup-config
Rtr1(config)#privilege exec level 7 show ip route
Rtr1(config)#privilege exec level 7 show ip int brief
Rtr1(config)#enable secret level 7 tESt7
```

The following lines show how the new privilege level would be accessed and a confirmation of the new level:

```
Rtr1>enable 7
Password:
Rtr1#show privilege
Current privilege level is 7
Rtr1#
```

Any attempt to run a command other than those specifically defined for this privilege level returns the same error message as any attempt to run a command from the wrong mode. As you will see in Chapter 4, AAA authentication provides some additional options for this feature.

Note that the privilege feature only limits user access if the user only knows the enable secret password for the defined level. If the user knows any other level password, then they can go there as well.

Using Access Control Lists to Secure the Network

Access control lists (ACLs) are powerful tools that are often at the heart of many other processes, including security processes. ACLs can be a good starting point for adding security and traffic management to a network, but they can't protect the network by themselves. Devices like firewalls and proxy servers, as well as good security practices—password management, physical security, and solid administrative policies—should be used to augment them.

STUDY TIP While ACLs aren't a specific exam objective, they're fundamental in many security implementations and processes. Newer operating system (OS) versions for devices like PIX firewall boxes and Catalyst switches use Cisco IOS–like features and commands and, therefore, incorporate access list capabilities. Make sure you're comfortable with how access lists are constructed and how they work. If you need a refresher on ACLs, go to Appendix A.

Standard ACLs

This section includes examples of standard IP ACLs with various wildcard Mask options to perform the following security-related tasks.

- Traffic filtering
- Limiting access to Telnet sessions
- Limiting access to HTTP sessions
- Filtering routing updates
- Limiting the **debug ip packet** analysis and, therefore, CPU use

Before looking at these examples, let's look at a technique that can be used in test labs to make a limited pool of devices look more diverse. In some of the next exercises, having more than one LAN interface on the routers would be useful, but this isn't required. The next section looks at using loopback interfaces as a simple work-around in the lab.

Using Loopback Interfaces on Lab Routers

What do you do if your router has only a single Ethernet interface? Or worse, what if one of your lab routers has only a Token Ring LAN interface and you don't have a Multistation Access Unit (MAU) unit to plug it in to, so the interface won't come up? Short of getting another device, the only solution is to create loopback interface(s). What are loopback interfaces?

- They are logical interfaces, in that they don't connect to any cables and, therefore, they have no hosts. Loopback interfaces are logical interface only.

- They are "up" by default and remain up unless administratively shut down with the **shutdown** command.
- They can be configured as a host, a subnet, a network, or a supernet, like any other interface.
- They can be included in the routing updates like any other interface.
- They can be used as a ping, a Telnet destination, or a source.
- Some routing protocols, such as OSPF and BGP, use loopbacks as router IDs.

Configuring a loopback interface is similar to creating a subinterface. The syntax is

router(config)#interface loopback *num*
router(config-if)#ip address *ip-address subnet-mask*

The *num* is any whole number between 0 and 2147483647. It won't accept module-type notation, like 0/1.

Note, the **no shutdown** command isn't required and would only be used with loopback interfaces following a **shutdown** command.

Here is an example:

```
router(config)#interface loopback 0
router(config-if)#ip address 192.168.3.1 255.255.255.255
```

The Layer 3 addresses can be either an IP or an IPX interface configuration (AppleTalk isn't supported on all router platforms). The example is a 32-bit mask, called a *host mask*, because it specifies a single host, not a network or a subnetwork. Some routing protocols, such as Open Shortest Path First (OSPF), use loopback interfaces as router IDs and prefer the host mask to avoid problems. With protocols such as RIP, IGRP, and EIGRP, you can use any mask that serves your purpose.

Traffic Filtering

Traffic filtering is a common use for standard access lists used when the plan is to block all packets from a specific source host or group of hosts from reaching a portion of the network. Figure 2-3 shows a simple two-router network that might represent two branch locations of a smaller business.

The following code represents two access lists created on the Rtr1 router. ACL 10 allows only a single host (192.168.2.20) from the Rtr2 LAN and all hosts from the Rtr1 LAN to go out to the Internet. By not adding a permit any statement, all other hosts are denied.

```
Rtr1(config)#access-list 10 permit host 192.168.1.20.0.0.0.255
Rtr1(config)#access-list 10 permit 192.168.1.20 0.0.0.255
Rtr1(config)#access-list 20 permit host 192.168.2.20
Rtr1(config)#access-list 20 deny 192.168.2.0 0.0.0.31
Rtr1(config)#access-list 20 permit any
Rtr1(config)#int s0
Rtr1(config-if)#ip access-group 10 out
Rtr1(config-if)#int e1
Rtr1(config-if)#ip access-group 20 out
```

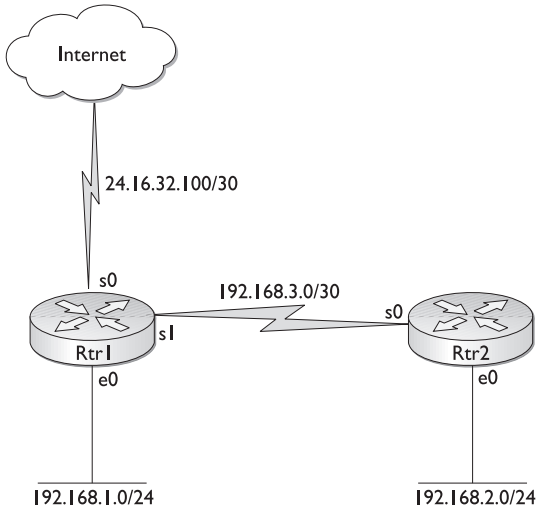

Figure 2-3 Two-router network with two ACLs

The first line of ACL 20 permits host (192.168.2.20) from the Rtr2 LAN to access the LAN on Rtr1. The second line blocks the rest of 192.168.2.0 subnet mask 255.255.255.224, or addresses 192.168.2.0 to 192.168.2.31 from the Rtr1 LAN. The final line allows the rest of Rtr2 LAN and anything coming in over the Internet.

Placement of Standard ACLs Most administrators deduce quickly that if packets can be discarded as early as possible, this will reduce bandwidth requirements on links and router CPU cycles used to process packets that are going to be dumped anyway. Unfortunately, with standard ACLs, the only criterion for making the determination to permit or deny is the source address—the destination is unknown. Standard access lists are generally placed as close to the destination network segment as possible to exercise the most control. For example, in Figure 2-3 in the preceding section, a standard ACL blocking a certain Internet address from entering the Rtr2 LAN would need to be placed on the Rtr2 e0 interface outbound. While s0 on Rtr1 would save more router resources, it would also block access to the Rtr1 LAN. Similarly, applying it to s0 on Rtr2 would work for now, but what if additional interfaces are used later? They would also be blocked.

Looking at Figure 2-3, if ACL 20 had been placed inbound on the Serial 1 interface of Rtr1 or outbound on s0 of Rtr2, this would have saved resources, but this would also have blocked access to the Internet for hosts 192.168.2.1 to 192.168.2.19 and 192.168.2.21 to 192.168.2.30.

Log Option Since v11.3 of the IOS, the Log option at the end of an ACL statement results in logging packets that meet the ACL criterion. How and where the logging is stored is controlled with the **logging console** command. The simplest form of the Log option causes a message to be printed to the console screen the first time the ACL is activated, and then every five minutes while the ACL is still being used. The first report indicates only the first packet, while the subsequent reports summarize the number of occurrences.

The following output lines show the result of adding the Log option to an ACL that blocks the access of host 192.168.1.10 to a LAN.

```
Rtr1(config)#access-list 50 deny   192.168.1.10 log
Rtr1config)#access-list 50 permit any
Rtr1(config)#int e0
Rtr1(config-if)#ip access-group 50 out
Rtr1(config-if)#^Z
Rtr1 #
11:29:37: %SEC-6-IPACCESSLOGS: list 50 denied 192.168.1.10 1 packet
Rtr1 #
11:34:53: %SEC-6-IPACCESSLOGS: list 50 denied 192.168.1.10 27 packets
Rtr1#
```

In the example, the last two lines indicate the workstation attempted to ping a LAN host (192.168.5.1) seven times from the Command window. The first packet was recorded in the first log entry. The other 27 packets, 4 per ping, were reported five minutes later.

The **logging console** command in Global Configuration mode can be used to modify the Log options, which includes specifying a host address running Syslog server software to forward all log entries for permanent storage. Chapter 5 has a section demonstrating the various logging options.

Limiting Access to Telnet Sessions

Using Figure 2-4, assume the company wants to limit Telnet access to Rtr1 to users on the Rtr2 LAN and a single outside address 42.17.17.45, which is the router administrator's home computer. Just as it's possible to filter the physical interfaces, such as Ethernet 0 and serial 1, standard access lists can be used to filter the virtual ports, thereby limiting telnet access to the routers.

Five virtual ports, or vty lines, typically are designated as vty 0 4, allowing up to five telnet sessions to be established. Similar to setting Telnet passwords, you can set identical restrictions on all vty lines at one time.

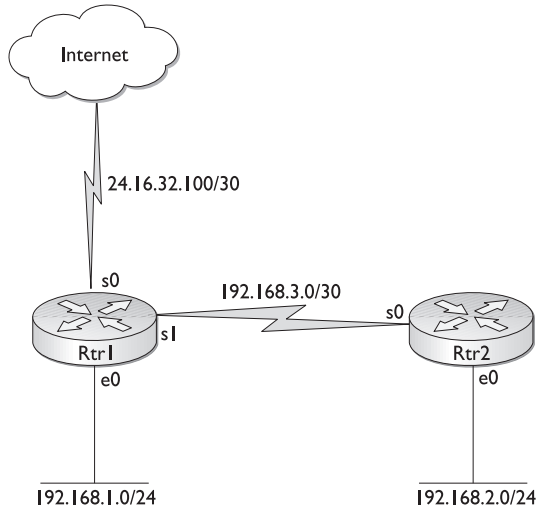

Figure 2-4 Two-router network for two branch locations

In the following output, inbound access list 15 limits access to the router to users specified by the company. The lack of a permit any statement limits the access to those hosts defined in lines one and two. ACL 16 prevents any user who successfully Telnetted into the router to Telnet to another device. They can only exit out after doing their assigned tasks.

```
Rtr1(config)#access-list 15 permit 192.168.2.0 0.0.0.255
Rtr1(config)#access-list 15 permit 42.17.17.45
Rtr1(config)access-list 16 deny any
Rtr1(config)#line vty 0 4
Rtr1(config-line)#access-class 15 in
Rtr1(config-line)#access-class 16 out
Rtr1(config-line)#password cisco
Rtr1(config-line)#login
```

The important keyword here is "access-class," which is used in place of access-group when you're applying an ACL to a virtual interface.

The following code shows the result when another router, or a host on any other subnet/network, tries to Telnet into Rtr1. It's important to realize that this implementation of the access-class applies to the entire router, not only to one direction on an interface.

```
Rtr2>telnet 192.168.3.1           <-could have used 192.168.1.1
Trying 192.168.3.1...
% Connection refused by remote host
```

The following lines show the result of a permitted user (who cleared access-list 15 and supplied the password) trying to telnet onto another router. The output assumes an IP HOST table was set up with both IP addresses associated with the name Rtr2 (Rtr1(config)#ip host rtr2 192.168.3.2 192.168.2.1). Telnet attempts both interfaces, but is denied at each attempt.

```
Rtr1>telnet Rtr2
Trying Rtr2 (192.168.3.2)...
% Connections to that host not permitted from this terminal
Trying Rtr2 (192.168.2.1)...
% Connections to that host not permitted from this terminal
Rtr1>
```

Show Line VTY Command The **show line vty 0 4** command lists the virtual terminals and shows that the access lists are applied. The **Acc0** (ACL outbound) and **AccI** (ACL inbound) display the ACL numbers. Running the command to verify that your access lists are, in fact, applied to all virtual terminals is always a good idea.

```
Lab-X#show line vty 0 4
 Tty Typ   Tx/Rx  A Modem  Roty AccO AccI   Uses   Noise  Overruns   Int
   2 VTY          -    -      -   16   15      0       0     0/0       -
   3 VTY          -    -      -   16   15      0       0     0/0       -
   4 VTY          -    -      -   16   15      0       0     0/0       -
   5 VTY          -    -      -   16   15      0       0     0/0       -
   6 VTY          -    -      -   16   15      0       0     0/0       -
```

Limiting Access to HTTP Sessions

Many Cisco devices allow access to the User and Privilege modes via a web browser. Figure 2-5 shows an example of the opening screen for a 2500 model router. While not as extensive or friendly as some manufacturer's web interfaces, it does enable the user to view the interfaces, view the diagnostic log, ping the device, execute commands, and open a telnet session.

For this feature to work, the **ip http server** command must be issued in Global Configuration mode because the default is for the service to be off. The command allows the device to act as an HTTP server. Many company security policies require the feature to be turned off by issuing a **no ip http server** command to reduce the exposure to unauthorized access.

Another alternative is to use a standard ACL to limit which host(s) can access the device from a web interface. The basic steps include the **ip http server** command, any

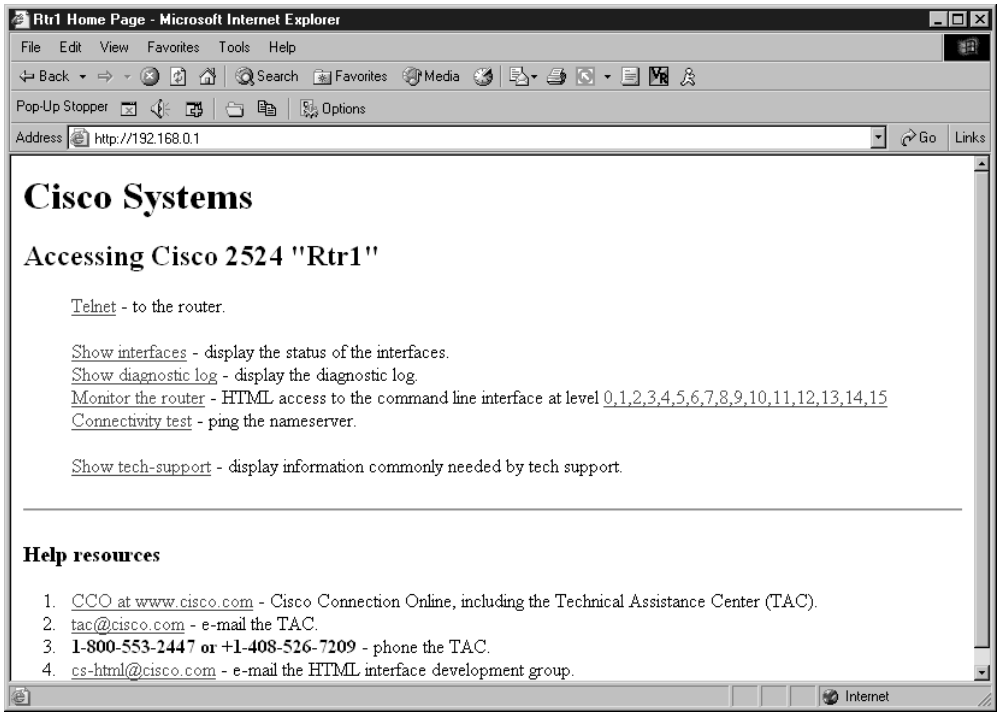

Figure 2-5 Web browser access to a 2500 router

ACL statement(s) required to define the authorized host, and the **ip http access-class {*acl-num* | *acl-name*}** command. The following lines show an example of the commands:

```
Rtr1#conf t
Rtr1(config)#ip http server
Rtr1(config)#access-list 90 permit 192.168.0.10
Rtr1(config)#access-list 90 permit 192.168.5.0 0.0.0.31
Rtr1(config)#access-list 90 permit 192.168.45.0 0.0.0.255
Rtr1(config)#ip http access-class 90
Rtr1(config)#^Z
Rtr1#
```

The lack of a permit any statement limits access to those specifically defined in the permit statement(s). Any time the device is accessed with a web browser, a login screen appears asking for a user name and password. With the previous configuration, the device host name (Rtr1) would be the user name and the enable secret password would be the password.

Controlling the access further is possible with the **ip http authentication** command, which supports several methods of establishing the user name/password combinations. The syntax would be:

ip http authentication {aaa | enable | local | tacacs}

AAA	The AAA authentication feature covered in Chapter 4
Enable	Uses the enable password for authentication (the default HTTP)
Local	Uses the local user database (user name/password) defined in Global mode
Tacacs	The TACACS or XTACACS server authentication covered in Chapter 5

Limiting the debug ip packet Analysis and, Therefore, CPU Use

The **debug ip packet** command displays in real-time all IP packet activity passing through a router. This is, unfortunately, extremely hard on CPU resources. In fact, a router running at 50 percent capacity can be buried by the command dominating the CPU use. Many organizations ban all use of the command for this reason. A second problem with the command is the screen is often overwhelmed by output, making it difficult to see the expected information.

Fortunately, the command syntax, **debug ip packet [*acl-num*]**, allows an ACL to be created to filter the specific traffic of interest. In the following example, the output is limited to any traffic originating in the 192.168.0.0 network. Any valid host address, subnet, or network address can be used.

```
Rtr1#conf t
Rtr1(config)#access-list 20 permit 192.168.0.0 0.0.0.255
Rtr1(config)#^Z
Rtr1#debug ip packet ?                          <-shows the options
  <1-199>      Access list
  <1300-2699>  Access list (expanded range)
  detail       Print more debugging detail
  <cr>
```

```
Rtr1#debug ip packet 20
IP packet debugging is on for access list 20
Rtr1#
00:05:28: IP: s=192.168.0.1 (local), d=224.0.0.10 (Ethernet0),
    len 60, sending broad/multicast
00:05:28: IP: s=192.168.0.10 (Ethernet0), d=192.168.0.1 (Ethernet0),
    len 56, rcvd 3
```

The options shown in the **debug ip packet ?** indicate extended ACLs could be used to filter, based on destination address, protocol, or port numbers.

Extended Access Lists

Extended access lists provide a higher level of traffic control by being able to filter packets based on the protocol, source and/or destination IP address, and source and/or destination port number. For example, an extended access list can block an address (or group of addresses) in a particular network from accessing the FTP services on a specific server, while still allowing other services.

Restrict Application Traffic

Implicit in using TCP and UDP access list statements is that it becomes possible to support or suppress certain higher level applications. Because port numbers are associated with applications, allowing or denying access to a specific port number determines if that application can be used and which devices can access it. Setting up network segments to carry only one or two types of traffic—say, NNTP news or SMTP mail—becomes possible. Access lists can be established to prevent any other traffic from entering that particular segment. These ACLs not only focus on the source and destination of the packets, but also on the service offered. The following table shows some common port numbers and the related services (applications).

ftp	File Transfer Protocol	20, 21
Telnet	Telnet, Remote Terminal	23
smtp	Electronic Mail (servers)	25
pop3	POP Mail (users)	110
nntp	Network News	119

More examples are included in the sections dealing specifically with TCP and UDP filtering.

TCP's Established Option The Established option is a TCP-only feature that can use the connection-oriented attributes of the TCP to limit traffic coming into a network or network segment to those sessions that originated from within that network. The established condition is only true if the ACK (acknowledge) or RST (reset) bits are set to one in the TCP header, indicating an already established connection. A packet with no ACK or RST bit set, but a SYN (synchronize) bit set to one, is used to establish a new connection and can then be denied. Figure 2-6 demonstrates the three-step "handshake" that TCP uses to establish a connection.

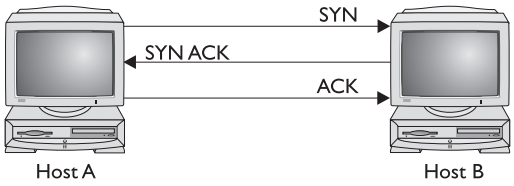

Figure 2-6 TCP three-way handshake to establish a session

The following output demonstrates allowing any host to respond to web and Telnet requests that originated within the 192.168.1.0 network, but blocks all other TCP packets.

```
access-list 101 permit tcp any 192.168.1.0 0.0.0.255 established
access-list 101 permit tcp any 192.168.1.0 0.0.0.255 any eq telnet
access-list 101 deny   tcp any any
access-list 101 permit ip  any any
```

The Established option can help reduce the risk of a common type of hacker attack that buries a host in SYN requests, preventing it from handling normal business. Because a sophisticated hacker can manipulate the TCP header bits, this tool needs support from other tools to protect against that threat.

Look over the following ACL statement using the Established option. This is a common first effort when trying to limit web activity to those sessions originating within the network. So what's wrong with the statement?

```
access-list 101 permit tcp any 192.168.1.0 0.0.0.255 eq www established
```

Remember, www is an "alias" for port 80. A web session originating inside would use port 80 as the destination, but would designate a port above 1024, such as 1065, as the source port. This means the returning packet would have port 80 as the source and port 1065 as the destination. The ACL is looking for port 80 as the destination. The following output might work better:

```
access-list 101 permit tcp any eq www 192.168.1.0 0.0.0.255 established
```

When working with the Established feature, it's important to make sure you understand what the mnemonic and any port numbers stand for, if you use the port numbers. Another approach is represented in the following code lines, which allow any established sessions, but block all other TCP traffic.

```
access-list 101 permit tcp any 192.168.1.0 0.0.0.255 established
access-list 101 deny   tcp any any
```

NOTE *Source-port filtering*, the process of filtering data on the source port of a packet, isn't secure because a skilled hacker could easily change a source port on a packet, which could then pass through the filter.

Named Access Lists

Since Cisco IOS Release 11.2, it's possible to use a text name for access lists in many cases. Some new features were added to named ACLs that make them more than just text names. The benefits of using named access lists are described in the following paragraph:

Time-Based Access Lists

Until recently, an access list was either implemented or not. An administrator would have to change the configuration to include a feature, such as blocking Internet access, and then would have to physically remove either the ACL or the access-group statement to allow access. This wasn't convenient if the goal was to block employee access during work hours, but to allow access at other times of the day.

Since version 12.0 of the IOS, it's possible to create access lists that can incorporate date- and time-sensitive features. This allows filtering to occur only during certain hours of the day or on certain days of the week. IP and IPX extended access lists are currently the only functions that can use time ranges. *Time ranges* enable the network administrator to specify when the permit or deny statements in the access list are in effect. Both named or numbered access lists can use the time range feature.

Two separate tasks are required to implement time ranges:

- Define a time range.
- Reference the time range.

Define a Time Range

Three basic commands are used to implement the time-based features in extended ACLs:

- time-range
- periodic
- absolute

The **time-range** command is used to create (name) a time range, while the absolute and periodic commands are used to define the days and times included in the named range. While multiple periodic entries are allowed per time-range command, only one absolute entry is allowed.

If a **time-range** command has both absolute and periodic values specified, then the periodic items are evaluated only after the absolute start time is reached, and aren't evaluated any further after the absolute end time is reached.

The time-range Command

The **time-range** command is used to enable Time-Range Configuration mode and name time ranges for functions, such as extended access lists. The command is issued in Global Configuration mode. The syntax is:

```
Router(config)#time-range time-range-name
```

time-range-name	User-specified name of the time range

The following output demonstrates defining a time range named NoDayAccess. Remember, the name *is* case sensitive.

```
Router(config)#time-range NoDayAccess
```

The periodic Command

The **periodic** command is used to specify a recurring range of days and times to be included in the time-range created in the previous command. Using this command allows a feature to be implemented on only certain days of the week and within certain hours. The command is issued in the Time-Range Configuration mode, which is activated by the **time-range** command. The syntax is

Router(config-time-range)#periodic *days-of-the-week hh:mm* to [*days-of-the-week*] *hh:mm*

days-of-the-week	Used to define the days included in the time-range
hh:mm	Defines the beginning or ending time of the time range. *hh* is the 24-hour time and *mm* is the minute.

The first occurrence of the *hh:mm* argument is the starting hours:minutes that the time range is in effect. The second occurrence is the ending hours:minutes the statement is in effect. Times are expressed using the 24-hour clock; therefore, 8:00 P.M. would be typed as **20:00**. The keyword **to** is required to complete the range, as in ***start-time*** *to* ***end-time***.

NOTE Remember, all time specifications are interpreted as local time. To ensure the time range entries take effect at the desired times, the software clock should be synchronized using the Network Time Protocol (NTP) or some other authoritative time source.

The *days-of-the-week* argument can seem complicated until you become used to it. Used properly, this can be used to represent a single day or combination of days using the following keywords:

Monday, **Tuesday**, **Wednesday**, **Thursday**, **Friday**, **Saturday, Sunday**—day names

daily Monday through Sunday

weekdays Monday through Friday

weekend Saturday and Sunday

The following examples should help you understand the process. The following lines create a time range for a single day named MondayDays, and then define the day and hours (8:00 A.M. to 5:30 P.M.).

```
Router(config)#time-range MondayDays
Router(config-time-range)#periodic monday 8:00 to 17:30
```

When multiple days can be defined by the keywords daily, weekday, and weekend, the keyword is used. The next lines creates a time range named WeekdayEves, and then

define hours (5:00 P.M. to 10:00 P.M.). In this example, the feature is to be implemented Monday through Friday.

```
Router(config)#time-range WeekdayEves
Router(config-time-range)#periodic weekdays 17:00 to 22:00
```

When multiple days other than those defined by the keywords daily, weekday, and weekend are required, the days are listed separated by spaces. The next lines create a time range named TuThuEves and then defines the days and hours (5:00 P.M. to 10:00 P.M.).

```
Router(config)#time-range TuThuEves
Router(config-time-range)#periodic tuesday thursday 17:00 to 22:00
```

With a recurring time range on consecutive days, other than those defined by the keywords, the days are also listed separated by spaces. The next lines create a time range named SunMonTue and then define the hours (5:00 P.M. to 10:00 P.M.) included each day.

```
Router(config)#time-range SunMonTue
Router(config-time-range)#periodic sunday monday tuesday thursday 6:00 to 9:00
```

With a time range that "spans" consecutive days, the optional *days-of-the-week* after the **to** is used. The next lines create a time range named MonToWed, and then define the start and ending time as 8:00 A.M. Monday to Wednesday at 5:00 P.M.

```
Router(config)#time-range MonToWed
Router(config-time-range)#periodic Monday 8:00 to Wednesday 17:00
```

Multiple periodic entries are allowed per **time-range** command.

The absolute Command

Use the **absolute time-range** command to specify a time range to be in effect each day. The command is issued in the Time-Range Configuration mode, which is activated by the **time-range** command. The syntax is

Router(config-time-range)#absolute [start *time date*] [end *time date*]

The start and end times use the 24-hour clock. The dates are in a text form, such as June 1, 2003.

The following sample defines a time range from 8:00 December 18, 2003, until 8:00 January 3, 2004:

```
Router(config)#time-range holidays-closed
Router(config-time-range)#absolute 8:00 18 december 2003 8:00 3 january 2004
```

Only one absolute entry is allowed per **time-range** command. If a **time-range** command has both absolute and periodic values specified, then the periodic items are evaluated only after the absolute start time is reached, and aren't evaluated any further after the absolute end time is reached.

Reference the Time Range

An extended access control list can be either a numbered or a named ACL. In each case, the Time-Range option is added to provide an additional qualifier for the permit|deny statement. The basic syntax for a named ACL would be

> Router(config)#ip access-list extended *name*
> Router(config-ext-nacl)# {permit | deny} {*protocol* | *protocol-keyword*}{*source wildcard* | any} [*operator source-port*] {*destination wildcard* | any} [*operator destination-port*] [precedence *precedence*] [tos *tos*] [log | log-input] [options] **[time-range *time-range-name*]**

The basic syntax for a numbered ACL would be

> Router(config)#access-list *acl#* {permit | deny} {*protocol* | *protocol-keyword*}{*source wildcard* | any} [*operator source-port*] {*destination wildcard* | any} [*operator destination-port*] [precedence *precedence*] [tos *tos*] [log | log-input] [options] **[time-range *time-range-name*]**

The following time range example with a periodic statement denies web traffic to employees on the Ethernet LAN for Monday through Friday during business hours (8:00 A.M. to 6:00 P.M.).

```
time-range no-web
  periodic weekdays 8:00 to 18:00
!
ip access-list extended block-web
  deny tcp any any eq www time-range no-web
  permit ip any any
!
interface ethernet 0
  ip access-group block-web in
!
```

The following example uses a time-based access list to allow a LAN to begin accessing the network beginning at 8:00 A.M. on January 1, 2003. The access will continue until stopped. The access list and time range together permit traffic on Ethernet interface 0 starting.

```
time-range new-lan
  absolute start 8:00 1 January 2003
!
ip access-list extended start-service
  permit ip 192.168.15.0 0.0.0.255 any time-range new-lan
!
interface ethernet 0
  ip access-group start-service in
```

The following example uses a time-based access list to block a LAN from accessing the network beginning at midnight on December 31, 2003.

```
time-range stop-lan
  absolute end 23:59 31 December 2003
!
ip access-list extended stop-service
  permit ip 192.168.15.0 0.0.0.255 any time-range stop-lan
!
interface ethernet 0
  ip access-group stop-service in
```

The following example uses a time-based access list to permit weekend employees to browse the Internet for a two month test period from 8:00 A.M. on June 1, 2003, to 6:00 P.M.on July 31, 2003.

```
time-range web-test
absolute start 8:00 1 June 2003 end 18:00 31 July 2003
periodic weekends 00:00 to 23:59
!
ip access-list extended lan-web
  permit tcp 192.168.15.0 0.0.0.255 any eq www time-range web-test
!
interface ethernet 0
  ip access-group lan-web in
```

The show time-range Command

The **show time-range** command can be used to display the current status of the time ranges and list the periodic and absolute attributes.

```
Rtr1#show time-range
time-range entry: WeekdayEves (inactive)
   absolute start 08:00 01 December 2004 end 08:00 03 January 2005
   periodic weekdays 17:00 to 22:00
time-range entry: weekdayeves (active)
   used in: IP ACL entry
time-range entry: worktime (active)
   absolute start 08:00 01 December 2004 end 08:00 03 January 2005
   periodic weekdays 8:00 to 18:00
   used in: IP ACL entry
Rtr1#conf t
```

The output shows two defined time ranges **WeekdayEves** and **worktime**, but it also shows another undefined time range, **weekdayeves**. The output shows it was used in an ACL, indicating someone tried to use **WeekdayEves** and forgot that the names are case sensitive. The result in this case will be that since no times are defined, the ACL statement will be active all of the time.

Chapter Review

Many simple device configuration techniques can add to the security of the network. To a great extent, these often fall into the category of commonsense practices, such as using administrative access passwords on all device access points.

As Cisco moves more and more devices to IOS-based command structures, access lists remain a need-to-know technology. While not a complete security solution, access lists are an integral part of any security program.

Standard access lists filter based on source address alone, creating a simple, yet powerful, tool for blocking all traffic or access to a host, subnet, or network. Standard ACLs can be used for traffic filtering, limiting access to Telnet sessions, limiting access to Web browsers trying to access a Cisco router or switch, filtering routing updates, and focusing commands like **debug ip packet** to conserve router resources.

Extended access lists can be used to filter on protocol, source address, destination address, source and destination port identifiers for TCP and UDP traffic, and various powerful options. The TCP Established option can be used to limit TCP traffic only to what originated within the network.

Named access lists are a variation on the numbered ACLS supporting for standard and extended versions. Named ACLs are easier to create than numbered lists, and allow limited editing and deletion of specific statements that can't be done with numbered lists. They can be descriptive of their purpose and, therefore, easier for follow-up support to work with. Some IOS features and all IOS versions prior to 11.2 don't support named ACLs, requiring some thought in mixed environments. Some newer features like reflexive ACLs only work with named lists, so it's probably safe to say they're going to be a bigger, rather than smaller, part of the future.

Questions

1. Which of the following interface types is least likely to be on a firewall appliance?
   - **A.** Fast Ethernet
   - **B.** Serial
   - **C.** Ethernet
   - **D.** Token Ring
2. Which of the following would not be considered a basic security step in a router configuration?
   - **A.** Setting access privilege levels
   - **B.** Setting an MOTD banner to welcome the user to the device
   - **C.** Encrypting passwords in the configuration files
   - **D.** Setting all passwords
3. Which of the following is *not* true about numbered access lists?
   - **A.** An ACL is made up of one or more permit or deny statements.
   - **B.** If an ACL doesn't have at least one permit statement, it will deny everything.
   - **C.** All ACL statements with the same number are part of the same ACL.
   - **D.** New statements are always added to the top of the list statements.
   - **E.** ACL statements must be entered sequentially to be processed properly.
   - **F.** An ACL can be added to (appended), but not edited. Any attempt to edit an item will delete the entire ACL.
4. Which one of the following will deny access to a class C network?
   - **A.** Rtr1(config)#access-list 15 deny 192.168.1.0 255.255.255.0
   - **B.** Rtr1(config)#access-list 15 deny 192.168.1.0 0.0.0.255

C. Rtr1(config-acl)#access-list 15 deny 192.168.1.0 255.255.255.0

D. Rtr1(config-acl)#access-list 15 deny 192.168.1.0 0.0.0.255

5. What is the ACL line to deny the subnet 192.168.1.16 subnet mask 255.255.255.240?

A. access-list 15 deny 192.168.1.16 0.0.0.255

B. access-list 15 deny 192.168.1.0 0.0.0.16

C. access-list 15 deny 192.168.1.16 0.0.0.15

D. access-list 15 deny 192.168.1.16 0.0.0.31

6. With the Log option for ACLs, a message appears when the first match occurs, and then at what interval as long as matches continue?

A. One minute

B. Five minutes

C. Ten minutes

D. Thirty minutes

7. When limiting access to Telnet sessions, which command would work?

A. access-group 15 in

B. access-group 15 out

C. access-class 15 in

D. access-class 15 out

E. access-session 15 in

8. Which two commands could be used to secure the web browser access to a device?

A. ip http server

B. no ip http server

C. ip http access-class 90

D. ip http access-group 90

9. Which of the following protocols uses the established option?

A. UDP

B. ICMP

C. TCP

D. IGRP

10. Numbered extended ACLs are created in which mode?

A. Privilege mode

B. Global Configuration mode

C. Local Configuration mode

D. Access Configuration mode

11. Which statement is *not* true about named access lists?

A. Named access lists aren't compatible with older IOS releases (pre-11.2).

B. A standard access list and an extended access list can't have the same name.

C. Names must begin with an alphanumeric character and are case sensitive.

D. All processes that use access lists can use a named ACLs

12. Which statement will create a named extended ACL?

A. Rtr1(config)#ip extended access-list tcp-control

B. Rtr1(config)#ip access-list named extended tcp-control

C. Rtr1(config)#ip access-list extended tcp-control

D. Rtr1(config-ext-nacl)#ip access-list extended tcp-control

13. The time-based ACL statements are relative to which one of the following?

A. The computer clock

B. The world clock

C. The router clock

D. The day, month, and year

14. Which command will define a periodic time range?

A. Router(config-time-range)#periodic tuesday thursday 17:00 to 22:00

B. Router(config-time)#periodic tuesday thursday 17:00 to 22:00

C. Router(config-time-range)#periodic tuesday, thursday 17:00 to 22:00

D. Router(config-time)#periodic tuesday, thursday 17:00 to 22:00

15. Which statement is true about defining a time range?

A. A time range can have either periodic or absolute times.

B. A time range can have one periodic and multiple absolute times.

C. A time range can have multiple periodic and one absolute time.

D. A time range can have one periodic and one absolute time.

16. Which one of the following is true about the Established option in a TCP access list?

A. Outbound traffic is limited to established customers.

B. Outbound traffic is limited to sessions originating outside the network.

C. Inbound traffic is limited to sessions originating outside the network.

D. Inbound traffic is limited to sessions originating inside the network.

17. In the following ACL, what is the impact of the third statement?

```
access-list 101 deny tcp 192.168.3.0 0.0.0.255 192.168.1.0 0.0.0.255
access-list 101 deny tcp 192.168.3.0 0.0.0.255 any eq ftp
access-list 101 permit tcp 192.168.3.0 0.0.0.255 192.168.1.0 0.0.0.255 eq www
access-list 101 deny tcp any 192.168.1.0 0.0.0.255 any eq telnet
access-list 101 permit ip any any
```

A. It allows network 192.168.3.0 to access 192.168.1.0 for web access.

B. It allows network 192.168.1.0 to access 192.168.2.0 for web access.

C. It allows network 192.168.3.0 to access any network for web access.

D. The line does nothing at all.

Answers

1. **B.** Serial. Firewall device, such as the PIX box, use LAN interfaces
2. **B.** Setting a MOTD banner to welcome the user to the device
3. **D.** New statements are always added to the top of the list statements (They're actually appended to the bottom of the list.)
4. **B.** Rtr1(config)#access-list 15 deny 192.168.1.0 0.0.0.255
5. **C.** access-list 15 deny 192.168.1.16 0.0.0.15
6. **B.** Five minutes
7. **C.** access-class 15 in
8. **B.** no ip http server; c. ip http access-class 90
9. **C.** TCP
10. **B.** Global Configuration mode
11. **D.** All processes that use access lists can use a named ACL.
12. **C.** Rtr1(config)#ip access-list extended tcp-control
13. **C.** The router clock. If the router clock is wrong, the statements will be implemented wrong.
14. **A.** Router(config-time-range)#periodic tuesday thursday 17:00 to 22:00
15. **C.** A time range can have multiple periodic and one absolute time
16. **D.** Inbound traffic is limited to sessions originating inside the network.
17. **D.** The line does nothing at all. All TCP traffic from 192.168.3.0 to 192.168.1.0 was denied in the first statement

PART II

Securing the Network Perimeter

- **Chapter 3** Cisco AAA Security Technology
- **Chapter 4** CiscoSecure ACS and TACAS+/Radius Technologies
- **Chapter 5** Securing Cisco Perimeter Routers
- **Chapter 6** IOS Firewall Feature Set—CBAC
- **Chapter 7** IOS Firewall—Intrusion Detection System
- **Chapter 8** IOS Firewall—Authentication Proxy

CHAPTER 3

Cisco AAA Security Technology

In this chapter, you will learn how to:

- Describe the Cisco AAA model
- Describe and configure the AAA system and each of its three components
- Test the NAS AAA configuration using applicable **debug** and **show** commands

Cisco IOS software provides features for simple access control, such as local user name authentication, line password authentication, and enable password authentication. These methods, represented in the following output, are commonly used when you first learn to work with the Cisco IOS.

```
no service password-encryption
enable secret 5 $1$s3/7$C4ngFihNBDwqlmdj1
!
username xyzdotcom password cisco
!
line con 0
 password cisco
 login
line aux 0
 password cisco
 login
line vty 0 4
 login local
end
```

These features don't provide the same degree of access control that's possible by using AAA, however. In this chapter, you look at the methods used by Cisco's AAA security technology to control remote access to the network. While AAA is supported on many devices and can be expected on each of the certification exams, this chapter introduces AAA within the context of securing the Network Access Server (NAS). This traditionally has been called "securing the dial-up connection," using NAS routers to handle the incoming modem and ISDN sessions. As you see in the chapters that cover PIX and PIX IOS, AAA uses the same methods to control remote access from the Internet.

The Cisco AAA Model

Cisco's AAA network security services provide the principal structure to manage who has access to the network, what services they're permitted to use, and a record of what they did while logged on. The three components of AAA include the following:

Authentication	Identifies and verifies the remote user. This process can take the form of a simple password or user name/password combination verified locally by the NAS, or it could use one or more authentication servers and technologies, such as one-time passwords (OTPs) or tokens.
Authorization	Determines what devices, features, or services a specific remote user is authorized to access in the network, such as network resources or services. This concept is much like that of user permissions in the Windows server model.
Accounting	Allows the network administrator to define a process for tracking the services remote users are accessing. The data collected can be used for client billing, auditing, or network planning.

NAS Servers

NAS servers, sometimes simply called *access servers,* are routers with interfaces designed to service the remote users of the company. A simple example would be the older 2501s with two serial interfaces connected to two modems, or the 2503 with two serial and an ISDN interface to provide access service to a small number of dial-in users, such as employees dialing in from their motel rooms.

Other combinations include the 2509–2512 series that have one or two special 68-pin interfaces that can each support up to eight modem connections with an octal cable. As modem service is being replaced by DSL and cable connections, these units are dropping in price and are becoming popular today as terminal servers. As a terminal server, a single console connection to an access server can be used to provide up to 16 console connections to other devices.

NOTE While this isn't a CCSP exam requirement, a lab (Terminal Server.pdf) on this book's CD-ROM demonstrates configuring and using a terminal server to save time in the lab or production network.

Most of the common modular routers have, or have had, interfaces with multiple serial and/or multiple ISDN interfaces to provide modem and ISDN access service to larger organizations and early ISPs. Figure 3-1 shows a four-port serial interface for a 4000 series router that could connect to four modems or four ISDN devices. The 4000s also offered similar modules with four or eight ISDN interfaces.

Figure 3-2 shows an NM-16A module for a 2600 or 3600 series router with two of the 68-pin serial interfaces mentioned earlier. With octal cables, these two connectors support 16 connections. Four of the interfaces on the NM-32A allow 32 connections.

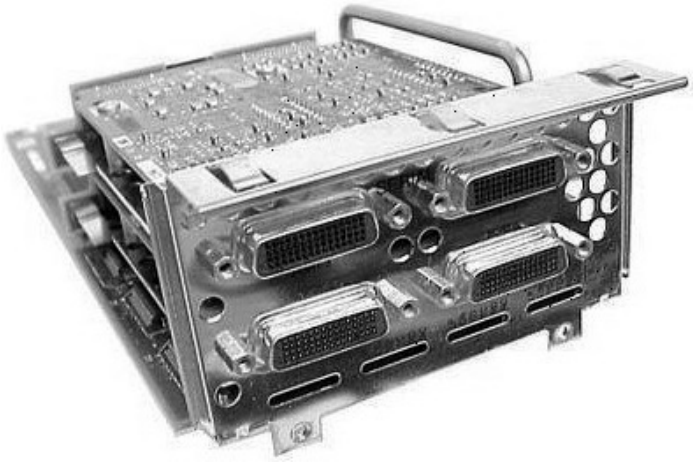

Figure 3-1
Four-port serial module for a 4000 series router

At the top end, the AS5x00 series routers can support dozens to thousands of outside connections in a variety of combinations. The AS5400 and AS5800 routers provide the capability to terminate ISDN, 56K analog modem, fax, and VoIP calls on the same interface. The AS5400 can support over 600 connections, while the AS5800 can support 2,000.

For more information, go to http://www.cisco.com on the Web and choose the link for Products. Access Servers are listed under product categories. No CCO account is needed for much of the information.

Why Authenticate?

Not too long ago, PCs didn't require a user to type in their user name and/or a password. The computer powered up, and whoever was sitting at the keyboard had full access to

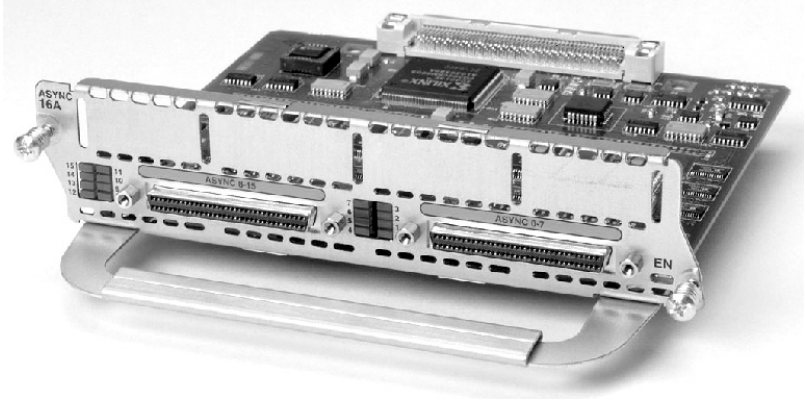

Figure 3-2
NM-16A module for 2600/3600 routers

anything stored on the machine. Even many networked computers were configured without a login requirement because, after all, "everybody was just like family." This would be the ultimate "open" system.

It didn't take long to determine that possibly everyone in the family doesn't need to see the checkbook program or read each other's e-mails, but, to protect those resources, it was necessary to know who was sitting at the keyboard. Some applications experimented with using passwords, and people learned to use password protection on shared resources like folders or printers. Once a password was created, it had to be shared with anyone who needed access. If more than one shared resource existed, one of two problems arose. Either unique passwords were necessary for each resource, requiring some users to keep track of multiple passwords and the resource they were associated with, or the same password was used, which meant any user who had legitimate access to one protected resource now had the password for all protected resources.

If someone left under less-than-favorable circumstances, all passwords known to that person had to be changed, and then the other users had to be told about the changes. These early efforts showed right away that leaving security up to the individual users wasn't reliable and wouldn't scale well as the network group.

Many organizations have this same problem with the user names/and or passwords used to access their network devices, such as routers and switches, by storing the access user name and/or password locally on the device. The following output demonstrates access requiring only a password (line con 0 and line aux 0) and access using the local database (line vty 0 4).

```
!
username xyzdotcom password cisco
!
line con 0
 password cisco
 login
line aux 0
 password cisco
 login
line vty 0 4
 login local
end
```

The preceding techniques were covered in Chapter 2. While both methods work, they have the same shortfalls and security issues as the early networking administrators faced with user data and applications. The example could be made more secure by using techniques covered in Chapter 2, including the following:

- Using a different password for each access point
- Using complex passwords containing eight or more characters, incorporating uppercase and lowercase letters, plus numbers and symbols
- Requiring routine password changes
- Using the **service password-encryption** command to hide the passwords from wandering eyes that could see the configuration

The encryption also reduces the chances of someone being able to capture the passwords if the configuration is included in a Telnet session. Telnet transmits all communications in cleartext.

Centralized Authentication

The lack of scalability and reliability limitations associated with locally stored authentication can only be overcome by using some form(s) of centralized authentication server. Network operating systems (OSs) typically have this feature at their core. In addition to allowing access to the network, this system usually is linked to a mechanism for matching login IDs with permissions to use protected resources. While all authenticated users might have permission to use any of the network printers, only members of the "accounting" group might be granted "permission" to access to the Accounting Department servers. And maybe only those accounting group members who are also part of the "payroll" group might be granted access to the payroll software and data.

If a payroll accountant leaves the company under any circumstances, it's only necessary to delete their user account, which also removes them from the accounting and payroll groups, maintaining security for those assets.

AAA is a technology that can work independently or with the network security system to provide centralized authentication, authorization, and accounting security for network devices and remote user access. Figure 3-3 demonstrates a simple example of AAA and a NAS server providing secure access to a company network. With only authentication features, the remote user could have access to both the server (web server possibly)

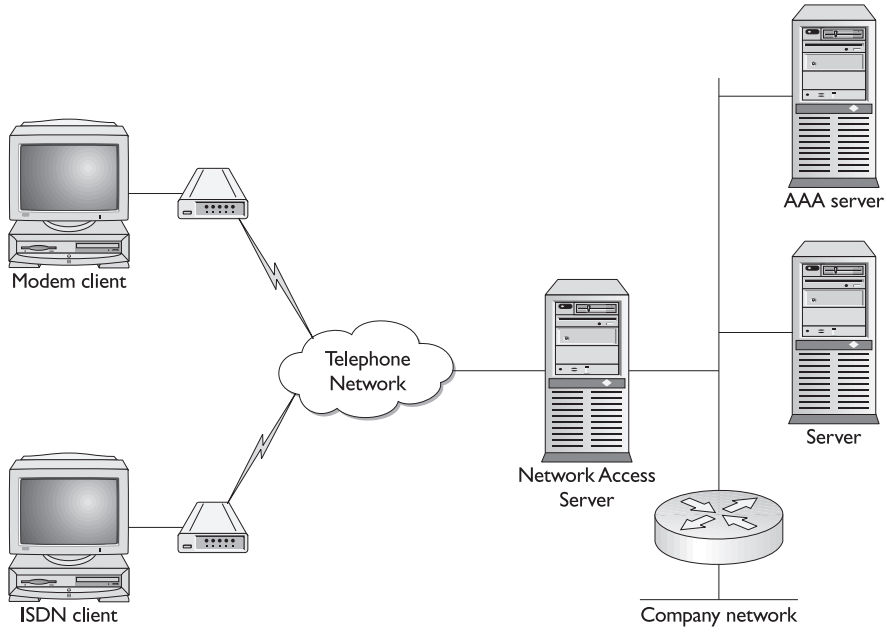

Figure 3-3 AAA and NAS server providing secure remote access to a network

and the company network. With authorization implemented, it would be possible to allow the modem user access to both, while limiting the ISDN user to one or the other.

AAA Benefits

Cisco's AAA technology centralized network access security provides many benefits to the organization and network administrator, including the following:

- Increased flexibility
- Increased security
- Scalability
- Standard authentication methods
- Multiple backup systems

Increased Flexibility

AAA's support of authorization, in addition to authentication, means access can be maintained on a "need to have" basis, without having to maintain multiple passwords. The accounting support means that user auditing and cost-allocation policies can be implemented, as well as providing a trail that might be useful in troubleshooting network problems.

Increased Security

Multiple devices with the *same* locally administered user name/password offer a low level of security. Everyone having access to everything, without regard for need, also unnecessarily increases risk. Multiple locally administered passwords would appear to increase security, but might lead to employees writing down passwords if too many exist to remember. This situation would be exacerbated if complex passwords were implemented.

AAA, with its centralized security database and authorization features, allows a single secure user name/password combination for each employee and yet allows restricting access to a "need to have" basis. At the same time, AAA allows for rapid resolution of compromised passwords or terminated employees.

Scalability

AAA is a template approach to security management that remains reliable and flexible as the network grows larger and more complex. By centralizing the security databases and supporting authorization, AAA avoids the nightmare of managing many user name/password combinations in a growing environment or the alternative "weak" security of using a small number of combinations. Locally stored authentication means any time there's a potential of a compromised user name/password or termination of an employee, each device "should" be reconfigured. The more devices that exist, the greater the amount of effort. AAA avoids this in much the same way that server security is maintained under the same circumstances.

Standard Authentication Methods

AAA supports RADIUS, TACACS+, and Kerberos security protocols for securing dial-in sessions. These protocols provide secure authentication, including encrypted communications and interaction with network server security systems. The next section compares these three systems.

Multiple Backup Systems

AAA supports multiple security servers, such as TACACS+, on the same network to provide redundancy in case of device failure or link congestion. In addition, AAA allows for multiple authentication methods to be specified so, if the first one is unavailable, then a second or third option could be used. For example, if the specified TACACS+ server is offline, the locally stored user name/password database could possibly be used or maybe even the enable password. These alternatives must be defined in advance or access could be blocked until the specified service is restored.

TACACS+, RADIUS, and Kerberos Support

AAA supports all three of these security protocols to control dial-up access into networks. You look, in turn, at each, but note that Cisco supports Kerberos as a legacy security protocol for those networks already committed to it. Cisco Secure Access Control Server (ACS), covered in the next chapter, only implements TACACS+ and RADIUS databases.

At the most obvious level, each of these three protocols does the same thing. Each provides a secure authentication process that allows remote users to access an organization's network resources. At the nuts and bolts level, these are quite different systems, requiring several chapters to detail. The good news is this: that detailed information exists in many places, including Cisco's web site, which is where it's going to stay. This chapter covers those features and differences that might be on the certification exams and would allow a person to choose among them for implementation, or at least to move ahead with intelligent research.

Kerberos is covered first, and then TACACS+ and RADIUS are compared to help determine which should be implemented as part of Cisco Secure ACS.

NOTE It's important to make sure that TACACS+, RADIUS, or Kerberos server services are properly configured before adding the client features to the NAS. Otherwise, you could lock yourself out and require a password recovery.

Kerberos

Kerberos derives its name from the three-headed dog that guards the gates of Hades in Greek mythology. *Kerberos,* the security protocol, is an authentication system developed at the Massachusetts Institute of Technology (MIT), which uses the Data Encryption Standard (DES) cryptographic algorithm for encryption and authentication. Kerberos is based on the assumption that internal users are no more inherently trustworthy than external users and, therefore, applies security and encrypted communications for services like Telnet for all users.

Kerberos is designed to enable two parties to exchange private information across an otherwise open network like the Internet by assigning a unique key, called a *ticket*, to each user that logs on to the network. The ticket is then embedded in messages to identify the sender of the message. Kerberos is based on the concept of a trusted third party, called the *key distribution center (KDC)*, a Kerberos authentication server (AS), performing secure verification of all users and services on the network. This verification involves exchanging encrypted authentication messages without transmitting the user password.

The "lite" explanation of using Kerberos to gain network access might help.

1. A remote user opens a PPP connection to an organization's NAS router.
2. The router prompts the user for a user name and a password.
3. The router then uses only the user name to request a ticket or credential verifying the user identity (TGT) from the key distribution center (KDC).
4. Assuming the user name is known to the KDC, an encrypted TGT is sent back to the router, which includes (among other things) the user's identity. The ticket has a limit life of eight hours to reduce the exposure to an unauthorized capture and replay attempt.
5. The router uses the password from Step 2 to decrypt the TGT. If the decryption is successful, the remote user is authenticated to the router and granted access to the network.

When a remote user successfully authenticates at a boundary router, the user and the user's machine become part of the network. Another TGT from the KDC is necessary to access network services because the original TGT is stored on the router and isn't used for additional authentication, unless the user physically logs on to that router.

Securing a network service involves a double encryption. When a user requests access to a network service, such as a printer or Telnet access to a host, the KDC creates a service credential. This service credential contains the client's identity, the desired network service's identity, and a time limit. The service credential is encrypted first using a password shared by the KDC and the service. The result is encrypted again using the user's TGT as the key.

The target service uses the TGT supplied by the user to perform the first decryption. If the result can be successfully decrypted using the password shared by the KDC and the service, then the service is available.

Since Cisco IOS 11.2, organizations already using Kerberos 5 security can use their existing authentication servers to authenticate their routers and switches. The IOS software Kerberos authentication capabilities support the following network services:

- Telnet
- rlogin
- rsh
- rcp

The previous applications are said to be *Kerberized*, applications and services that have been modified to support the Kerberos credential infrastructure and encrypted communications.

Driver's License Analogy An analogy commonly used for Kerberos is a state driver's license, wherein the state is the KDC that issues a TGT, the license. The license contains information that can authenticate the user: the picture and description items. In addition, the permissions, such as motorcycle endorsement, and restrictions, such as glasses required, associated with the license are also included. Like a TGT, the license also has an expiration time after which it's no longer valid.

Some states include an authentication code made up of portions of the key supplied data, such as name and birth date. Any crude attempt to alter part of the key data makes the authentication code not match.

To complete the analogy, a third party accepts the TGT, license, and, after confirming the picture and description, trusts that the state did a reasonable job in confirming the identity before issuing the license, and thereby accepts it as proof of identity, age, or the right to drive.

For more information, go to http://www.cisco.com on the Web and perform a search for Kerberos. If you'll be working in a Kerberos environment, add a search for Configuring Kerberos to get assistance on using the Kerberos commands. No CCO account is needed for much of the information.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is an access server authentication, authorization, and accounting protocol developed by Livingston Enterprises, Inc., now a division of Lucent Technologies. RADIUS is a system of distributed security that secures remote access to networks and network services against unauthorized access.

RADIUS is a fully open protocol, distributed in source code format that can be modified to work with any security system currently available on the market. Numerous implementations of RADIUS server code are commercially and freely available. Cisco's servers include CiscoSecure ACS for Windows, CiscoSecure UNIX, and Cisco Access Registrar. Several IETF RFCs define the RADIUS protocol, but currently it's still a proposed standard.

A RADIUS implementation is made of the three following components:

- Protocol using UDP/IP communication.
- RADIUS server—a service running on a central Windows or UNIX server, typically at the customer's site.
- RADIUS client—a service residing in the dial-up access servers that can be distributed throughout the network. Cisco added RADIUS client support beginning with IOS v11.1.

Client/Server Model A NAS router running the RADIUS client service passes user information to defined RADIUS server(s), and then acts, based on the response received. The RADIUS server responds to all user connection requests, authenticating the

user, and then forwarding all configuration information needed for the client to provide the service to the user. RADIUS servers can act as proxy clients for other types of authentication servers. The RADIUS server, after reviewing the credentials, replies to the client with one of the following responses:

- **Accept** The user is authenticated.
- **Reject** The user is not authenticated and access is denied. The user will be prompted to reenter the user name/password.
- **Challenge** A request for more information from the user to confirm identity.
- **Change Password** The user must change their password.

Network Security Technology RADIUS clients and servers use a shared secret technology for all authentications transactions. The shared secret is never sent over the network. To eliminate the possibility of someone successfully snooping a user's password on an unsecured network, any transmitted user passwords are sent encrypted between the client and RADIUS server.

Flexiblity RADIUS server supports a variety of methods to authenticate a user, including PPP PAP or CHAP, MS-CHAP, UNIX login, and others.

Interoperability The other side of any open standard is that it doesn't always guarantee interoperability with other vendors' implementations'. RADIUS's implementations between different vendors can vary because the RADIUS standard specifically allows vendors to offer custom features or attributes in what are called *AV pairs.* If a vendor device doesn't recognize the AV pair, it ignores it. To demonstrate the variations, the IETF implementation supports attributes, while Ascend supports 254. Cisco currently supports 58 attributes on Cisco Secure ACS, access servers, Ethernet switches, PIX firewalls, and VPN 3000 concentrators.

TACACS+

Terminal Access Controller Access Control System Plus (TACACS+) is an authentication protocol that allows an NAS to communicate with an authentication server to determine if a user has access to the network.

TACACS+, a new protocol developed by Cisco, replaces two earlier industry standards: TACACS and XTACACS (Extended Terminal Access Controller Access Control System). TACACS+ isn't compatible with the two older protocols. Cisco has submitted TACACS+ protocol specification in a draft RFC to the IETF for development of a standard and for those customers interested in developing their own TACACS+ software.

TACACS+ server services are maintained in a database on a TACACS+ daemon running on a Windows 2000/NT or UNIX host. Cisco's servers supporting TACACS+ include CiscoSecure ACS for Windows, CiscoSecure UNIX, and Cisco Access Registrar. Cisco Access Servers (Cisco Secure ACS) can implement both TACACS+ and RADIUS. The underlying architecture of TACACS+ protocol complements the AAA architecture.

TACACS+ and RADIUS Compared

This section compares several key features of TACACS+ and RADIUS to help understand the strengths and weaknesses of each. While counterpointing many of these comparisons is possible and the information is useful, remember, in the long run, it's Cisco's certification and their comparison.

Authentication and Authorization RADIUS combines the authentication and authorization services. The access-accept packets sent by the RADIUS server in cleartext to the client contain authorization information. Accounting is a separate service on the RADIUS server.

TACACS+ fully supports the AAA architecture by separating the authentication, authorization, and accounting. This allows the flexibility of using another service, such as Kerberos, for authentication, while still using TACACS+ for authorization and/or accounting.

UDP vs. TCP TACACS+ uses TCP for connection-oriented transport between clients and servers. TCP port 49 is reserved for TACACS+. RADIUS uses UDP for best-effort delivery, requiring additional variables to be defined, such as retransmit attempts and time-outs to compensate.

The acknowledgements (TCP ACK) provide indications that a request has been received within (approximately) a network round-trip time (RTT). This same TCP process uses RST (reset) packets to provide immediate indication of a failed (or offline) authentication server. UDP can't tell the difference between a failed server, a slow server, and a nonexistent server.

TCP keepalive packets can be used to watch for failed servers and to facilitate rapid failover between multiple connected authentication servers.

TCP scales better and adapts better to growing and/or congested networks.

Challenge/Response RADIUS supports only unidirectional challenge/response from the RADIUS server to the RADIUS client. TACACS+ supports bidirectional challenge/response like CHAP between the two NASs.

Multiprotocol Support Both TACACS+ and RADIUS support SLIP and PPP encapsulation protocols, but RADIUS doesn't support the following TACACS+ supported protocols:

- Novell Asynchronous Services Interface (NASI)
- X.25 PAD connection
- Net BIOS Frame Protocol Control protocol
- AppleTalk Remote Access protocol (ARAP)

Packet Encryption RADIUS encrypts only the password in the access-request packet from the client to the server, using MD5 hashing for security. The remainder of the packet is in cleartext exposing information such as user name, authorized services, and accounting to be captured by snooping.

TACACS+ encrypts the entire data payload of the packet leaving only the standard TACACS+ header in cleartext. While leaving the body of the packets unencrypted is possible for debugging purposes, normal operation fully encrypts the body for more secure communications. A field in the header indicates whether the body is encrypted.

Router Management RADIUS doesn't support limiting the user access to specific router commands as a tool for router management or terminal services.

TACACS+ supports two methods for controlling the authorization of router commands on either a per-user or a per-group basis.

- Assign commands to privilege levels and have the router use TACACS+ to verify that the user is authorized at the specified privilege level.
- Explicitly define the commands allowed on a per-user or a per-group basis on the TACACS+ server.

AAA System Components

Access control is the way to manage who is allowed access to the NAS and what services they're allowed to use once they have access. Authentication, Authorization, and Accounting (AAA) network security services provide the primary framework through which to set up access control on the router or access server. AAA is comprised of three independent security functions included in the name:

- Authentication
- Authorization
- Accounting

AAA as Facilitator

AAA is designed for flexibility by enabling the administrator to configure the type of authentication and authorization on a per-line (per-user) or per-service basis. During configuration, the types of authentication and authorization to be allowed are defined by creating method lists, and then applying those method lists to specific services or interfaces. The method lists are used to authenticate dial-in users. These lists create an ordered list of security protocols to be used for authentication, thus creating a backup system for authentication in case the initial method fails. For example, the preferred authentication method might be TACACS+, but if the TACACS server isn't available, then use the local user name/password database. Finally, if the user name/password entries have been removed, then use the enable password.

AAA uses protocols, such as RADIUS, TACACS+, and Kerberos, to administer its security functions. If a router or access server is acting as an NAS, then AAA is the means through which the NAS communicates with the RADIUS, TACACS+, or Kerberos security server.

Steps to Configure AAA

Configuring AAA is relatively simple once the basic process is understood. The basic steps to configure AAA security on a Cisco router or access server are the following:

1. Enable AAA by using the **aaa new-model** global configuration command.
2. If you decide to use a separate security server, such as RADIUS, TACACS+, or Kerberos, configure security protocol parameters to use the appropriate server(s).
3. Define the method lists for authentication by using an **aaa authentication** command.
4. Apply the method lists to a particular interface or line, if required.
5. (Optional) Configure authorization using the **aaa authorization** command.
6. (Optional) Configure accounting using the **aaa accounting** command.

You will look at Steps 3 through 6 in the next sections. For now, you will concentrate on starting AAA and telling it how to find the authentication server, assuming one will be used.

NOTE The exam and this book only use the TACACS+ and RADIUS features. For more information on configuring Kerberos, go to http://www.cisco.com/ on the Web and perform a search for configuring Kerberos. No CCO account is needed for much of the information.

Enable AAA

To enable the AAA access control model on the router or NAS, in global configuration mode, type the **aaa new-model** command. Use the no form of the command to disable the AAA. The syntax is

```
Rtr1(config)#aaa new-model
Rtr1(config)#no aaa new-model
```

There are no variables. The words simply turn the feature on and off. Once entered, the CON, VTY, AUX, and TTY lines require a user name and a password for access. Any previous password configuration for those lines is automatically removed.

Define the Security Server—TACACS+

If a TACACS server host is used for authentication, use the **tacacs-server host** global configuration command. Use the no form of this command to delete the specified name or address. The syntax is

```
Rtr1(config)#tacacs-server host {hostname | ip address} [single-connection] [port port#]
    [timeout seconds] [key string]
Rtr1(config)#no tacacs-server host hostname
```

hostname	Name of the TACACS+ server.
ip-addresss	IP address of the TACACS+ server.
single-connection	(Optional) Specify that the router maintain a single open connection for confirmation from the AAA/TACACS+ server. This command contains no autodetect feature and fails if the host isn't running a CiscoSecure daemon.
port	(Optional) Overrides the default, which is port 49.
port#	(Optional) Port number of the server (range 1 to 65535).
timeout	(Optional) Overrides the global timeout value set with the global **tacacs-server timeout** command for this server only.
seconds	(Optional) New timeout interval in seconds.
key	(Optional) Specify an authentication and encryption key. This must match the key used by the TACACS+ daemon. Specifying this key overrides the key set by the global command **tacacs-server key** for this server only.
string	(Optional) Character string to use as the authentication and encryption key.

Multiple **tacacs-server host** commands can be used to specify additional host servers. The Cisco IOS software searches for hosts in the order in which they're specified. Use the single-connection, port, timeout, and key options only when running AAA/TACACS+ server.

Because some of the parameters of the **tacacs-server host** command override global settings made by the **tacacs-server timeout** and **tacacs-server key** commands, this command can be used to enhance security on a network by uniquely configuring individual TACACS+ connections.

The following example specifies a TACACS host named tac-serv1:

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host tac-serv1
```

The next AAA example specifies that the router first try the CiscoSecure TACACS+ host 192.168.1.4. If 192.168.1.4 is unavailable, then use tac-serv1:

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server host tac-serv1
```

The next AAA example specifies that the router consult the CiscoSecure TACACS+ host named 192.168.1.4 on port number 51. The timeout value for requests on this connection is three seconds. The encryption key is a_secret.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4 single-connection port 51
     timeout 3 key a_secret
```

Define TACACS+ Server Key Option After enabling AAA with the **aaa new-model** command, the authentication and encryption key must be set using the **tacacs-**

server key global configuration command. Use the no form of this command to disable the key.

Rtr1(config)#tacacs-server key *key*
Rtr1(config)#no tacacs-server key [*key*]

key	A character entry, this key must match the key used on the TACACS+ daemon. Any leading spaces are ignored, *but spaces within and at the end of the key aren't ignored.* Don't enclose the key in quotation marks unless they're part of the key.

The following example shows an AAA TACACS+ implementation with the authentication and encryption key set to seattle19:

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
```

Define the Security Server—RADIUS

To specify a RADIUS server host, use the **radius-server host global** configuration command. Use the no form of this command to delete the specified RADIUS host. The syntax is

Rtr1(config)#radius-server host {*hostname* | *ip-address*} [auth-port *port-number*]
[acct-port *port-number*] [timeout *seconds*] [retransmit *retries*] [key *string*]
[alias{*hostname* | *ip-address*}]
Rtr1(config)#no radius-server host {*hostname* | *ip-address*} [auth-port *port-number*]
[acct-port *port-number*] [timeout *seconds*] [retransmit *retries*] [key *string*]

hostname	DNS name of the RADIUS server.
ip-address	IP address of the RADIUS server.
auth-port	(Optional) To specify a UDP destination port for authentication requests.
port-number	(Optional) Port number for authentication requests. The host isn't used for authentication if it's set to 0.
acct-port	(Optional) To specify a UDP destination port for accounting requests.
port-number	(Optional) Port number for accounting requests. The host isn't used for accounting if it's set to 0.

Use multiple radius-server host commands to specify multiple hosts. The software searches for hosts in the order in which they're specified.

The following example specifies host1 as the RADIUS server and uses default ports for both accounting and authentication.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host host1.domain.com
```

The next example defines port 12 as the destination port for authentication requests and port 16 as the destination port for accounting requests on a RADIUS host named 192.168.1.4. Because entering a line resets all the port numbers, you must specify a host and configure both the accounting and authentication ports on a single line.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4 auth-port 12 acct-port 16
```

To use separate servers for accounting and authentication, use the zero port value, as appropriate. The following example specifies that RADIUS server192.168.1.4 be used for accounting, but not for authentication, and that RADIUS server host1 be used for authentication, but not for accounting:

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4 auth-port 0
Rtr1(config)#radius-server host host1.domain.com acct-port 0
```

Define RADIUS Server Key Option After enabling AAA with the **aaa new-model** command, the authentication and encryption key must be set using the **radius-server key** global configuration command. Use the no form of this command to disable the key:

Rtr1(config)#radius-server key *key*
Rtr1(config)#no radius-server key [*key*]

key	A character entry, this key must match the key used on the RADIUS daemon. Any leading spaces are ignored, *but spaces within, and at the end of the key, aren't ignored.* Don't enclose the key in quotation marks unless they're part of the key.

The following example shows an AAA RADIUS implementation the authentication and encryption key set to seattle19:

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4
Rtr1(config)#radius-server key seattle19
```

Authentication

Authentication is the way a user is identified prior to being allowed access to the network and network services. The authentication process can include a login and password dialog box, challenge and response, messaging support, and any encryption used.

You configure AAA authentication by defining a named list of authentication methods, and then applying that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they'll be performed. The method list must be applied to a specific interface before any of the defined authentication methods will be performed. The only exception is the default method list (which, by coincidence, is named "default"). The default method list is automatically applied to all interfaces if no other method list is defined. A defined method list overrides the default method list.

Steps to Configure AAA

Configuring AAA is relatively simple once the basic process is understood. The basic steps to configure AAA security on a Cisco router or access server are as follows:

1. Enable AAA by using the **aaa new-model** global configuration command.
2. If you decide to use a separate security server, such as RADIUS, TACACS+, or Kerberos, configure security protocol parameters to use the appropriate server(s).
3. Define the method lists for authentication by using an **aaa authentication** command.
4. Apply the method lists to a particular interface or line, if required.
5. (Optional) Configure authorization using the **aaa authorization** command.
6. (Optional) Configure accounting using the **aaa accounting** command.

You will look at Steps 3 and 4 in this section.

Define the Method Lists

After identifying the authentication server and defining an associated encryption key, now it's time to define method lists for authentication. Use the **aaa authentication** command to name the list and define the authentication method in the order they're to be tried.

The **aaa authentication *purpose*** command has several options, including the following implementations:

- **aaa authentication login** Authentication at login
- **aaa authentication ppp** Authentication methods for interfaces running PPP
- **aaa authentication nasi** Authentication for Netware Asynchronous Services Interface
- **aaa authentication arap** Authentication method for AppleTalk Remote Access protocol

The book and the exams will concentrate on the login and PPP options. The configuration is similar except for the method choices, which can be overcome by using the **?** to display the options while configuring.

Define the Method Lists—Login

To set AAA authentication at login, use the **aaa authentication login** global configuration command. Use the no form of this command to disable AAA authentication. The syntax is

```
Rtr1(config)#aaa authentication login {default | list-name} method1 [method2…]
Rtr1(config)#no aaa authentication login {default | list-name} method1 [method2…]
```

default	Uses the listed authentication methods as the default list to be used when a user logs in
list-name	A character string used to name the list of authentication methods that can be specified to be used when a user logs in
method	At least one of the keywords described in the next table

The *method* argument identifies the list of methods the authentication algorithm tries, in the stated sequence. Method keywords are described in the following table.

Keyword	Description
group tacacs+	Use the list of all TACACS+ servers to authenticate services.
group radius	Use the list of all RADIUS servers to authenticate services.
group *group-name*	Use a subset of RADIUS or TACACS+ servers for authentication, as defined by the server group *group-name*.
local	Use the local user name database for authentication.
local-case	Use the case-sensitive local user name database for authentication.
Enable	Use enable password for authentication.
Line	Use the line password for authentication.
krb5	Use Kerberos 5 for authentication.
krb5-telnet	Use Kerberos 5 Telnet authentication protocol when using Telnet to connect to the router.
none	Use no authentication—no security.

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

Implementing Authentication Method Lists

The authentication method lists are implemented on interfaces with the **login authentication {default | *list-name*}** command. If no list is implemented on an interface with the **login authentication** command, a default list to be used can be specified with the default keyword, followed by the methods.

If authentication isn't specifically set for a line, the default is to deny access and no authentication is performed. Use the **show running-config** command to display currently configured lists of authentication methods.

The following example creates an AAA authentication list called *XYZ-access*. This authentication first tries to contact a TACACS+ server. If no server is found, TACACS+ returns an error and AAA tries to use the enable password. If this attempt also returns an error (because no enable password is configured on the server), the user is allowed access with no authentication.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
```

```
Rtr1(config)#aaa authentication login XYZ-access group tacacs+ enable none
Rtr1(config)#line vty 0 4
Rtr1(config-line)#login authentication XYZ-access
```

The following example creates the same list, but sets it as the default list that's used for all login authentications if no other list is specified. It adds the local user name/password database as the second choice.

```
Rtr1(config)#username last password hope
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
Rtr1(config)#aaa authentication login default group tacacs+ local enable none
```


Define the Authentication Method Lists—PPP

To set AAA authentication at login, use the **aaa authentication ppp** global configuration command. Use the no form of this command to disable AAA authentication. The syntax is

Rtr1(config)#aaa authentication ppp {default | *list-name*} *method1* [*method2*…]
Rtr1(config)#no aaa authentication ppp {default | *list-name*} *method1* [*method2*…]

default	Uses the listed authentication methods as the default list to be used when a user logs in
list-name	Character string used to name the list of authentication methods that can be specified to be used when a user logs in
method	At least one of the keywords described in the next table

The *method* argument identifies the list of methods the authentication algorithm tries, in the stated sequence. Method keywords are described in the following table.

Keyword	Description
group tacacs+	Use the list of all TACACS+ servers to authenticate services.
group radius	Use the list of all RADIUS servers to authenticate services.
group *group-name*	Use a subset of RADIUS or TACACS+ servers for authentication, as defined by the server group *group-name*.
local	Use the local user name database for authentication.
local-case	Use the case-sensitive local user name database for authentication.
Enable	Use enable password for authentication.
Line	Use the line password for authentication.
krb5	Use Kerberos 5 for authentication.
if-needed	Does not authenticate if the user has already been authenticated on a TTY line.
none	Use no authentication—no security.

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

Implementing Authentication Method Lists

The PPP authentication method lists are implemented on interfaces with the **ppp authentication {default | *list-name*}** command. *These lists contain up to four authentication methods that are used when a user tries to log in to the serial interface.* If no list is implemented on an interface with the **ppp authentication** command, a default list to be used can be specified with the default keyword followed by the methods.

If authentication isn't specifically set for a line, the default is to deny access and no authentication is performed. Use the **show running-config** command to display currently configured lists of authentication methods.

The following example uses the previous named list example and adds a PPP authentication list named PPP-access. This authentication first tries to contact a TACACS+ server. If no TACACS+ server is found, the user is allowed access with no authentication.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
Rtr1(config)#aaa authentication login XYZ-access group tacacs+ enable none
Rtr1(config)#aaa authentication ppp PPP-access group tacacs+ none
Rtr1(config)#interface serial 0/1
Rtr1(config-if)#encapsulation ppp
Rtr1(config-if)#ppp authentication ppp-access
```

The following example creates the same list, but sets it as the default list used for all login authentications if no other list is specified, and adds the if-needed method.

```
Rtr1(config)#username last password hope
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
Rtr1(config)#aaa authentication login default group tacacs+ local enable none
Rtr1(config)#aaa authentication ppp default if-needed group tacacs+ none
```

Authorization

AAA authorization works by assembling a set of attributes that describe what the user is authorized to perform or access. These attributes are compared to the information contained in a database for a given user, and the result is returned to AAA to determine the user's actual capabilities and restrictions. The database can be located locally on the access server or the router, or it can be stored remotely on a RADIUS or TACACS+ security server. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights by associating attribute-value (AV) pairs, which define those rights, with the appropriate user. All authorization methods must be defined through AAA.

As with authentication, you configure AAA authorization by defining a named list of authorization methods, and then applying that list to various interfaces.

Steps to Configure AAA

Configuring AAA is relatively simple once the basic process is understood. The basic steps to configure AAA security on a Cisco router or access server are

1. Enable AAA by using the **aaa new-model** global configuration command.
2. If you decide to use a separate security server, such as RADIUS, TACACS+, or Kerberos, configure security protocol parameters to use the appropriate server(s).
3. Define the method lists for authentication by using an **aaa authentication** command.
4. Apply the method lists to a particular interface or line, if required.
5. (Optional) Configure authorization using the **aaa authorization** command.
6. (Optional) Configure accounting using the **aaa accounting** command.

You will look at Step 5 in this section.

Define the Authorization Method Lists

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

A *method list,* as in authentication, is simply a named list describing the authorization methods to be queried (such as RADIUS or TACACS+), in sequence. Method lists enable one or more security protocols for authorization to be designated, thus ensuring a backup system in case the initial method fails. Cisco IOS software uses the first method listed to authorize users for specific network services. If that method fails to respond, the Cisco IOS software selects the next method listed in the method list. This process continues until successful communication occurs with a listed authorization method or all methods defined are exhausted.

NOTE The Cisco IOS software attempts authorization with the next listed method only when no response occurs from the previous method. If authorization fails at any point in this cycle (meaning the security server or local user name database responds by denying the user services), the authorization process stops and no other authorization methods are attempted.

Use the **aaa authorization** global configuration command to set parameters that restrict a user's network access. Use the no form of this command to disable authorization for a function. The basic syntax is

Rtr1(config)#aaa authorization {*authorization-type*} {default | *list-name*} *method1* [*method2*…]Rtr1(config)#no aaa authorization {*authorization-type*}

The first step is to choose which of the eight authorization types AAA supports is to be validated. The actual syntax and choices include

Rtr1(config)#aaa authorization {**network** | **exec** | **commands** ***level***| **reverse-access** | **configuration** | **config-commands** | **auth-proxy** | **ipmobile**}
{default | *list-name*} *method1* [*method2*…]

network	All network-related service requests, including SLIP, PPP, PPP NCPs, and ARAP.
exec	Is the user allowed to run an EXEC shell?
commands *level*	Specific command level to be authorized (0 through 15).
reverse-access	Reverse access connections, such as reverse Telnet.
configuration	Downloads the configuration from the AAA server.
config-commands	Configuration mode commands.
auth-proxy	Authentication Proxy Services.
Ipmobile	Mobile IP services.

Once the authorization type is selected, the rest is just like the authentication process.

default	Uses the listed authorization methods that follow this argument as the default list of methods for authorization
list-name	Character string used to name the list of authorization methods
method1 [*method2*…]	One of the keywords listed in the following table

The actual method lists are specific to the type of authorization being requested. The six methods Cisco IOS software supports for authorization are described in the following table.

Method	Description
group tacacs+	Uses the list of all TACACS+ servers to provide authorization services. TACACS+ authorization defines specific rights for users by associating attribute-value (AV) pairs, which are stored in a database on the TACACS+ security server, with the appropriate user.
group radius	Uses the list of all RADIUS servers to provide authorization service. RADIUS authorization defines specific rights for users by associating attributes, which are stored in a database on the RADIUS server.
if-authenticated	Allows the user to access the requested function if the user is authenticated.
local	Uses the local database for authorization, as defined by the user name command, to authorize specific rights for users. Only a limited set of functions can be controlled by the local database.
krb5-instance	Uses the instance defined by the **Kerberos Instance Map** command.
none	The NAS doesn't request authorization information. Authorization isn't performed over this line/interface.

When creating a named method list, a particular list of authorization methods for the indicated authorization type is defined. Once defined, method lists must be applied to specific lines or interfaces, as with authentication, before any of the defined methods will be performed. The **authorization** command causes a request packet, containing a series of AV pairs, to be sent to the RADIUS or TACACS+ daemon as part of the authorization process. The daemon can do one of the following:

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

Implementing Authorization Method Lists

To configure AAA authorization using named method lists, perform the following tasks beginning in Global Configuration mode:

Task	Command
Create an authorization method list for an authorization type and enable authorization.	aaa authorization {*authorization-type*} *list-name* [*method1* [*method2*…]]
Enter Line Configuration mode on which you want to apply the authorization method list. or Enter the Interface Configuration mode for the interfaces to which you want to apply the authorization method list.	line [aux | console | tty | vty] *line-number* [*ending-line-number*] or interface *interface-type interface-number*
Apply the authorization list to the line(s). or Apply the authorization list to the interface.	authorization {*authorization-type*} *list-name* or ppp authorization *list-name*

The following example defines the network authorization method list named XYZ-auth, which specifies TACACS+ authorization will be used on serial lines using PPP. If the TACACS+ server fails to respond, the local network authorization will be performed.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
Rtr1(config)#aaa authentication login XYZ-access group tacacs+ enable none
Rtr1(config)#aaa authentication ppp PPP-access group tacacs+ none
Rtr1(config)#aaa authorization network xyz-auth group tacacs+ local
Rtr1(config)#interface serial 0/1
Rtr1(config-if)#encapsulation ppp
Rtr1(config-if)#ppp authentication ppp-access
Rtr1(config-if)#ppp authorization xyz-auth
```

Accounting

Accounting enables you to track the services users are accessing, as well as the amount of network resources they're consuming. When AAA accounting is activated, the NAS

reports user activity to the TACACS+ or RADIUS security server (depending on which security method you have implemented) in the form of accounting records. Each accounting record is comprised of accounting AV pairs and is stored on the access control server. This data can then be analyzed for network management, client billing, and/or auditing. All accounting methods must be defined through AAA. As with authentication and authorization, you can configure AAA accounting by defining a named list of accounting methods, and then applying that list to various interfaces.

Steps to Configure AAA

Configuring AAA is relatively simple once the basic process is understood. The basic steps to configure AAA security on a Cisco router or access server are the following:

1. Enable AAA by using the **aaa new-model** global configuration command.
2. If you decide to use a separate security server, such as RADIUS, TACACS+, or Kerberos, configure security protocol parameters to use the appropriate server(s).
3. Define the method lists for authentication by using an **aaa authentication** command.
4. Apply the method lists to a particular interface or line, if required.
5. (Optional) Configure authorization using the **aaa authorization** command.
6. (Optional) Configure accounting using the **aaa accounting** command.

You will look at Step 6 in this section.

Define the Accounting Method Lists

When **aaa accounting** is activated, the NAS monitors either RADIUS accounting attributes or TACACS+ AV pairs pertinent to the connection, depending on the security method you implemented. The NAS reports these attributes as accounting records, which are then stored in an accounting log on the security server. The **aaa accounting** command enables accounting and enables you to create named method lists defining specific accounting methods on a per-line or per-interface basis. Use the no form of this command to disable accounting. The basic syntax is

```
Rtr1(config)#aaa accounting {accounting-type} {default | list-name} {accounting-method}
    method1 [method2…]

Rtr1(config)#no aaa accounting {accounting-type}
```

The first step is to choose which of the eight authorization types that AAA supports should be validated. The syntax and choices include

Rtr1(config)# aaa accounting {**system** | **network** | **exec** | **connection** | **commands** ***level*** | **nested** | **update**} {default | *list-name*} {***accounting-method***} *method1* [*method2…*]

system	All system-level events not associated with users, such as reloads. Don't use named accounting lists, only use the default list for system accounting.
network	All network-related service requests, including SLIP, PPP, PPP NCPs, and ARAP.
exec	Creates accounting records about user EXEC terminal sessions on the NAS, including user name, date, start, and stop times.
connection	All outbound connections from the NAS, such as Telnet, local-area transport (LAT), TN3270, packet assembler/disassembler (PAD), and rlogin.
commands *level*	Specific command level to track for accounting (0 through 15).
nested	Provides accounting when starting PPP from EXEC, generate NETWORK records before EXEC-STOP record.
update	Enables periodic interim accounting records to be sent to the accounting server.

Once the accounting type is selected, you must specify the accounting method to be used in recording the results. The syntax and choices include

Rtr1(config)#aaa accounting {*accounting-type*} {default | *list-name*} {**start-stop** | **wait-start** | **stop-only** | **none**} *method1* [*method2…*]

start-stop	Sends a *start accounting notice* at the beginning of a process and a **stop** *accounting notice* at the end of a process. The *start accounting record* is sent in the background. The requested user process begins, regardless of whether the start accounting notice was received by the accounting server.
wait-start	Sends both a start and a stop accounting notice to the accounting server. But, if the *wait-start* keyword is used, the requested user service does *not* begin until the start accounting notice is acknowledged.
stop-only	Sends a *stop accounting notice* at the end of the requested user process.
None	Disables accounting services on this line or interface.

For minimal accounting, use the stop-only keyword. For more detailed accounting, include the start-stop keyword. For even more accounting control, include the wait-start keyword, which ensures the start notice is received by the RADIUS or TACACS+ server before granting the user's process request.

When **aaa accounting** is activated, the network access server monitors either RADIUS accounting attributes or TACACS+ AV pairs pertinent to the connection, depending on the security method implemented. The network access server reports these attributes as accounting records, which are then stored in an accounting log on the security server. The default | *list-name* and *method* options are just like in the authentication and authorization process. Individual named accounting method lists are specific to the indicated

accounting type. *System accounting* doesn't use named accounting lists. Only the default list for system accounting can be defined.

default	Uses the listed accounting methods that follow this argument as the default list of methods for accounting
list-name	Character string used to name the list of accounting methods
method1 [*method2*…]	One of the keywords listed in the following table

Accounting method keywords are described in the following table.

Keyword	Description
group tacacs+	Uses the list of all TACACS+ servers to provide authorization services. The NAS reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server.
group radius	Uses the list of all RADIUS servers to provide authorization service. The NAS reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server.
group | *group-name*	Uses a subset of RADIUS or TACACS+ servers for accounting, as defined by the server group *group-name*.

Method lists for accounting define the way accounting will be performed. Named accounting method lists enable the option to designate a particular security protocol to be used on specific lines or interfaces for particular types of accounting services.

Implementing Accounting Method Lists

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

To configure AAA accounting using named method lists, perform the following tasks beginning in Global Configuration mode:

Task	Command
Create an accounting method list and enable accounting.	aaa accounting {*accounting-type*} *list-name* {*accounting-method*} *method1* [*method2*…]
Enter Line Configuration mode on which you want to apply the accounting method list. or Enter the Interface Configuration mode for the interfaces to which you want to apply the accounting method list.	line [aux | console | tty | vty] *line-number* [*ending-line-number*] or interface *interface-type interface-number*

Task	Command
Apply the accounting list to the line(s). or Apply the accounting list to the interface.	aaa accounting {*accounting-type*} *list-name* or ppp accounting *list-name*

The following example defines the network accounting method list named XYZ-acct, where command accounting services are provided by a TACACS+ security server with a stop-only restriction.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
Rtr1(config)#aaa authentication login XYZ-access group tacacs+ enable none
Rtr1(config)#aaa authentication ppp PPP-access group tacacs+ none
Rtr1(config)#aaa authorization network xyz-auth group tacacs+ local
Rtr1(config)#aaa accounting network xyz-acct group start-stop tacacs+
Rtr1(config)#interface serial 0/1
Rtr1(config-if)#encapsulation ppp
Rtr1(config-if)#ppp authentication ppp-access
Rtr1(config-if)#ppp authorization xyz-auth
Rtr1(config-if)#ppp accounting xyz-acct
```

Testing AAA Configuration

Troubleshooting AAA can be rather comple—because it's used so often with other features, such as PPP—so remembering to use the troubleshooting commands associated with any protocols or technologies working with AAA is important. This section looks at some common commands for confirming AAA configuration and activity.

When working with AAA in a specific environment, such as dial-up modems, ISDN, or PPP, don't overlook Cisco's web site for more information. Go to http://www.cisco.com on the Web and perform a search for AAA or ISDN AAA. No CCO account is needed for much of the information.

The show Commands

Two **show** commands useful in debugging AAA are

show running-config	To verify that local AAA is configured correctly
show tacacs	To verify network connectivity between NAS and AAA server

The debug Commands

Cisco **IOS debug** command output provides a valuable source of information and feedback concerning state transitions and functions within the AAA environment. In addition to **debug** command output gathered directly from devices running Cisco IOS, the

Cisco AAA server can be configured to collect operational diagnostics. Use the following debug commands to capture AAA-related transitions and functions:

debug condition user *username*	Sets conditional debugging for a specific user and generates output debugs related to the user
debug aaa authentication	Displays authentication information with TACACS+ and RADIUS client/server interaction
debug aaa authorization	Displays authorization information with TACACS+ and RADIUS client/server interaction
debug aaa accounting	Displays accounting information with TACACS+ and RADIUS client/server interaction
debug tacacs	Displays TACACS+ interaction between the IOS client and the AAA server
debug radius	Displays RADIUS interaction between the IOS client and the AAA server
debug ppp negotiation	Shows if a client is passing PPP negotiation
debug ppp authentication	Shows if a client is passing PPP authentication
debug ppp error	Displays protocol errors and error statistics associated with PPP connection negotiation and operation

Chapter Review

Remote dial-up connections to the corporate network are made up of several dial-in technologies, including modem and ISDN connections, as well as virtual connections via the Internet. Access control is the process of controlling who can access the network and what resources they're allowed to use. Cisco's Authentication, Authorization, and Accounting (AAA) network security services configured on a router or network server implement this access control.

The three security components of AAA are designed to let you define and configure the type of authentication, authorization, and accounting in a detailed and consistent manner through the use of method lists, and then apply those method lists to specific services or interfaces according to your security plan. Method lists define a sequence of implementation processes that allow backup methods in case the initial method fails.

Authentication is the process of identifying users through user name and password verification methods that allow only approved individuals to access the network.

Authorization is the process of matching authenticated users with the permissions or privileges to use network resources

Accounting is the process of tracking or logging the different types of resources or services the remote users are accessing. This data can then be analyzed for auditing, troubleshooting, network management, and network planning client billing. The AAA databases can be remotely stored on one or more TACACS+ or RADIUS servers. Authentication and authorization databases can be stored locally on the access server.

AAA benefits include

- Scalability
- Flexibility and granularity
- Multiple implementation methods, which provide redundancy
- Support of standard authentication methods, such as RADIUS, TACACS+, and Kerberos

Questions

1. Which of the following is *not* one of the three components of AAA?
   - **A.** Accounting
   - **B.** Acknowledgement
   - **C.** Authorization
   - **D.** Authentication
2. Which one of the following is the process of determining what devices, features, or services a specific remote user has permission to access in the network, such as network resources or services?
   - **A.** Accounting
   - **B.** Acknowledgement
   - **C.** Authorization
   - **D.** Authentication
3. Which of the following is a term for the router with interfaces designed to service the remote users of the company?
   - **A.** Remote server
   - **B.** NAS
   - **C.** Access point
   - **D.** Authentication server
4. Which one of the following is *not* one of the three security protocols to control dial-up access into networks supported by AAA?
   - **A.** TACACS+
   - **B.** Kerberos
   - **C.** RADIUS
   - **D.** ASICS

5. Which of the following security protocols is considered legacy and is supported for those organizations already implementing it?
   - **A.** TACACS+
   - **B.** Kerberos
   - **C.** RADIUS
   - **D.** ASICS
6. Which of the following is a security protocol developed by Livingston Enterprises, Inc., now a division of Lucent Technologies?
   - **A.** TACACS+
   - **B.** Kerberos
   - **C.** RADIUS
   - **D.** ASICS
7. Which of the following is a security protocol developed by MIT?
   - **A.** TACACS+
   - **B.** Kerberos
   - **C.** RADIUS
   - **D.** ASICS
8. Which of the following is a security protocol developed by Cisco and submitted to IETF as a proposed standard?
   - **A.** TACACS+
   - **B.** Kerberos
   - **C.** RADIUS
   - **D.** ASICS
9. Which of the following is *not* an advantage of TACACS+ over RADIUS?
   - **A.** Uses TCP for connections
   - **B.** Supports multiple protocols, including ARAP, NASI, and X.25 PAD
   - **C.** Fully supports the AAA architecture by separating the components
   - **D.** Supports server-based security databases
10. Which command enables the AAA access control model on the router?
    - **A.** tacacs-server host
    - **B.** radius-server host
    - **C.** aaa new-model
    - **D.** tacacs-server key *key*

11. Which command identifies the TACACS+ server host to be used for authentication?
    - **A.** Rtr1(config)#tacacs-server key seattle19
    - **B.** Rtr1(config-if)#tacacs-server key seattle19
    - **C.** Rtr1(config)#tacacs-server host Seattle
    - **D.** Rtr1(config-if)#tacacs-server host Seattle
12. In the following command, what is the first authentication method?
    aaa authentication login XYZ-access group tacacs+ enable none
    - **A.** Group servers
    - **B.** TACACS+
    - **C.** enable password
    - **D.** None
13. Which of the following combines the authentication and authorization into a single database?
    - **A.** TACACS+
    - **B.** Kerberos
    - **C.** RADIUS
    - **D.** None of the above
14. Which command verifies network connectivity between the NAS and the AAA server?
    - **A.** show running-config
    - **B.** show tacacs
    - **C.** debug tacacs
    - **D.** debug aaa authentication
15. Which of the following is *not* an AAA benefit?
    - **A.** Scalability.
    - **B.** Automatic installation and configuration.
    - **C.** Flexibility and granularity.
    - **D.** Multiple implementation methods provide redundancy.

Answers

1. **B.** Acknowledgement
2. **C.** Authorization
3. **B.** NAS

4. **D.** ASICS
5. **B.** Kerberos
6. **C.** RADIUS
7. **B.** Kerberos
8. **A.** TACACS+
9. **D.** Supports server-based security databases (They both support this feature.)
10. **C.** **aaa new-model**
11. **C.** **Rtr1(config)#tacacs-server host Seattle**
12. **B.** TACACS+
13. **C.** RADIUS
14. **B.** **show tacacs**
15. **B.** Automatic installation and configuration.

CHAPTER 4

Cisco Secure ACS and TACACS+/RADIUS Technologies

In this chapter, you will learn how to:

- Describe Cisco Secure ACS
- Use the features and architecture of CiscoSecure ACS for Windows
- Understand the features of CiscoSecure ACS for UNIX
- Install CiscoSecure ACS 3.0 for Windows NT or Windows 2000
- Administer and troubleshoot CiscoSecure ACS for Windows
- Make use of TACACS+ overview and configuration
- Configure Cisco Secure ACS and TACACS+
- Verify TACACS+

This chapter looks at the features and architecture of the Cisco Secure Access Control Server (ACS) application. The chapter also covers the general process of preparing to install and configure the current versions of ACS support for Windows 2000/NT and UNIX (Solaris). Before starting installation, it's important for you to have the correct instructions for Cisco Secure ACS, for the appropriate network OS and OS version.

For the most recent information, go to http://www.cisco.com on the Web and choose the link for Products | Security Components under the Product Categories heading. Access servers are listed under product categories. Or, perform a search on Cisco Secure ACS. This search should also reveal any recent security announcements and mitigation techniques. No CCO account is needed for much of the information. Be aware that some Windows references refer to NT or ACS NT, but also include Windows 2000.

NOTE Special Cisco Secure ACS upgrade and reinstallation requirements and steps must be followed to make sure everything goes smoothly. The process varies, depending on whether the original configuration settings are to be incorporated in the new installation. Always get the latest upgrade instructions.

Describe Cisco Secure ACS

The *Cisco Secure ACS* is a server-based application offering a centralized user-access control strategy for the various access gateways to your network, as shown in Figure 4-1. Cisco Secure ACS supports the user Authentication, Authorization, and Accounting (AAA) features covered in the preceding chapter and is covered in several chapters pertaining to PIX firewalls.

ACS allows the administrator to manage user access for Cisco IOS routers, firewalls, Cisco Catalyst switches, and virtual private networks (VPNs), as well as newer technologies, such as dial-up and broadband DSL, cable access solutions, voice over IP (VoIP), Cisco wireless implementations, and third-party vendor Terminal Access Controller Access Control System Plus (TACACS+) offerings.

ACS is a strategic product at Cisco, incorporating new, advanced features as they become available to increase the diversity and granularity of implementations. Current support includes user and administrative access reporting, database synchronization and user importation tools, Lightweight Directory Access Protocol (LDAP) user authentication support, dynamic quota generation, restrictions such as time of day and day of week, and user and device group profiles.

CiscoSecure ACS for Windows and UNIX

There's no exam or practical purpose for this chapter to include detailed steps for installing, reinstalling, and upgrading Cisco Secure ACS. The network OS, ACS version, and new install versus upgrade/reinstall variables make using the most recent instructions downloaded from Cisco's site the only prudent solution.

Experience with installing server applications, careful reading of the correct installation instructions, and cautious progress can save you the hassle of having to start all over.

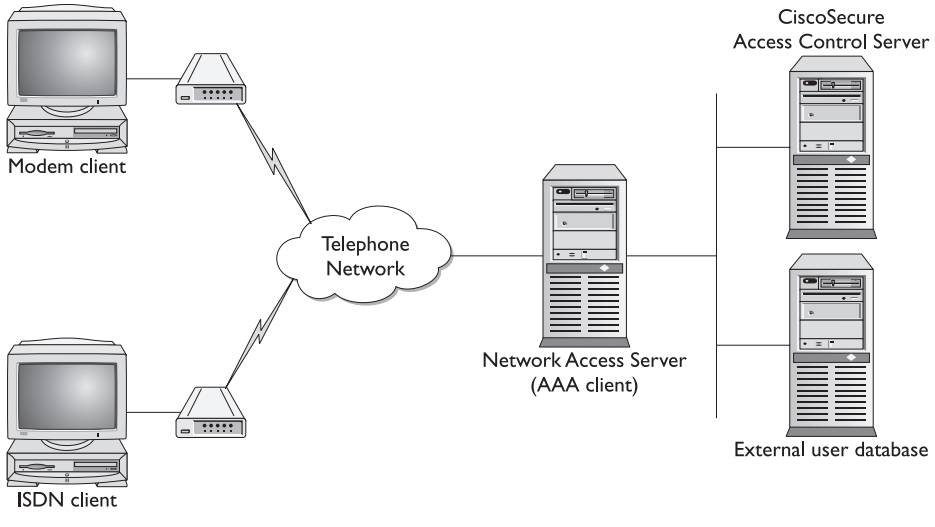

Figure 4-1 Cisco Secure ACS with an NAS AAA client

This section concentrates on system requirements and preparation that would be beneficial when you choose to install Cisco Secure ACS.

Features and Architecture of Cisco Secure ACS for Windows

The latest Windows version is Cisco Secure ACS version 3.0 for Windows 2000 and NT, supporting both RADIUS and TACACS+ server systems. New features included in Cisco Secure ACS version 3.0 allow network administrators to scale and deploy secure network services with centralized control, access management, and accounting within the Cisco Secure ACS framework. Using Cisco Secure ACS, network administrators can control

- Which users can access the network from either wired or wireless connections
- What privileges each user will have while in the network
- What accounting information is kept for capacity planning, account billing, or security audits
- What access and command controls are enabled for each configuration administrator

Features and Benefits

Specific features included in Cisco Secure ACS version 3.0 are the following.

IEEE 802.1x—Access Control for Switched LAN Users

The IEEE 802.1*x* standard brings new security services to the local area network (LAN) by letting network administrators control which users can access their switched LAN environment. The 802.1*x* standard manages port-level access control by using the Cisco Secure ACS Extensible Authentication Protocol (EAP), which is carried on RADIUS.

EAP Message Digest 5 (EAP-MD5) and EAP Transport LAN Services (EAP-TLS)

EAP is an IETF RFC standard for various authentication methods carried over any Point-to-Point Protocol (PPP) connection. *EAP-MD5* is a user name/password method incorporating MD5 hashing for security. *EAP-TLS* uses X.509 digital certificate technology to provide both Cisco Secure ACS server and client authentication. EAP-TLS also supports dynamic session key negotiation.

Microsoft Challenge Authentication Protocol (MSCHAP) Support

Cisco Secure ACS now supports MSCHAP version 2.0 exchanges and MSCHAP password change service for Microsoft Dial-Up Networking clients, Cisco VPN clients, and any other desktop client supporting MSCHAP version 2.0 password change service. A user whose password has expired now is automatically prompted to change their password after their next login.

Multiple LDAP Support

Cisco Secure ACS supports user authentication using records kept in a directory server through the LDAP, including Novell and Netscape, using a generic LDAP interface. New in Cisco Secure ACS version 3.0 is the capability to define multiple, different LDAP sources for user lookups. This lets you define a different LDAP repository to search for users. Another new feature is the capability to define secondary, backup LDAP servers, for times when a primary LDAP server is unavailable.

Device Command Sets (DCS)

Device Command Sets (DCS) is a new TACACS+ administration tool that uses a central Cisco Secure ACS graphical user interface (GUI) mechanism to control the authorization of each command on each device via per-user, per-group, or per-network device group mapping. DCS provides a method to group and name command profiles, which can then be paired with users, groups of users, or device groups. The new features and tools provide greater granularity, scalability, and manageability in setting authorization restrictions for network administrators.

Per-User Access Control Lists (ACL)

Per-user Access Control Lists (ACL) is a Cisco PIX Firewall Solution service that allows administrators to define access control lists (ACL) of any length, for users or groups of users using the Cisco Secure ACS GUI.

New NAS Wildcard, Multi-NAS, and Named Access Filters Features

NAS wildcard support for device, device group entry, and NAS filtering enables easier device entry and management in the Cisco Secure ACS system.

Multi-NAS allows the administrator to create shared NAS profile templates that define a group of network devices with the same attributes: shared key, authentication method, or login/accounting parameters. Multi-NAS also enables administrators to provide multiple IP addresses or ranges of IP addresses.

Named access filters simplify and facilitate assigning the same access filter to multiple devices or device groups.

User-Extensible Vendor-Specific Attributes (VSAs)

The Cisco Secure ACS now supports user-defined outbound Vendor-Specific Attributes (VSAs) using the web GUI, including support for Broadband Service Manager (BBSM) implementations.

Cisco Secure ACS Benefits

Cisco Secure ACS is a powerful access control system with many high-performance, flexibility, and scalability features for the growing WAN or LAN. Some of the benefits of the Windows 3.0 version include the following.

Ease of use	A web-based user interface simplifies and distributes configuration for user profiles, group profiles, and ACS configuration.
Product flexibility	AAA support is integrated in the Cisco IOS Software, allowing Cisco Secure ACS to be implemented across virtually any Cisco NAS.
Scalability	Built to support large networks with support for redundant servers, remote databases, and backup user database services.
Extensibility	LDAP authentication forwarding supports user authentication using profiles stored in directories or databases from leading vendors, such as Microsoft, Netscape, and Novell.
Administration	The capability to assign different access levels to each Cisco Secure administrator, plus the option to group network devices, facilitates easier control, flexibility, and granularity in defining, changing, and enforcing security policy administration.
Management	Shares Windows user name/password management by using Windows 2000 Active Directory and NT database support, as well as using the Windows Performance Monitor for real-time statistics viewing.
Protocol flexibility	Simultaneous TACACS+ and RADIUS support allows flexible implementation of VPN or dial support at both ends of Internet Protocol Security (IPSec) and Point-to-Point Tunneling Protocol (PPTP) tunnels.
Token server support	Token server support for RSA SecurID, Passgo, Secure Computing, ActiveCard, Vasco, and CryptoCard.
Control	Supports dynamic quotas for time-of-day, network usage, number of logged sessions, and day-of-week access restrictions.

Cisco Secure ACS for Windows Internal Architecture

Cisco Secure ACS for Windows NT/2000 version 3.0 servers is designed to be modular and flexible to scale from simple to complex networks. Cisco Secure ACS includes the following service modules:

CSAdmin

CSAdmin is the service for the Cisco Secure ACS internal web server that eliminates the need for a third-party web server. Once installed, Cisco Secure ACS must be configured from its HTML interface, which requires that CSAdmin be running. CSAdmin is a multithreaded application allowing multiple administrators to access it at the same time. CSAdmin is best for distributed, multiprocessor, and clustered environments.

While starting and stopping the other services from within the Cisco Secure ACS HTML interface is possible, this doesn't include starting or stopping CSAdmin. If CSAdmin stops abnormally through an external action, Cisco Secure ACS is only accessible from the Windows NT/2000 server on which it's running. CSAdmin can be started or stopped from the Windows NT/2000 Service menu.

CSAuth

CSAuth is the authentication and authorization service used to permit or deny access to users. CSAuth is the database manager that determines whether access should be granted and defines the privileges for a particular user. Cisco Secure ACS can access several different databases for authentication purposes. When a request for authentication arrives, Cisco Secure ACS checks the database configured for that user. If the user is unknown, Cisco Secure ACS checks the database(s) configured for unknown users. The database options include the following:

- **Cisco Secure ACS user database** The fastest option involves locating the user name and checking the password against the internal Cisco Secure ACS user database, as depicted in Figure 4-2. This avoids any delay while Cisco Secure ACS waits for a response from an external user database.

- **Windows NT/2000 user database** CSAuth passes the user name and password to Windows NT/2000 for authentication using its user database. Windows NT/2000 then provides a response approving or denying validation. Figure 4-3 represents Cisco Secure ACS using the network OS security database to authenticate users.

- **Novell NDS option** Uses the Novell NDS service to authenticate users. Cisco Secure ACS supports one tree, but the tree can have multiple Containers and Contexts. The Novell requester must be installed on the same Windows server as Cisco Secure ACS.

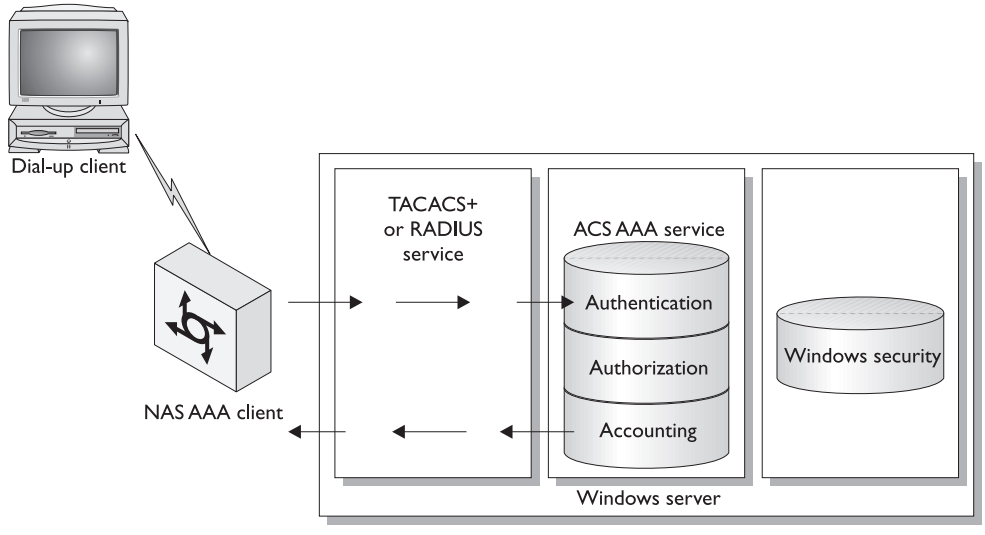

Figure 4-2 Cisco Secure ACS using its own database to authenticate users

- **ODBC** Open Database Connectivity (ODBC)–compliant SQL databases use the ODBC standardized API developed by Microsoft and are now used by most major database vendors. A benefit of ODBC in a web-based environment is easy access to data storage programs, such as Microsoft Access and SQL Server.
- **UNIX passwords** Cisco Secure ACS includes a password import utility to import passwords from a UNIX database.
- **Generic LDAP** Cisco Secure ACS supports authentication of users against records kept in a directory server through the LDAP. Both PAP and CHAP passwords can be used when authenticating against the LDAP database.
- **Token Card servers** Cisco Secure ACS supports token servers, such as RSA SecurID, and SafeWord AXENT, and any hexadecimal X.909 Token Card, such as CRYPTOCard. Cisco Secure ACS either acts as a client to the token server or, in other cases, uses the token server's RADIUS interface for authentication requests. Figure 4-4 shows the Token Card server interacting with Cisco Secure ACS.

When the user authenticates using one of the defined methods, Cisco Secure ACS obtains a set of authorizations from the user profile and any groups the user belongs to. This information is stored with the user name in the Cisco Secure ACS user database. Some authorizations are the services the user is entitled to, such as IP over PPP, IP pools from which to draw an IP address, access lists, and password aging information. The authorizations, with the authentication approval, are then passed to the CSTacacs or CSRadius modules to be sent to the requesting device.

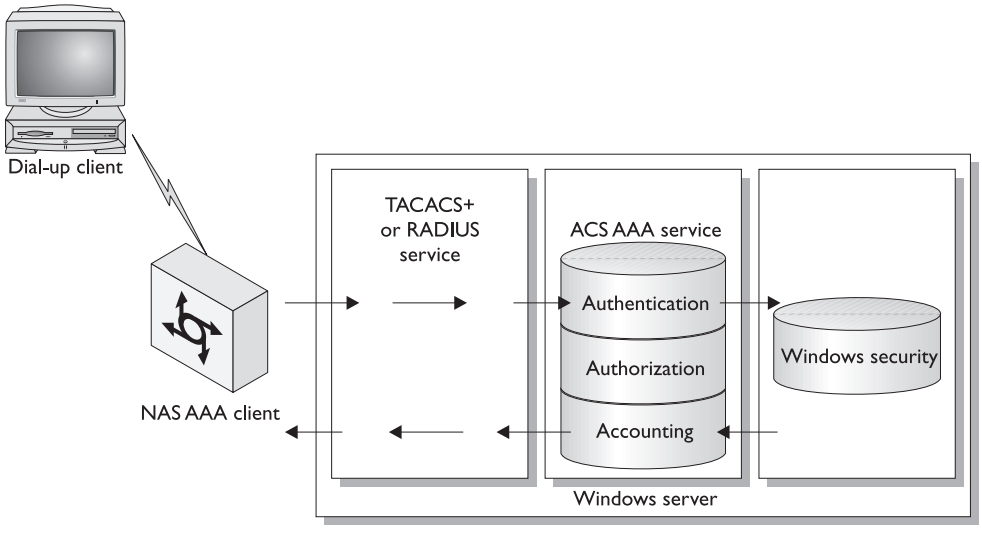

Figure 4-3 Cisco Secure ACS using Windows security database for authentication

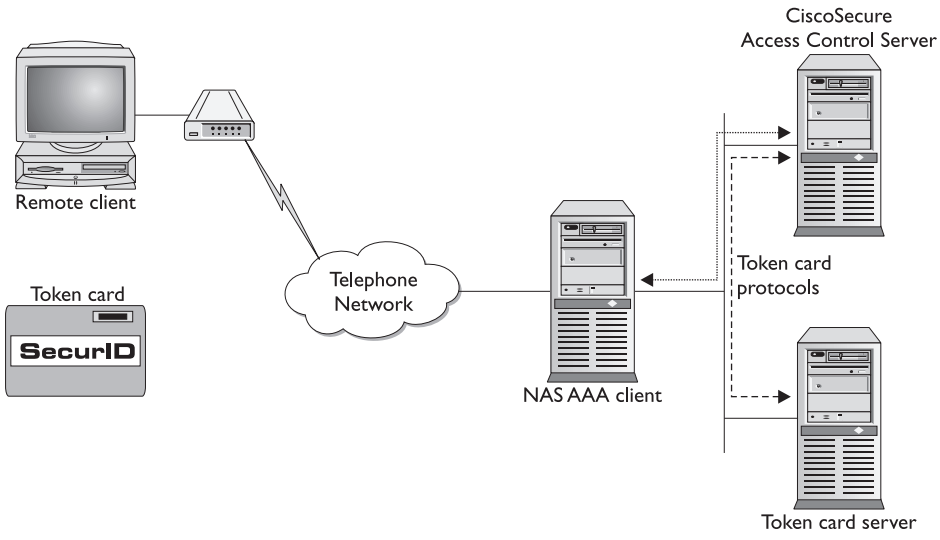

Figure 4-4 Remote user authentication using Token Card

CSDBSync

CSDBSync is an alternative to using the ODBC dynamic link library (DLL) to synchronize the Cisco Secure ACS database with third-party RDBMS systems. Because version 2.4, CSDBSync synchronizes AAA client, AAA server, network device groups (NDGs), and proxy table information.

CSLog

CSLog is the service that captures and places logging information. CSLog gathers data from the TACACS+ or RADIUS packet and CSAuth, and formats the data into the comma-separated value (CSV) files that can be imported into spreadsheets supporting the format.

CSMon

CSMon minimizes downtime in a remote access network environment. CSMon works for both TACACS+ and RADIUS by automatically detecting which protocols are in use. CSMon performs four basic activities:

- **Monitoring** Monitors the overall status of Cisco Secure ACS and the host system it's running on. It uses the Windows Event Log and Performance Monitor to monitor overall system health, including disk, CPU, and memory utilization.
- **Recording** Records and reports all exceptions to a special log file that can be used to diagnose problems.

- **Notification** Alerts the administrator to potential problems and real events regarding Cisco Secure ACS, and records all such problems. The default notification method is Simple Mail Transfer Protocol (SMTP) e-mail, but scripts can be written to enable other methods, such as pager notification.
- **Response** Can be configured to attempt to fix detected problems automatically and intelligently, such as running scripts to restart stopped services.

CSTacacs and CSRadius

The *CSTacacs* and *CSRadius* services communicate between the CSAuth module and the access device requesting authentication and authorization services. CSTacacs is used to communicate with TACACS+ devices, and CSRadius is used to communicate with RADIUS devices. Both services can run at the same time. CSTacacs and CSRadius services must be configured from CSAdmin.

Each module can be started and stopped individually from within the Microsoft Service Control Panel; or, with the exception of CSAdmin, each can be stopped from within the Cisco Secure ACS HTML interface.

System Performance

Cisco Secure ACS's performance capabilities, like most server services, are largely dependent on the resources of the Windows server it's installed on. Other factors include network topology, network management, and the selection of user authentication databases. Common sense would rightly indicate that a faster processor, increased memory, and high-speed connectivity will increase both the speed and volume of authentications per second.

The following items are general indicators of system performance, but the actual Cisco Secure ACS performance on a particular network could vary based on the environment and AAA configuration.

- **Maximum users supported** Technically, no limit exists to the number of users the Cisco Secure ACS user database can support if disk space is available. Cisco has successfully tested Cisco Secure ACS with databases greater than 100,000 users. While a single Cisco Secure ACS server using multiple databases might be able to support 300,000 to 500,000 users, using replicated multiple Cisco Secure ACS servers would increase that number substantially.
- **Transaction processing** A single minimal ACS server with a 10,000 user database might be able to process 80 RADIUS logins, plus approximately 40 TACACS+ logins per second. Increasing memory and/or the number and size of the processors would increase these numbers, while increasing the size of the database will reduce performance.
- **Maximum number of AAA client devices** Approximately 2,000 network devices running any AAA client.

Features of CiscoSecure ACS for UNIX

The most recent UNIX version is *CiscoSecure Access Control Server v2.3 for UNIX (Solaris)* to control the authentication, authorization, and accounting of users accessing the corporate network, Internet, or intranet. Using Cisco Secure ACS, network administrators can control the following:

- Which users can access the network from either wired or wireless connections
- What privileges each user can have while in the network
- What accounting information is kept for capacity planning, account billing, or security audits

Features and Benefits

CiscoSecure ACS v2.3 for UNIX provides the following features that allow network administrators to scale and deploy secure network services with centralized control, access management, and accounting within the Cisco Secure ACS framework:

- Simultaneous TACACS+ and RADIUS support for flexibility in implementation.
- HTML/JAVA GUI simplifies and speeds configuration for user and group profiles. SSL is also supported for secure server configuration.
- Administration of users using groups for maximum flexibility and to facilitate enforcement and changes of security policies.
- Token caching of passwords.
- Local and remote domain declaration.
- Virtual private dial-up network (VPDN) allows dial-up users to connect securely to a corporate network through a third-party ISP. VPDN can use AAA servers such as RADIUS and TACACS+ and services for better scalability of VPDN.
- Import mechanism to rapidly import a large number of users.
- Relational database support using Oracle, Sybase, or the included SQL Anywhere.
- Password support that includes Cleartext, DES encrypted, Bellcore S/Key, UNIX /etc/passwd file, Challenge Handshake Authentication Protocol (CHAP), Password Authentication Protocol (PAP), and AppleTalk Remote Access (ARA).
- Token server support for CryptoCard, Secure Computing, and Security Dynamics.
- Time-of-day and day-of-week access restrictions.
- User restrictions based on NAS name, port name, or remote address, including calling line ID (CLID.)
- Disabling of an account on a specific date.
- Disabling of an account after *n* failed attempts to thwart brute force attacks.
- Accounting information stored in the relational database.

Preparing to Install UNIX ACS

Cisco Secure ACS operates on UNIX Server as a service. Remembering the performance issues covered earlier, the UNIX server computer must meet the following minimum hardware and software requirements.

Hardware Requirements

The server must meet the following minimum hardware requirements:

- Sun SPARCstation 20
- CD-ROM drive
- 128MB of RAM
- 256MB of disk swap space
- 500MB of disk space

Software Requirements

The server must meet the following software requirement:

- Solaris V2.51 or V2.6, V7, V8

Third-Party Software Requirements

The server must have a compatible web browser installed. Both Java and JavaScript must be enabled for any web browsers to be used to administer Cisco Secure ACS. Cisco Secure ACS has been tested with the following browsers:

- Microsoft Internet Explorer versions 5.0 and 5.5
- Netscape Communicator version 4.76

External Databases

- Oracle v7.33
- Sybase v11.1

NAS Minimum IOS Requirements

- Cisco IOS v11.1 (TACACS+)
- Cisco IOS v11.2 (RADIUS)

Installing Cisco Secure ACS 3.0 for Windows

Cisco Secure ACS operates on Windows 2000 Server or NT 4.0 Server as a service. Cisco Secure ACS can operate on Windows 2000 Advanced Server and Windows 2000 Datacenter Server implementations *if Microsoft Clustering Services are not installed*. Cisco Secure ACS can run on a domain controller or a member server.

Hardware Requirements

Remembering the performance issues just covered, the Windows server computer must meet the following minimum hardware, operating system (OS), and third-party software requirements.

- Intel class Pentium 550 MHz PC or compatible
- 256MB of RAM
- 250MB of free disk space or more if you're running your database on the same machine
- Minimum resolution of 256 colors at 800 × 600 or lines

Operating System Requirements

The server must be running a functioning English-language version of one of the following Microsoft Windows OSs:

- Windows 2000 Server with Service Pack 1 or Service Pack 2 installed
- Windows 2000 Advanced Server with Service Pack 1 or Service Pack 2 installed (Microsoft Clustering Services must *not* be installed)
- Windows 2000 Datacenter Server with Service Pack 1 or Service Pack 2 installed (Microsoft Clustering Services must *not* be installed)
- Windows NT Server 4.0 with Service Pack 6a installed

While, technically, Windows service packs can be applied either before or after installing Cisco Secure ACS, if the service packs are installed before the Cisco Secure ACS software, the process will go more smoothly. If not, the Cisco Secure ACS installation program displays warnings that the required service pack isn't present. If a service pack message is displayed, continue the installation, and then install the required service pack *before* starting user authentication with Cisco Secure ACS.

Third-Party Software Requirements

The server must have a compatible web browser installed. Both Java and JavaScript must be enabled for any web browsers to be used to administer Cisco Secure ACS. Cisco Secure ACS has been tested with the following browsers:

- Microsoft Internet Explorer versions 5.0 and 5.5
- Netscape Communicator version 4.76

NAS Minimum IOS Requirements

- Cisco IOS v11.1 (TACACS+)
- Cisco IOS v11.2 (RADIUS)

Network Requirements

The following network requirements should be in place before you begin to install Cisco Secure ACS:

- A web browser meeting the previous third-party software requirements must be properly installed on any Windows server to be used to administer Cisco Secure ACS.
- For full TACACS+ and RADIUS support on any Cisco IOS devices, make sure the AAA clients are running Cisco IOS Release 11.2 or later.
- Make sure any non-Cisco IOS AAA clients can be configured with TACACS+ or RADIUS support, or both.
- Make sure the Windows server can ping the AAA clients.
- For Cisco Secure ACS to use the Grant Dial-In Permission To User feature when authorizing Windows network users, this option must be selected in the Windows NT User Manager or Windows 2000 Active Directory Users And Computers for the applicable user accounts.
- Make sure all network cards in the server that will host the Cisco Secure ACS software are enabled. Disabled NICs will slow installation because of delays caused by Microsoft CryptoAPI.
- Make sure any dial-in, VPN, or wireless clients can successfully connect to the applicable AAA clients.

Back Up Server Data

As with any major change to a server, backing up the Windows server installation, including the Windows Registry, makes sense. When upgrading or reinstalling Cisco Secure ACS, note the following issues:

- Back up the Cisco Secure ACS configuration and database, and then copy the backup file to a drive other than one that's local to the Cisco Secure ACS server.
- When upgrading Cisco Secure ACS, the backup created can't be used after a successful upgrade. This backup only provides a recovery option for the previous Cisco Secure ACS installation.
- The ACS Backup feature temporarily stops any Cisco Secure ACS services during the backup.

Gathering Information Required During Installation

During new Cisco ACS installations or any upgrades and reinstallations that don't preserve the existing configuration, specific information about your Windows server and an AAA client on your network need to be entered. Collect the applicable information before beginning the installation procedure. Upgrades or Cisco Secure ACS reinstallations intended to preserve the existing configuration and database don't require this information.

To collect information required during the installation of Cisco Secure ACS, follow these steps:

1. For the first AAA client to be configured to use Cisco Secure ACS's AAA services, determine which AAA protocols and vendor-specific attributes you want to implement.
   - TACACS+ (Cisco IOS)
   - RADIUS (Cisco Aironet)
   - RADIUS (Cisco BBSM)
   - RADIUS (Cisco IOS/PIX)
   - RADIUS (Cisco VPN 3000)
   - RADIUS (Cisco VPN 5000)
   - RADIUS (IETF)
   - RADIUS (Ascend)
   - RADIUS (Juniper)
   - RADIUS (Nortel)
2. Record the name of the AAA client.
3. Record the IP address of the AAA client.
4. Record the Windows 2000/NT server IP address.
5. Record the TACACS+ or RADIUS key (shared secret).

Administering and Troubleshooting Cisco Secure ACS for Windows

Once installed, Cisco Secure ACS is configured and administered using a web browser through the HTML interface, enabling easy configuration from any host on the LAN or WAN. The Cisco Secure ACS HTML interface uses HTML and some Java functions for ease of use. This program design keeps the interface straightforward and responsive. Figure 4-5 shows the Cisco Secure ACS HTML interface.

From the HTML interface, you can easily view and edit user and group information, stop and restart services, add or change remote administrators, modify AAA client information, back up the system, view reports, and more. Reports track connection activities,

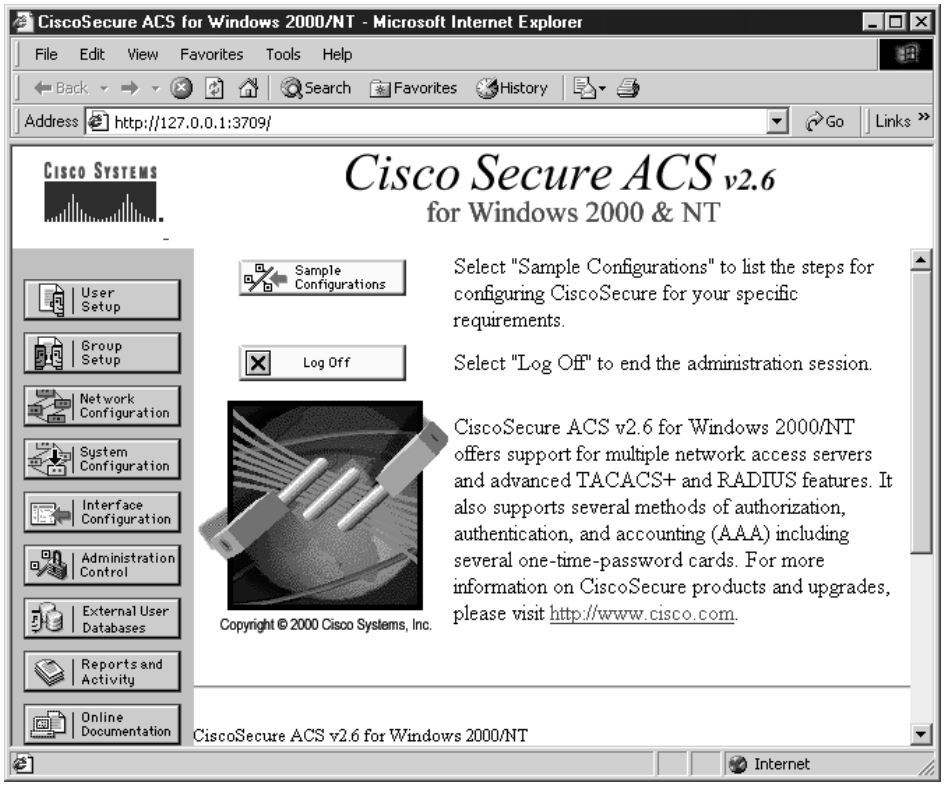

Figure 4-5 Cisco Secure ACS HTML interface

show users who are logged in, list any failed authentication and authorization attempts, and show a history of the recent tasks of administrators.

The HTML interface has three vertical frames that perform the following functions:

- Navigation bar
- Configuration area
- Display area

Figure 4-6 shows the three frames of the ACS HTML configuration screen.

Navigation Bar

The *navigation bar,* the gray frame on the left of the browser window, contains the task buttons. Each button changes the configuration area (second panel) to a section of the

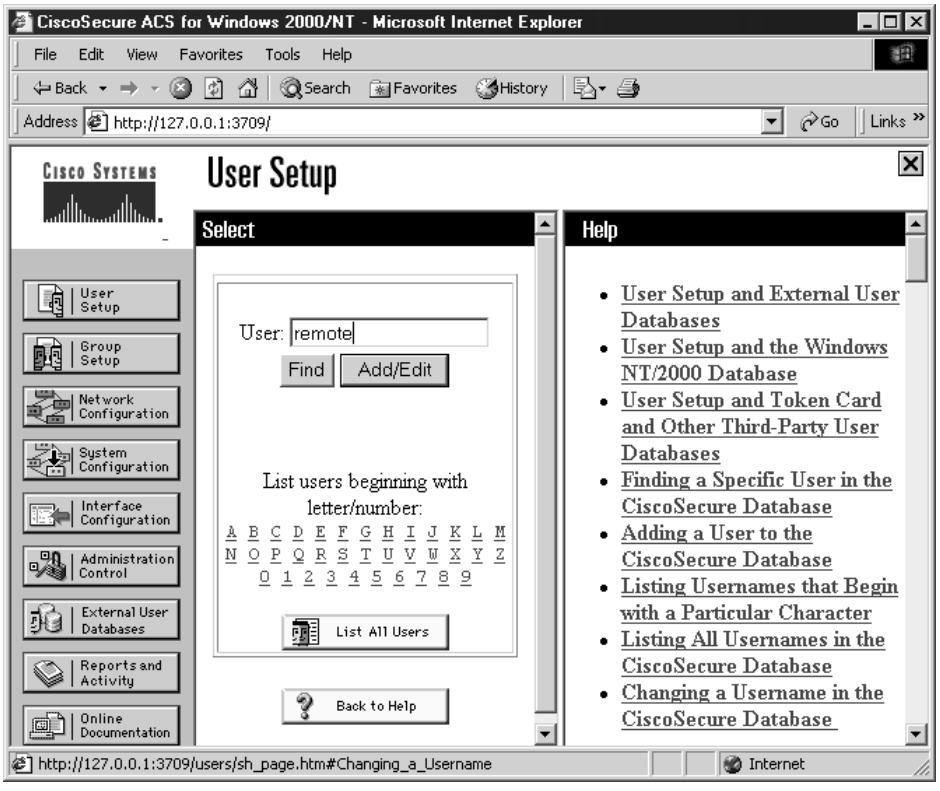

Figure 4-6 ACS configuration screen

Cisco Secure ACS application. This frame doesn't change; it always contains the following buttons:

User Setup	Add and edit user profiles.
Group Setup	Configure network services and protocols for user groups.
Network Configuration	Add and edit network access devices.
System Configuration	Configure database information and accounting.
Interface Configuration	Display or hide product features and options.
Administration Control	Define and configure access policies.
External User Databases	Configure external databases for authentication.
Reports And Activity	Display various accounting and logging summaries.
Online Documentation	View the Cisco Secure ACS User Guide.

Configuration Area

The *Configuration Area,* middle frame, displays web pages that belong to one of the sections represented by the buttons in the navigation bar. The configuration area is where you add, edit, or delete program setup information. For example, in Figure 4-6, it's possible to find a particular user and change their user information.

Most configuration pages have an appropriate Submit button at the bottom that's used to confirm your changes. Figure 4-7 shows the Submit button options. The Submit + Restart button is used for those configuration changes that require stopping and restarting the services. If you don't click the Submit button or click the Cancel button, the changes won't be saved.

Display Area

The Display Area, right-side frame, shows one of the following options, depending on the button selected in the navigation bar.

Online Help	Basic help about the page currently shown in the Configuration Area. This help isn't intended to be in-depth information but, instead, basic information about the topic in the middle frame. For more detailed information, click Section Information at the bottom of the page to go to the applicable part of Online Documentation.
Reports Or Lists	Displays available lists or reports, including accounting summary reports. Most listings are hyperlinks to the specific configuration views, so clicking the link enables you to edit that item.
System Messages	Displays system and error messages after you submit your changes, indicating the nature of the problem. Any incorrect information remains in the Configuration Area for easy review and correction.

Accessing the HTML Interface

The HTML interface can be reached by using a web browser from anywhere on the network at either of the following URLs. The first two lines show the syntax, while the last two are examples.

```
http://windows-server-IP-address:2002
http://windows-server-host-name:2002
http://192.168.1.3:2002
http://wilson:2002
```

Figure 4-7 Configuration Area Submit and Cancel buttons

If Cisco Secure ACS is installed on the local server you're accessing with a web browser, it's possible to use either of the following two URLs. These commands take advantage of default naming and addressing standards.

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

Remote administrative sessions always require a login using a valid administrator name and password, as configured in the Administration Control section.

To access the Cisco Secure ACS configuration HTML interface, follow these steps:

1. Open a web browser.
2. In the Address or Location bar in the web browser, type the applicable URL.
3. If the Cisco Secure ACS for Windows 2000/NT Login page appears, follow these steps:
   a. Type a valid Cisco Secure ACS administrator name in the User name box.
   b. Type a valid administrator password in the Password box.
   c. Click the Login button.

The Cisco Secure ACS for Windows 2000/NT initial page appears, as shown earlier in Figure 4-5.

Logging Off the HTML Interface

Click the Logoff button to end the ACS session. Failure to do so might allow unauthorized access by someone using the web browser after you, or even unauthorized access through the HTTP port left open to support the administrative session. Cisco Secure ACS can timeout unused administrative sessions.

Remote Administrative Session Issues

Many organization security policies wouldn't allow remote access to resources like the Cisco Secure ACS server. For this reason, various technologies are designed to keep outsiders safely outside the network. The nature of these technologies is often inconsistent with the ACS architecture.

While this isn't always practical, the recommendation is that remote administration sessions not have any of the following technologies between the administration browser and the Cisco Secure ACS server.

- HTTP proxy server
- NAT gateway
- Firewall devices

Each of these technologies can interfere with or prevent connection between the browser and the ACS server. Because these are common technologies, particularly when accessing the ACS server from outside the network, the following explanations and suggestions might be useful.

HTTP Proxy Servers Because of the way proxy servers work between clients and servers, a web browser configured to use a proxy server for a remote administrative session will appear to Cisco Secure ACS server as originating from the IP address of the proxy server, not the address of the remote workstation. ACS remote administrative session tracking requires that each browser resides on a workstation with a unique IP address. ACS administrative sessions using a proxy-enabled web browser is neither tested nor supported.

Suggestion: If the web browser is configured to use a proxy server, disable HTTP proxying before attempting the remote Cisco Secure ACS administrative session.

NAT Gateway If a remote session uses a web browser on a workstation behind a NAT gateway, the gateway substitutes a global "public" address for the workstation's real local IP address. When ACS receives the HTTP requests, the NAT device's public IP address conflicts with the workstation's private IP address *included in the content of the HTTP requests.* Cisco Secure ACS won't allow this.

If the Cisco Secure ACS server is behind a NAT gateway, you could try to configure the NAT gateway to forward all connections to port 2002 to the Cisco Secure ACS server, retaining the port 2002. In addition, all the ports allowed using the HTTP port allocation feature must be similarly mapped. Cisco hasn't tested this and it doesn't support this implementation.

Firewall Devices Firewalls implementing NAT would fall into the last section. For firewalls not performing NAT, remote ACS administrative sessions conducted across the firewall will require additional configuration of both the ACS software and the firewall. This is necessary because ACS assigns a random HTTP port at the beginning of a remote administrative session.

The firewall must be configured to allow HTTP traffic across the range of ports ACS is configured to use. This can be configured using the HTTP port allocation feature to configure the range of TCP ports to be used by Cisco Secure ACS for remote administrative HTTP sessions. The firewall must also be configured to permit HTTP traffic through port 2002 because this is the port a remote web browser must access to initiate an administrative session.

To reduce the risk of malicious discovery of an active administrative port by an unauthorized user, keep the HTTP port range as narrow as possible. Any unauthorized user would have to impersonate, or "spoof," the IP address of the legitimate remote host to use the active administrative session HTTP port.

Suggested Configuration Sequence

No single set of steps, or even order of steps, exists for configuring ACS. The size and topology of the network, the types of network access supported, and even the technological skills of the administrative staff will all be factors. The CCO latest documentation, and security alerts, as well as the ACS documentation feature, should always be consulted before and during the final implementation begins. The following sequence is keyed to the functions represented in the navigation toolbar.

Configure Administrators

Configure at least one administrator (Administration Control button), or there's no remote administrative access and all configuration activity must be done from the server. The administrative policy and security policy should dictate the number and details of administrator accounts.

Configure the ACS Web Interface

You can configure ACS HTML interface (Interface Configuration button) to show only those features and controls you plan to use. This streamlines ACS, making it less difficult to use, less intimidating, and easier to train new administrators to use. The downside, of course, is that features and controls aren't available and are possibly unknown to the staff because they aren't present.

Aspects of the web interface that can be configured in this section include the following:

- **User Data Configuration Options** You can add (or edit) up to five fields for recording information on each user.
- **Advanced Options** Over 20 options ranging from User-Level Network Access Restrictions to Voice-over-IP (VoIP) Accounting options.
- **Protocol Configuration Options for TACACS+** Settings enable the display or hiding of TACACS+ administrative and accounting options.
- **Protocol Configuration Options for RADIUS** Settings enable the display or hiding of RADIUS administrative and accounting options.

Figure 4-8 shows some of the configuration options available in setting up TACACS+.

Configure System

Figure 4-9 shows some of the dozen or more features that can be configured within the System Configuration section (System Configuration button). This is where the Logging options are made that can be used later to produce reports. ACS comprises several Windows NT/2000 services. The Service Control page provides basic status information about the services, and enables you to configure the service log files, and to stop or restart the services.

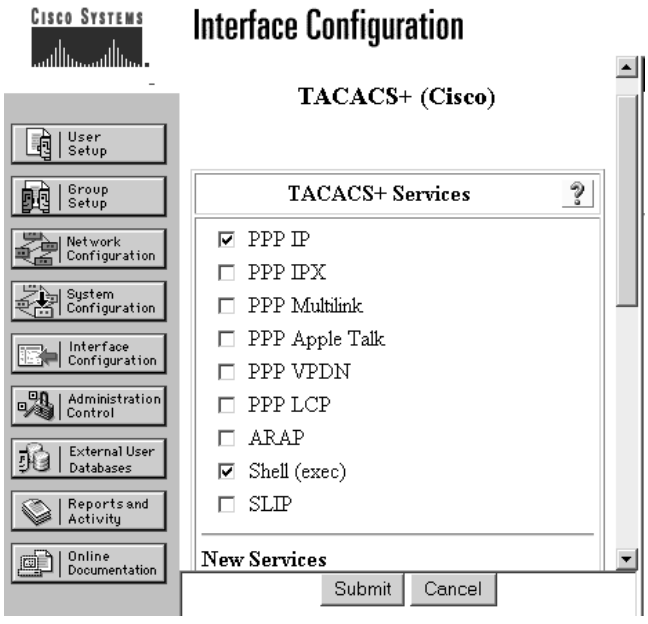

Figure 4-8 Interface configuration setting TACACS+ features

Configure Network

Figure 4-10 shows the network configuration (Network Configuration button) used to control distributed and proxied AAA functions. This section is used to establish the identity, location, and grouping of AAA clients and servers, and to determine what authentication protocols each is to employ.

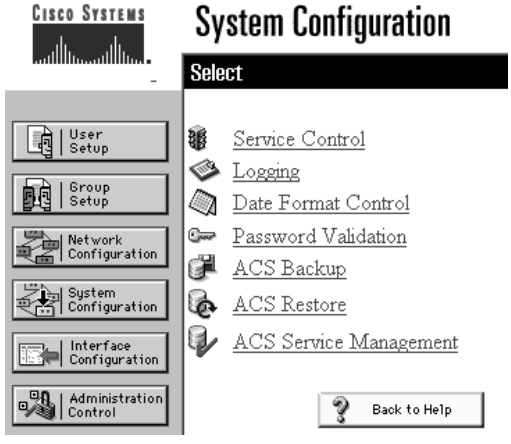

Figure 4-9 ACS System Configuration options

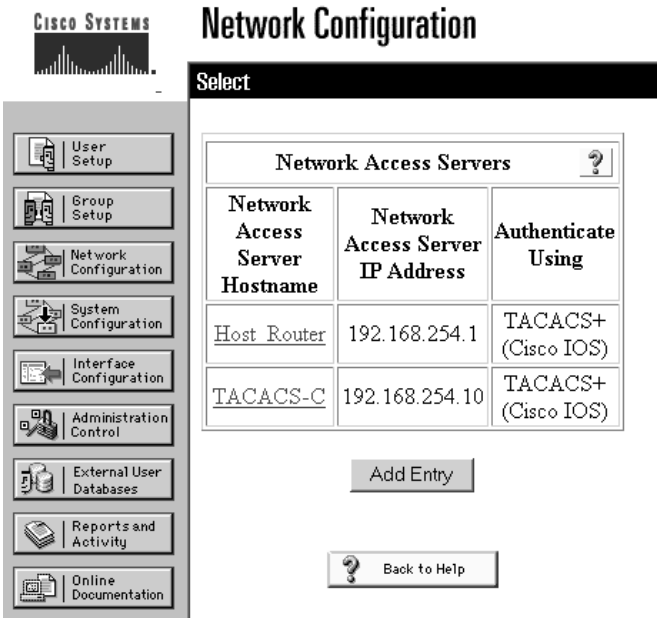

Figure 4-8 Network configuration showing the NAS router and the AAA server

Configure External User Database

The External User Databases button is used to implement an external database to establish and maintain user authentication accounts. This configuration usually is based on the existing network administration mechanisms. In addition to implementing an external user database (or databases), this section is used to define requirements for ACS database replication, backup, and synchronization.

Configure Shared Profile Components

The Shared Profile Components section enables administrators to develop and name reusable, shared sets of authorization components, which might be applied to one or more users, or groups of users, and referenced by the assigned name within their profiles. These include network access restrictions (NARs), command authorization sets, and downloadable PIX ACLs.

- NARs enable the administrator to define additional authorization conditions that must be met before a user can gain access to the network.
- Command authorization sets provide a central mechanism to control the authorization of each command on each network device.
- Downloadable PIX ACLs enable the creation of an ACL once, in Cisco Secure ACS, and then load that ACL to any number of PIX firewalls that authenticate using the Cisco IOS/PIX protocol.

These shared profile components enhance the scalability of the selective authorization feature. Shared profile components, once configured, can be applied to many users or groups, and they eliminate having to configure the authorization explicitly for each user group for each possible command on each possible device.

Configure Groups

Figure 4-11 shows some of the features that can be assigned to groups (Group Setup button), much the same way group permissions are used in the Windows authentication model. It's always easier to deal with group privileges, permissions, or features, and then make sure the appropriate users are in the correct groups. This not only facilitates setting up a new user, it also facilitates removing all of a user's permissions and features when they leave the organization.

Configure Users

Figure 4-12 shows defining a User account (User Setup button). Once the groups are defined, it's time to create the user accounts. Note, unlike the Windows groups methods, a user can belong to only one user group and user level settings always override group level settings.

Configure Reports

The Reports And Activity section (Reports or Lists button) is used to specify the nature and scope of logging that ACS performs, which, ultimately, determine the reports that can be generated.

Figure 4-9
Example of group setup options

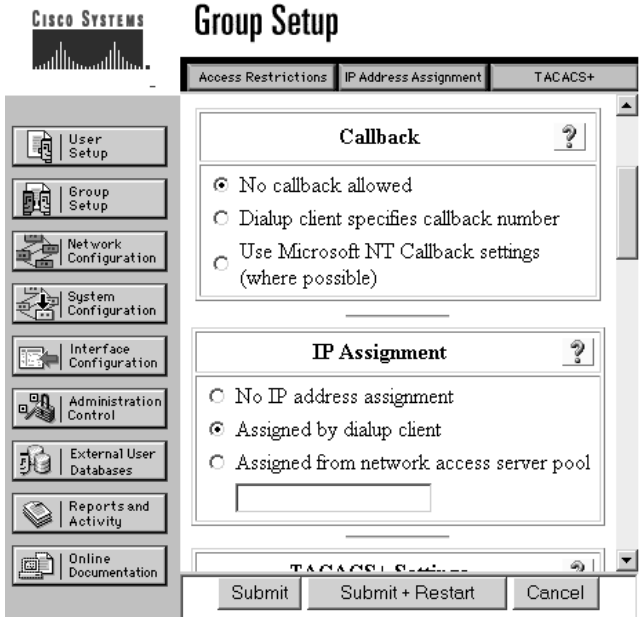

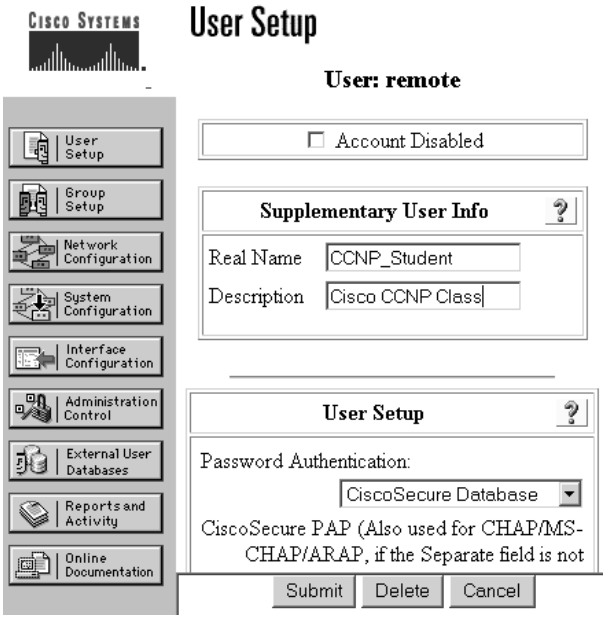

Figure 4-10 User account entry screen

TACACS+ Overview

TACACS+ is an authentication protocol that allows a network access server to communicate with an authentication server to determine if a user has access to the network.

TACACS+ is a new protocol developed by Cisco that replaces two earlier industry standards—TACACS and XTACACS (Extended). TACACS+ is *not* compatible with the two older protocols. Cisco has submitted TACACS+ protocol specification in a draft RFC to the IETF for development of a standard and for those customers interested in developing their own TACACS+ software.

TACACS+ server services are maintained in a database on a TACACS+ daemon running on a Windows 2000/NT or UNIX host. Cisco's servers supporting TACACS+ include CiscoSecure ACS for Windows, CiscoSecure UNIX, and Cisco Access Registrar. Cisco Access Servers (Cisco Secure ACS) can implement both TACACS+ and RADIUS. The underlying architecture of TACACS+ protocol complements the AAA architecture.

TACACS+ fully supports the AAA architecture by separating the authentication, authorization, and accounting. This allows the flexibility of using another service, such as Kerberos, for authentication, while still using TACACS+ for authorization and/or accounting.

TACACS+ uses TCP for connection-oriented transport between clients and servers. TCP port 49 is reserved for TACACS+. The acknowledgments (TCP ACK) provide indications that a request has been received. This same TCP process uses RST packets to provide immediate indication of a failed (or offline) authentication server. TCP keepalives can be used to watch for failed servers and to facilitate rapid failover between multiple connected

authentication servers. TCP scales better and adapts better to growing and/or congested networks.

TACACS+ supports bidirectional challenge/response, like CHAP, between the two network access servers.

In addition to supporting SLIP and PPP encapsulation protocols, TACACS+ supports the following protocols:

- Novell Asynchronous Services Interface (NASI)
- X.25 PAD connection
- Net BIOS Frame Protocol Control protocol
- AppleTalk Remote Access protocol (ARAP)

Packet Encryption TACACS+ encrypts the entire data payload of the packet, leaving only the standard TACACS+ header in cleartext. While it's possible for debugging purposes to leave the body of the packets unencrypted, normal operation will fully encrypt the body for more secure communications. A field in the header indicates whether the body is encrypted.

- TACACS+ supports two methods for controlling the authorization of router commands on either a per-user or per-group basis.
- Assign commands to privilege levels and have the router use TACACS+ to verify that the user is authorized at the specified privilege level.

Explicitly define on the TACACS+ server the commands allowed on a per-user or per-group basis.

Configuring Cisco Secure ACS and TACACS+

Basic Cisco Secure ACS configuration for TACACS+ support is performed using the Interface Configuration button on the ACS HTML interface. The interface settings enable you to display or hide TACACS+ administrative and accounting options. You can simplify the HTML interface by hiding the features you don't use. Figure 4-13 shows configuring TACACS+ options and features.

The TACACS+ section includes three areas:

TACACS+ Services Settings	This section lists the most commonly used TACACS+ services and protocols. Select each TACACS+ service to appear as a configurable option on either the User Setup page or the Group Setup page.
New Services	Specify any services or protocols unique to your network configuration.
Advanced Configuration Options	Advanced TACACS+ features.

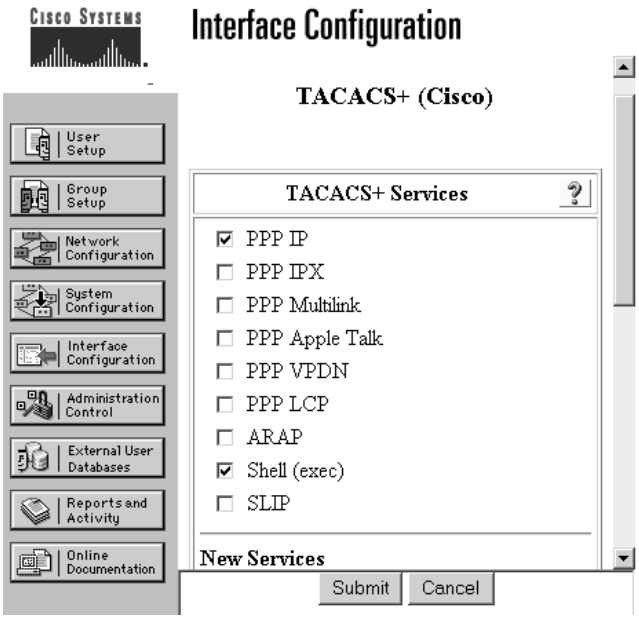

Figure 4-11 Configuring TACACS+ features

To configure the user interface for TACACS+ options, follow these steps:

1. Click the Interface Configuration button.
2. Click TACACS+. The TACACS+ Interface Configuration section appears.
3. In the TACACS+ Services table, select the check box for each TACACS+ service you want displayed on the applicable setup page.
4. To add new services and protocols, follow these steps:
   a. In the New Services section of the TACACS+ Services table, type in the Service and Protocol to add.
   b. Use the appropriate check box to select those to be displayed for configuration either under User Setup or Group Setup, or both.
5. In the Advanced Configurations Options section, select the check boxes of the display options you want to enable.
6. When you finish setting TACACS+ interface display options, click Submit. The selections made in this procedure determine what TACACS+ options Cisco Secure ACS displays in other sections of the HTML interface.

Configure NAS to TACACS+ Server Communication

Chapter 3 looked at the various commands to implement AAA features on the NAS. This chapter reviews the basic commands to allow the AAA client running on a NAS to locate and communicate with a Cisco Secure ACS TACACS+ server.

To define one or more TACACS servers, use the **tacacs-server host** global configuration command. Use the no form of this command to delete the specified server. The syntax is

Rtr1(config)#tacacs-server host {*hostname* | *ip address*} [single-connection] [port *port#*] [timeout *seconds*] [key *string*]
Rtr1(config)#no tacacs-server host *hostname*

hostname	Name of the TACACS+ server.
ip-addresss	IP address of the TACACS+ server.
single-connection	(Optional) Specify that the router maintain a single open connection for confirmation from the AAA/TACACS+ server. This command contains the no autodetect feature and fails if the host isn't running a CiscoSecure daemon.
port	(Optional) Overrides the default, which is port 49.
port#	(Optional) Port number of the server (range 1 to 65535).
timeout	(Optional) Overrides the global timeout value set with the global **tacacs-server timeout** command for this server only.
seconds	(Optional) New timeout interval in seconds.
key	(Optional) Specify an authentication and encryption key. This must match the key used by the TACACS+ daemon. This key overrides the key set by the global command **tacacs-server key** for this server only.
string	(Optional) Character string to use as the authentication and encryption key.

Multiple **tacacs-server host** commands can be used to specify additional host servers. The Cisco IOS software searches for hosts in the order in which they're specified. Use the single-connection, port, timeout, and key options only when running an AAA/TACACS+ server.

Because some of the parameters of the **tacacs-server host** command override global settings made by the **tacacs-server timeout** and **tacacs-server key** commands, this command can be used to enhance security on a network by uniquely configuring individual TACACS+ connections.

The following AAA example specifies that the router first try the CiscoSecure TACACS+ host 192.168.1.4. If 192.168.1.4 is unavailable, then use tac-serv1.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server host tac-serv1
```

The next AAA example specifies that the router consult the CiscoSecure TACACS+ host named 192.168.1.4 on port number 51. The timeout value for requests on this connection is three seconds. The encryption key is a_secret.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4 single-connection port 51
     timeout 3 key a_secret
```

Define TACACS+ Server Key Option

The authentication and encryption key must be set using the **tacacs-server key** global configuration command. Use the no form of this command to disable the key.

```
Rtr1(config)#tacacs-server key key
Rtr1(config)#no tacacs-server key [key]
```

key	A character entry, this key must match the key used on the TACACS+ daemon. Any leading spaces are ignored, *but spaces within and at the end of the key are not ignored.* Don't enclose the key in quotation marks unless those quotation marks are part of the key.

The following example shows an AAA TACACS+ implementation with the authentication and encryption key set to seattle19:

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server key seattle19
```

Verifying TACACS+

Troubleshooting and verifying TACACS+ can be rather complex because it's used so often with other features, such as PPP. Remember to use the troubleshooting commands associated with any protocols or technologies working with TACACS+.

The show Commands

Two show commands that are useful in debugging AAA are the following:

show running-config	To verify local TACACS+ is configured correctly, enter
show tacacs	To verify network connectivity between NAS and AAA server

The debug Commands

Cisco IOS debug command output provides a valuable source of information and feedback concerning state transitions and functions within the AAA environment. In addition to debug command output gathered directly from devices running Cisco IOS, the Cisco AAA server can be configured to collect operational diagnostics. Use the following debug commands to capture AAA-related transitions and functions.

debug condition user *username*	Sets conditional debugging for a specific user and generates output debugs related to the user
debug aaa authentication	Displays authentication information with TACACS+ client/server interaction
debug aaa authorization	Displays authorization information with TACACS+ client/server interaction

debug aaa accounting	Displays accounting information with TACACS+ client/server interaction
debug tacacs	Displays TACACS+ interaction between the IOS client and the AAA server
debug ppp negotiation	Sees if a client is passing PPP negotiation
debug ppp authentication	Sees if a client is passing authentication
debug ppp error	Displays protocol errors and error statistics associated with PPP connection negotiation and operation

Configure NAS to RADIUS Server Communication

Chapter 4 looked at the various commands to implement AAA features on the NAS. This chapter reviews the basic commands to allow the AAA client running on a NAS to locate and communicate with Cisco Secure ACS RADIUS+ server.

To define one or more RADIUS servers, use the **radius-server host global** configuration command. Use the no form of this command to delete the specified RADIUS host. The syntax is

Rtr1(config)#radius-server host {*hostname* | *ip-address*} [auth-port *port-number*] [acct-port *port-number*] [timeout *seconds*] [retransmit *retries*] [key *string*] [alias{*hostname* | *ip-address*}]
Rtr1(config)#no radius-server host {*hostname* | *ip-address*} [auth-port *port-number*] [acct-port *port-number*] [timeout *seconds*] [retransmit *retries*] [key *string*]

hostname	DNS name of the RADIUS server.
ip-address	IP address of the RADIUS server.
auth-port	(Optional) To specify a UDP destination port for authentication requests.
port-number	(Optional) Port number for authentication requests. The host isn't used for authentication if set to 0.
acct-port	(Optional) To specify a UDP destination port for accounting requests.
port-number	(Optional) Port number for accounting requests. The host isn't used for accounting if set to 0.

Use multiple **radius-server host** commands to specify multiple hosts. The software searches for hosts in the order they're specified.

The following example specifies host1 as the RADIUS server and uses default ports for both accounting and authentication.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host host1.domain.com
```

The next example defines port 12 as the destination port for authentication requests and port 16 as the destination port for accounting requests on a RADIUS host named

192.168.1.4. Because entering a line resets all the port numbers, you must specify a host and configure both the accounting and authentication ports on a single line.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4 auth-port 12 acct-port 16
```

To use separate servers for accounting and authentication, use the zero port value as appropriate. The following example specifies that RADIUS server 192.168.1.4 be used for accounting, but not for authentication, and RADIUS server host1 be used for authentication, but not for accounting:

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4 auth-port 0
Rtr1(config)#radius-server host host1.domain.com acct-port 0
```

Define RADIUS Server Key Option

The authentication and encryption key must be set using the **radius-server key** global configuration command. Use the no form of this command to disable the key.

Rtr1(config)#radius-server key *key*
Rtr1(config)#no radius-server key [*key*]

key	A character entry, this key must match the key used on the RADIUS daemon. Any leading spaces are ignored, *but spaces within and at the end of the key are not.* Don't enclose the key in quotation marks unless they're part of the key.

The following example shows an AAA RADIUS implementation with the authentication and encryption key set to seattle19:

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4
Rtr1(config)#radius-server key seattle19
```

Chapter Review

The CiscoSecure Access Control Server (ACS) is a server-based application offering centralized user access control support for the various network devices. CiscoSecure ACS can be installed on either Windows or UNIX servers that meet minimum resource requirements.

ACS allows the administrator to manage user access for Cisco IOS routers, firewalls, Cisco Catalyst switches, and virtual private networks (VPNs), as well as newer technologies like dial-up and broadband DSL, cable access solutions, voice over IP (VoIP), Cisco wireless implementations, and third-party vendor TACACS+ offerings.

ACS is a strategic product at Cisco, incorporating new advanced features as they become available to increase the diversity and granularity of implementations. Current support includes the use of network OS security databases and management tools, user and administrative access reporting, database synchronization and user importation

tools, Lightweight Directory Access Protocol (LDAP) user authentication support, dynamic quota generation, and restrictions such as time of day and day of week, as well as user and device group profiles.

Questions

1. Which two of the following are server platforms supported by Cisco Secure ACS?
   - **A.** Novell NetWare
   - **B.** Microsoft Windows
   - **C.** UNIX (Sun Solaris)
   - **D.** Linux
2. Which one of the following isn't controllable by a network administrator using Cisco Secure ACS for Windows?
   - **A.** What accounting information is kept for capacity planning, account billing, or security audits
   - **B.** What privileges each user will have while in the network
   - **C.** What encryption method will be used for interdevice communications
   - **D.** Which users can access the network from either wired or wireless connections
   - **E.** What access and command controls are enabled for each configuration administrator
3. Which IEEE standard is being introduced to support CiscoSecure ACS in the LAN?
   - **A.** 802.1q
   - **B.** 802.11b
   - **C.** 802.1*x*
   - **D.** 802.3au
4. True or False. Cisco Secure ACS for Windows does *not* support MSCHAP version 2.
   - **A.** True
   - **B.** False
5. True or False. Cisco Secure ACS for Windows supports per-user access control lists.
   - **A.** True
   - **B.** False
6. True or False. Cisco Secure ACS for Windows requires a choice between TACACS+ and RADIUS support.
   - **A.** True
   - **B.** False

7. Cisco Secure ACS supports Token Card authentication on which platform?
   A. Windows only
   B. UNIX only
   C. Both UNIX and Windows
   D. Neither one
8. Which Cisco Secure ACS for Windows service is the internal web server?
   A. CSAuth
   B. CSMon
   C. CSLog
   D. CSAdmin
9. Which Cisco Secure ACS-supported database would provide fastest authentication?
   A. Windows NT/2000 user database
   B. Novell NDS
   C. Token Card servers
   D. CiscoSecure user database
10. Which of the following does *not* impact Cisco Secure ACS server performance?
    A. Number and size of the CPUs
    B. Amount of server memory
    C. Number of users in the database
    D. Remote users connection speed
11. Which of the following won't work as a Cisco Secure ACS platform?
    A. Windows 2000 Server
    B. Windows 2000 Advanced Server
    C. Windows 2000 Datacenter Server
    D. Windows 2000 Professional
    E. Windows NT Server 4.0
12. True or False. Cisco Secure ACS for Windows and UNIX are identical, and they provide the same features and service.
    A. True
    B. False
13. True or False. One limitation of Cisco Secure ACS for UNIX is its lack of support for TACACS+.
    A. True
    B. False

14. Once installed, Cisco Secure ACS is configured and administered using what?
    - **A.** CLI only
    - **B.** Web browser only
    - **C.** CLI and/or web browser
    - **D.** Windows Management Interface
15. Which two of the following will only access the ACS HTML interface using a web browser from the server on which Cisco Secure ACS is installed?
    - **A.** http://*Windows-server-IP-address*:2002
    - **B.** http://localhost:2002
    - **C.** http://*Windows-server-host-name*:2002
    - **D.** http://127.0.0.1:2002

Answers

1. **B. and C.** Microsoft Windows and UNIX (Sun Solaris)
2. **C.** What encryption method will be used for interdevice communications
3. **C.** 802.1*x*
4. **B.** False. With ACS for Windows v3.0, it now supports MSCHAP.
5. **A.** True
6. **B.** False. It will support both simultaneously.
7. **C.** Both UNIX and Windows
8. **B.** CSMon
9. **D.** CiscoSecure user database because it's internal to ACS
10. **D.** The remote users connection speed
11. **D.** Windows 2000 Professional. It must be a version of the server.
12. **B.** False
13. **B.** False. Cisco Secure ACS for UNIX supports both.
14. **B.** Web browser only
15. **B.** http://localhost:2002 and d. http://127.0.0.1:2002

CHAPTER 5

Securing Cisco Perimeter Routers

In this chapter, you will learn about:

- Eavesdropping
- Limiting unnecessary services
- Denial of service attacks
- Unauthorized access
- Lack of legal IP addresses
- Rerouting attacks
- Lack of information about an attack

This chapter looks at those IOS features that can be used on the borders of the network to ward off unwanted and malicious traffic. Chapters 6 and 7 pick up with those additional features included in the Cisco IOS Firewall feature set. In many ways, this chapter pulls together information covered in the first three chapters. In Chapter 1, you learned that the four most common types of network attacks are reconnaissance attacks, access attacks, denial of service (DoS) attacks, and data manipulation attacks.

The *perimeter router* is the first line of defense against each of these threats. While the perimeter router can often thwart reconnaissance, access, and data manipulation attacks, it typically requires the assistance of the company's ISP to deal with DoS attacks.

Perimeter Router Terms and Concepts

A *network*, like a city or village, has an area where the network meets the outside world. Unlike a city, this interface is limited to those devices—hopefully routers—that have connections both to the inside of the network and with the outside world. Typically, an ISP is at the other end of that connection, and then the Internet.

In the preceding two chapters, you looked at network access servers (NAS), which also constitute a part of the network boundary. These NAS connections are a little different, in that security technologies like AAA can be set up to limit access to only those remote dial-in users with preestablished login authority. On the other hand, the network boundary, or perimeter, routers often must allow access to persons unknown to the corporate

shared resources, such as web and FTP servers. Furthermore, perimeter router connections also provide access for the Internet, which then exposes the network to other risks.

Simple Secure Network Design

In Chapter 1, you saw a simplified example of a secure network design. Figure 5-1 expands on that view. The perimeter router represents the border crossing or the demarcation point between the outside world and the internal network.

In a small branch office or in a telecommuter's home, a single perimeter router might be the only barrier between the inside and outside network. The features and configuration of that lone device would be the only line of defense. Many larger networks use a design similar to that shown in Figure 5-1, called screened subnet architecture. In *screened subnet architecture*, two devices work together with a router as the perimeter device and the firewall as the second line of defense. The result is three types of networks: inside, outside, and an area called the demilitarized zone (DMZ).

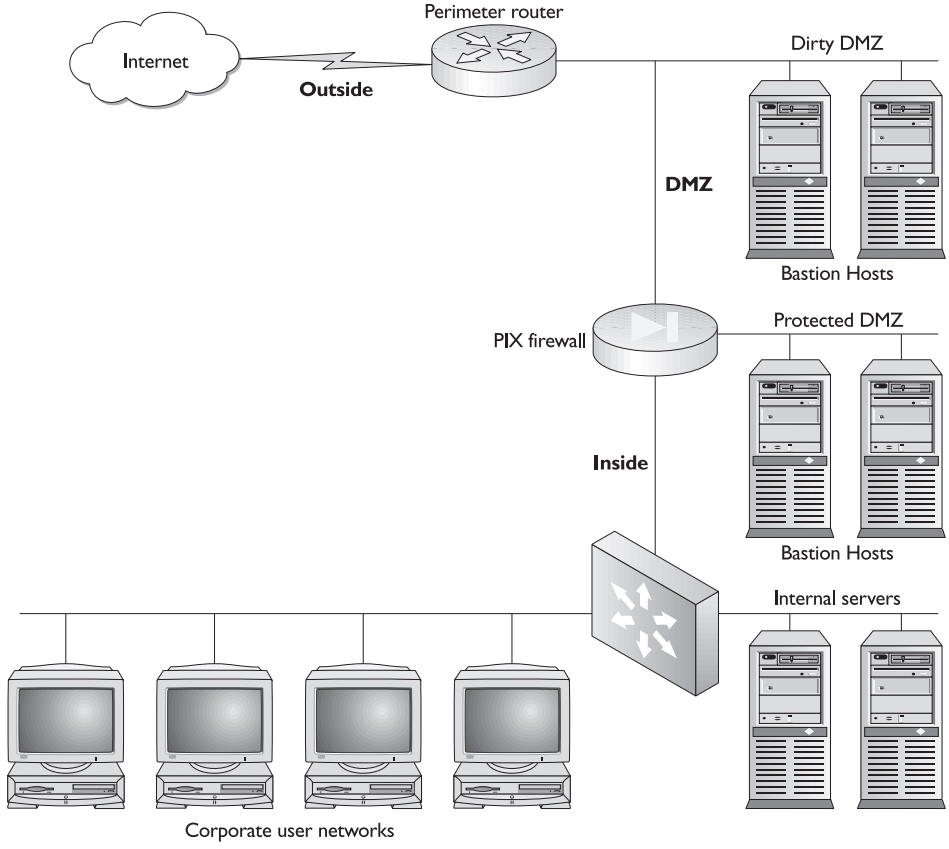

Figure 5-1 Network design with perimeter security

Inside Network

Inside networks are also referred to as the internal or private networks. The *inside area* is made up of the corporate network(s), including all workstations and, typically, any servers not shared with the outside world. These devices are considered trusted and can be accessed freely by other inside hosts. These *trusted* devices need to be protected from the outside world and even from attackers that might have compromised the DMZ area(s). The inside area is under one administrative authority and operates under a common security plan.

The inside area is often connected to one of the firewall interfaces, the *inside interface*. Additional firewalls might be used within the inside area to offer secure separation between two or more subnets. The company security policy might specify a firewall between the accounting/finance departments and the other operating units.

Outside Network

The *outside network* is also referred to as the external or public network. The outside area, or *untrusted area,* is considered to be all devices and networks beyond the direct control of the organization's administration and security policies. The outside area typically includes everything beyond the external interface of the perimeter router, the ISP, and the Internet and all networks attached to it.

The outside network wouldn't normally be attached directly to a firewall device because of media and data framing issues. A firewall device like the Cisco PIX devices have LAN interfaces only requiring another device between the firewall and the serial WAN connection. The perimeter router, possibly a lower-end model, provides the media conversion, framing transition, and any first-line security services.

Demilitarized Zone (DMZ)

The DMZ can be two or more areas inside the network perimeter, but not on the *inside* of the firewall device. The first type of DMZ, often called the *dirty DMZ* or *dirty net,* is the LAN segment between the perimeter router and the firewall. This area has only the protection of the perimeter router and the individual security features of any devices placed there. The second type of DMZ is made up of one or more additional LAN interfaces on the firewall. These areas are often called *protected DMZs* because they have the additional protection offered by the firewall device.

Not uncommonly, some firewall devices offer six or more interfaces, allowing for multiple protected DMZs with different security requirements. Special thought would have to be given to whether any performance benefits from the dirty DMZ only being "filtered" once is offset by the increased risk to whatever is placed out there.

DMZs contain shared server resources, such as web, DNS, and e-mail servers. These servers are available to the outside world. These shared servers are often called bastion hosts, bastion servers, or even sacrificial hosts. *Bastion hosts* must be hardened, and they receive the highest priority security maintenance because of their vulnerability to the outside world and increased likelihood of attacks. A *bastion server* typically runs only those specific services being shared, and all other services will be stopped or turned off.

The dirty DMZ is bordered by the outside interface of the firewall device and the internal interface of the perimeter router. The firewall must be configured to allow loose,

but regulated, access to the protected DMZ from the outside network, while at the same time protecting the inside network. Inside network users need access to the server resources in the DMZ and are typically allowed limited access, possibly restricting access to only those sessions originating within the inside network.

Firewall

A *firewall* is a device that separates or joins the inside network to the dirty DMZ and any optional protected DMZs. The firewall can be a router-running firewall feature set, a specialty server with two or more NICs in different networks, or a specialty device like the Cisco PIX that does nothing but provide firewall services. While suitable applications exist for each type of firewall, generally best is to use a dedicated device performing only security features, and leave routing and serving to other devices.

In a network like the example in Figure 5-1, the firewall would typically be configured to prevent access from the outside to the inside, possibly limiting access to those sessions originating from the inside network. The firewall configuration might allow inside users access to DMZ resources, while providing some defense for the inside from attackers who compromise a bastion host.

Unsolicited access from the outside directed to the inside would typically be blocked. Certain well-thought-out exceptions and configurations could be created, so e-mail server(s) residing on the inside network, instead of the DMZ, could still exchange e-mails. Securing this type of connection is covered in the firewall chapters.

The typical firewall device has two or more LAN interfaces: one each for the inside and outside networks. Optionally, an additional LAN interface can exist for each protected DMZ network. Today, the LAN interfaces are typically Fast Ethernet or Gigabit Ethernet, but there's no reason they couldn't be Ethernet, Token Ring, or FDDI.

Some small firewalls used in implementations like branch locations or telecommuter residences could only have two interfaces for separating the inside network from the outside world. In those small implementations, the inside interface could connect to a user machine via a crossover cable, or to a small hub or switch. The external interface would often connect to the DSL, cable modem, or ISDN device. The Cisco 806 router, shown in Figure 5-2, with an Ethernet interface, four-port hub, Cisco IOS, and supporting the firewall feature set, is an example.

While a firewall is normally used to separate the inside network from the outside world, also possible is to use a firewall to separate internal departments where additional

Figure 5-2 Cisco 806 router for telecommuter or small office firewalls

security is required. For example, a school might choose to place a firewall between the student network and the faculty network. In this case, the firewall might have only two interfaces, with the *inside* interface connected to the protected network and the *outside* interface connected to the network perceived as the potential threat.

Perimeter Router

The *perimeter router* is typically a standard router providing a serial connection to the outside world and a LAN connection to the internal network. The perimeter router should provide any filtering of outside traffic to implement basic security for the dirty DMZ and preliminary filtering for the inside network. This device could be running the firewall feature set for additional security options.

Because the perimeter router is often connected to a slower WAN interface on one side and it doesn't normally provide routing functions for internal networks, the LAN interface speed isn't as critical as making sure adequate memory and features exist to handle the outside connection. Even if the inside network is 100MB and all protected DMZ interfaces are full-duplex 100MB, if the Internet connection is a T1 (1.54MB), then a 10MB LAN interface on the perimeter router shouldn't impede traffic. Even most DSL or cable connections would be well below 10MB.

While bandwidth issues are important, feature sets are important on perimeter routers. Routers clear down to the 800 series support access list, firewall features, and so forth, making low-end devices attractive in some perimeter implementations. If intrusion detection features are needed, though, you should know that the firewall feature sets for devices below the 2600 devices don't include them. So, while a 1700 or 2500 device might handle the traffic, it won't provide intrusion detection services.

The next chapter looks at the Cisco IOS firewall feature set and the additional features it can add to the perimeter router.

Eavesdropping

Eavesdropping, using packet-sniffing tools, allows attackers to read transmitted information, including logon information and database contents. Eavesdropping can yield information, such as credit card information, which can be used in later transactions, damaging both the customer and the business that allowed the breach.

A related problem involves eavesdropping and session replay in which the attacker retransmits the captured data, such as a session logon sequence to gain access to network resources.

Router Solutions

If you're using a dynamic routing protocol that supports authentication, a good idea is to enable that authentication. This prevents some malicious attacks on the routing infrastructure and it can also help to prevent damage caused by misconfigured "rogue" network devices.

Configuring Authentication Example

Many routing protocols support authentication and the commands vary somewhat, but the following example demonstrates the feature using RIP version 2.

RIP version 1 doesn't support authentication, but RIP version 2 allows RIP authentication on a per-interface basis. RIP supports two modes of authentication: plaintext authentication and MD5 authentication. The default is plaintext authentication. Don't use plaintext authentication for security purposes because the unencrypted authentication key is sent in every RIP version 2 packet. Interfaces on both sides of the link must be configured for MD5 authentication using the same key number and key string.

The key chain determines the set of keys that can be used on the interface. If a key chain isn't configured, no authentication is performed on that interface, not even the default authentication.

The following commands configure RIP authentication. The first three are configured in Global Configuration mode and the last two are in Interface Configuration mode.

Command	Purpose
key chain *name*	Names a key chain. Could contain more than one key for added security.
key *number*	Defines the first key in the key chain.
key-string *string*	Defines the key value—must be identical on both ends of the link.
ip rip authentication key-chain *name*	Enables RIP authentication.
ip rip authentication mode {text | md5}	Uses MD5 digest authentication (or defaults to plain text authentication).

```
Rtr1:
key chain seattle
key 1
key-string 9631
!
interface Ethernet 0
  ip address 192.168.1.1 255.255.255.0
!
interface Serial 0
  ip address 192.168.2.1 255.255.255.252
  ip rip authentication mode md5
  ip rip authentication key-chain seattle
!
router rip
  version 2
  network 192.168.1.0
  network 192.168.2.0

Rtr2:
key chain seattle
key 1
key-string 9631
!
interface Ethernet 0
```

```
  ip address 192.168.3.1 255.255.255.0
!
interface Serial0
 ip address 192.168.2.2 255.255.255.252
 ip rip authentication mode md5
 ip rip authentication key-chain seattle
 clockrate 64000
!
router rip
 version 2
 network 192.168.3.0
 network 192.168.2.0
```


Verifying MD5 Authentication

By configuring the Rtr1 and Rtr2 as shown, all routing update exchanges are authenticated before being accepted. Verify this by observing the output obtained from the **debug ip rip** command:

```
Rtr1#debug ip rip
RIP: received packet with MD5 authentication
RIP: received v2 update from 192.168.2.2 on Serial0
     192.168.3.1/24 via 0.0.0.0 in 1 hops
```

MD5 authentication uses the one-way, MD5 hash algorithm in which the routing update doesn't carry the actual password (key) for the purpose of authentication. A128-bit message, generated by running the MD5 algorithm on the password, and the message are sent along for authentication. The recipient device runs the same hash using the key information configured and, assuming it generates the same 128-bit message, the update will be accepted. MD5 hashing is covered in greater detail in Chapter 8.

Encryption and Tunneling

Data separation using tunneling technologies, such as generic routing encapsulation (GRE) or Layer 2 Tunneling Protocol (L2TP), provides effective data privacy. When privacy requirements call for the use of digital encryption technology, protocols such as IPSec provide the added protection when implementing VPNs. The IPSec and VPN chapters, 9 through 16 and 21, cover these features in detail.

Hub and Switch Issues

While less-common than a few years ago, the shared media technology of hubs made them susceptible to sniffing attacks. All an intruder needed was access to any port in the collision domain and they could monitor all traffic. Just changing to switches reduced the threat because the only traffic visible on the port is traffic directed to the attached host or broadcasts.

Many switches support *port monitoring* or *port spanning*, which allows traffic from ports designated for capture to the port designated as the monitor port. The sniffer is attached to this port. Care should be taken to prevent port monitoring.

Limit Unneeded TCP/IP and Other Services

As a general rule, any unnecessary service should be disabled on perimeter routers. The following services are often useful, but should be disabled if they aren't being used.

TCP and UDP "Small Services"

By default, Cisco devices offer the small services: echo, chargen, and discard. These services, especially their UDP versions, are infrequently used for legitimate purposes, but can be used to launch DoS and other attacks, which would otherwise be prevented by packet filtering.

The small services are disabled by default in Cisco IOS 12.0 and later software. In earlier software, they can be disabled using the **no service tcp-small-servers** and **no service udp-small-servers** commands.

Finger

Cisco routers support the "finger" service used to identify which users are logged into a network device. This information can be useful to an attacker. The "finger" service can be disabled with the **no service finger** command.

NTP

The Network Time Protocol (NTP) isn't a particularly dangerous feature. If NTP is being used, be sure to configure a trusted time source and use proper authentication. Corrupting the time on network devices can subvert certain security protocols and cause some processes to fail to synchronize or function. If NTP isn't being used on a particular router interface, it can be disabled with the interface command **no ntp enable**.

CDP

Cisco Discovery Protocol (CDP) is a fairly useful feature, but on the network perimeter, it can be dangerous because it announces the following to any system on a directly connected segment: that the router is a Cisco device, the model number, and the Cisco IOS version being run. This information could be used to exploit vulnerabilities in the router. The CDP protocol can be disabled with the global configuration command **no cdp running**. CDP can be disabled on a particular interface with the **no cdp enable** command.

Denial of Service Attacks

The following Cisco features can be used to increase the basic security measures related to the way in which the router forwards IP packets.

Controlling Directed Broadcasts

An *IP directed broadcast* is a datagram sent to the subnet broadcast address. The directed broadcast is routed through the network as a unicast packet until it arrives at the target subnet, where it's converted into a link-layer broadcast. Only the last router, the one directly connected to the target subnet, can positively identify a directed broadcast.

IP directed broadcasts are used in *smurf* DoS attacks, in which the attacker sends ICMP echo requests from a falsified source address to a directed broadcast address. This causes all the hosts on the target subnet to send replies to the falsified source. By sending a continuous stream of these requests, the attacker can create a much larger stream of replies, burying the smurfed host and their link to their ISP.

The **no ip directed-broadcast** command on an interface causes discards directed broadcasts, such as 192.168.12.255, that would otherwise "explode" into link-layer broadcasts at that interface. The **no ip directed-broadcast** command is the default in Cisco IOS software version 12.0 and later.

Flood Management

As you saw in Chapter 1, many DoS attacks rely on floods of useless packets that congest network links, slow hosts, and overloaded routers. Being aware of where performance bottlenecks lie is important in flood management. If a DoS flood is burying a T1 line, then filtering the flood at the source end router can help, while filtering at the destination end will have little or no effect.

If an "underpowered" router is the bottleneck, then adding additional filtering will probably make things worse. In this case, increasing memory or replacing the device might have to be part of the solution.

Transit Floods

In some cases, Cisco's quality of service (QoS) features can be used against some kinds of floods on serial links. Using weighted fair queuing (WFQ), the default for low-speed serial lines in recent versions of Cisco IOS software, has proven effective against ping floods, but less effective against SYN floods. A *ping flood* appears to WFQ as a single traffic flow, whereas each packet in a SYN flood generally appears as a separate flow. A smurf reply stream falls somewhere between the two. Cisco QoS features are covered extensively on Cisco's web site.

TCP Intercept

The *TCP Intercept* feature is designed specifically to reduce the impact of SYN flooding attacks on hosts. TCP Intercept is available in some IOS versions for many routers with model numbers of 4000 or greater. A device supporting TCP Intercept can literally step in as a proxy and handle TCP session requests for a server that is under attack or heavy load. The device attempts to complete the TCP 3-way handshakes, forwarding successful attempts to the server and discarding the rest.

Antispoofing with RPF Checks

Cisco IOS versions that support Cisco Express Forwarding (CEF) can have the router check the source address of any packet against the interface through which the packet entered the router. If, according to the routing table, the input interface isn't a feasible path to the source address, the packet is then dropped. The feature is known as a *reverse path forwarding (RPF)* check and is enabled with the command **ip verify unicast rpf**.

Unauthorized Access

The perimeter router is the first line of defense against external threats. Features like using access lists to perform traffic filtering, covered in Chapter 2 and Appendix A, can help in the battle. In this section, you learn about the Dynamic and Reflexive access lists.

Address Filtering

RFC 2827 filtering involves filtering out any IP addresses from coming into a network segment, which should already be on that segment. If the entire 195.17.1.0 network is attached to a router interface, then no legitimate packets with source addresses in that network should be coming in through the router. This should be applied to perimeter routers for sure, but it can be used on internal routers to prevent spoofing within the network. Similarly, limiting any outbound packets, leaving the network to ones that have source addresses assigned to that network, can prevent a network's hosts from spoofing other networks. This could be the result of an attacker on the inside or a DoS bot on a local host participating in an attack on an outside host. If the company can get its ISP to perform RFC 2827 filtering on packets coming into the network, it would preserve the bandwidth of the link and kill some hack attempts.

RFC 1918 filtering involves filtering out RFC-defined private addresses from entering or exiting the network segment. Because they have no business on the Internet anyway, they shouldn't be there. If private addresses are used in the network, RFC 2827 filtering will include them. The following is a standard ACL approach, although an extended or named list would also work.

```
interface serial 0
ip access-group 10 in
access-list 10 deny 10.0.0.0 0.255.255.255
access-list 10 deny 172.16.0.0 0.15.255.255
access-list 10 deny 192.168.0.0 0.0.255.255
access-list 10 permit any
```

Dynamic (Lock-and-Key) Access Lists

Assume the network is secured with ACLs regulating traffic in, out, and through the network. Decisions were made to be conservative in allowing traffic to access the resources. So everything that can be blocked is being dealt with. What if you need some flexibility

to deal with some necessary exceptions? *Dynamic access lists,* often called *lock-and-key access lists,* can literally create temporary openings to your network for specific IP traffic.

Figure 5-3 shows an example of a company network that might have limited TCP access to the internal networks to those sessions that originated from within the LAN. To accomplish this, they might have used the TCP ACL established option. The organization still wants to allow the network administrators access from the outside to reduce the number of evening and weekend callouts, particularly for forgotten passwords, locked accounts, and so forth.

The point of lock-and-key is to grant temporary IP access to specific hosts that would normally be blocked. The process works like this:

1. A user Telnets to the router configured with the dynamic ACL. The router can be a perimeter router protecting the entire network or an internal router protecting certain segments.
2. The router challenges the user to authenticate. The authentication method used is whatever security has been applied to Telnet sessions (*line vty* settings) in the router configuration. Options include using standard passwords, local user name/password entries, or AAA.
3. Once successfully authenticated, the Telnet session is terminated and the router creates a temporary ACL that allows traffic between the specific host(s) and resources defined in the ACL statements.
4. The user(s) can then have temporary access through the router.
5. When the time limit defined in the ACL is reached, the temporary ACL list is removed and a new authentication is required to continue.

Properly configured, dynamic access lists provide the same benefits as standard and static extended access lists with the additional security benefit of authenticating users, thereby reducing the opportunity for network break-ins by hackers.

Creating a Lock-and-Key System

Creating a dynamic ACL is a two-step process that includes building or modifying an ACL that will be applied to a router interface, and then configuring the virtual terminal access to support the feature.

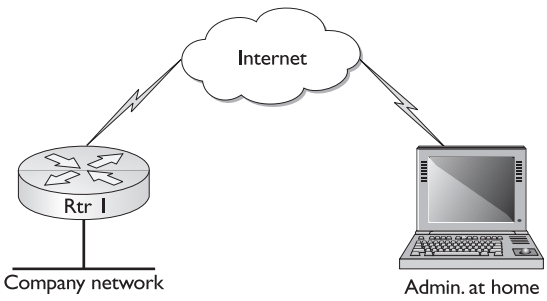

Figure 5-3 Simple example of lock-and-key access

The Dynamic ACL Statements The dynamic access list feature is an option for both numbered and named extended access lists for TCP, UDP, ICMP, and IGMP traffic. The [dynamic *dynamic-name* [timeout *minutes*]] options precedes the {permit | deny} choice.

The basic syntax for a numbered ACL would be the following:

Router(config)#access-list *acl#* **[dynamic *dynamic-name* [timeout *minutes*]]** {permit | deny} *protocol* | *protocol-keyword*}{*source wildcard* | any} [*operator source-port*] {*destination wildcard* | any} [*operator destination-port*] [precedence *precedence*] [tos *tos*] [log | log-input] [options]

The basic syntax for a named ACL would be as follows:

Router(config)#ip access-list extended *name*
Router(config-ext-nacl)#**[dynamic *dynamic-name* [timeout *minutes*]]** {permit | deny}{*protocol* | *protocol-keyword*}{*source wildcard* | any} [*operator source-port*] {*destination wildcard* | any} [*operator destination-port*] [precedence *precedence*] [tos *tos*] [log | log-input] [options]

dynamic	Identifies this access list as a dynamic access list, meaning it can create a temporary ACL.
dynamic-name	The name of the temporary ACL to be created.
timeout *minutes*	Specifies the absolute length of time (in minutes) a temporary ACL can exist. **Warning:** The default is infinite, which allows an entry to remain permanently until an administrator deletes it. Always set a timeout to close this "hole" in your security.

The following simple example creates a named access list with a dynamic access list statement. Line two allows any host to telnet to the router (IP address 201.5.45.87). Line three enables a dynamic list entry called **allow-in**. The statement allows router access to any IP traffic from the authenticated host or the host network. It also limits the temporary access to 60 minutes, regardless of activity.

```
Rtr1(config)#ip access-list extended filter-in
Rtr1(config-ext-nacl)#permit tcp any host 201.5.45.87 eq telnet
Rtr1(config-ext-nacl)#dynamic allow-in timeout 60 permit ip any any
```

Because the dynamic ACL will be applied to an interface and because only be one ACL can exist in each direction per protocol, the dynamic statements are typically going to be added to an existing ACL that's already regulating traffic through that interface.

The Virtual Terminal Statements The [dynamic *dynamic-name* [timeout *minutes*]] options precede the permit | deny choice.

Router(config-line)#autocommand access-enable [host] [timeout *minutes*]

autocommand	Executes the **access-enable** command when a user authenticates through a Telnet session into the router.
access-enable	Enables the lock-and-key access feature by enabling the router to create a temporary access list entry in a dynamic access list.
host	Limits access to the host that originated the Telnet. If not included, the temporary ACL allows access by all hosts on the network defined by the ACL **dynamic** statement network mask.
timeout	Specifies an **idle timer** (in minutes) for the temporary ACL entry. If the temporary ACL statement isn't accessed within this time period, the statement is deleted and requires the user to authenticate again. The default is for the entries to remain permanently.

The following example creates a single password—cisco—and then enables the dynamic ACL to create a temporary ACL limited to the host that authenticated. If five minutes pass between exchanges, the temporary entry is deleted.

This example could be one solution to the problem in Figure 5-3 to allow an administrator to access the network from home.

```
Rtr1(config)#username itadmin password cisco
Rtr1(config)#ip access-list extended filter-in
Rtr1(config-ext-nacl)#permit tcp any host 199.45.5.7 eq telnet
Rtr1(config-ext-nacl)#permit tcp any any established
Rtr1(config-ext-nacl)#permit icmp any any
Rtr1(config-ext-nacl)#dynamic allow-in timeout 60 permit ip any any log
Rtr1(config-ext-nacl)#int s0/0
Rtr1(config-if)#ip address 199.45.5.7 255.255.255.0
Rtr1(config-if)#ip access-group filter-in in
Rtr1(config-if)#line vty 0 4
Rtr1(config-line)#login local
Rtr1(config-line)#autocommand access-enable host timeout 5
Rtr1(config-line)#^Z
Rtr1#
```

Line 1 defines an acceptable user name/password combination.

Line 11 stipulates that the local database of user names and passwords will be used to authenticate. If multiple user name/password combinations had been configured, any set would be accepted.

Line 3, permit tcp any host 199.45.5.7 eq telnet, allows anyone to access the lock-and-key router. The "any" could be replaced by a host address, permit tcp host 24.12.117.91 host 199.45.5.7 eq telnet, if the admin has a permanent IP address. This reduces the opportunities for outsiders to access the network. Even if it can't be reduced to a single IP address, a network address and wildcard mask could limit exposure.

Line 6, dynamic allow-in timeout 60 permit ip any any, specifies the temporary ACL will be called allow-in, limits the temporary ACL statements to 60 minutes, and specifies the temporary ACL statements will allow any host access to any network for IP traffic. If the host option is left off line 12, the temporary ACL will be open to any host in the world. By replacing the first "any" with a network address and a wildcard mask, the exposure could be limited. Replacing the last "any" with a host address or a network address and a wildcard mask, all authenticated access would be limited to the defined

address(es). An example of changes to line six could be the following: dynamic allow-in timeout 60 permit ip 24.12.117.0 0.0.0.255 192.168.0.0 0.0.0.255.

Line 12 limits access to only the host address that authenticated and sets the idle timer to five minutes.

Line 10, line vty 0 4, specifies that all five virtual terminal sessions will be configured for lock-and-key. But, what if you have to Telnet into that router for administration purposes? The router is going to authenticate us, and then close the Telnet session just like everyone else. The following lines show an alternative that configures virtual lines 0 through 3 (4 total) to lock-and-key while saving the final session, line vty 4, for separate configuration.

```
!
line vty 0 3
  login local
  autocommand access-enable host timeout 10
line vty 4
  password cisco
  login
  rotary 1
```

The rotary Command Because normal virtual sessions are assigned in round-robin fashion, knowing which session is being accessed is impossible. The **rotary 1** command, shown in the example, makes it possible to specify a port number (3001) when setting up the Telnet session, and thereby specify the virtual session. The numbers in the 3000 range are reserved for this purpose, so adding 3000 to a rotary value used will access the correct line. The Telnet 199.45.5.7 3001 command would access the vty session 4. If the rotary statement had been rotary 17, the Telnet command would have used 3017.

Display Dynamic Access Lists

Displaying a temporary access list is much like any other access list using the **show access-list [acl# | acl-name]** or **show ip access-list [acl# | acl-name]** commands from Privilege Exec mode. The trick is this: the temporary access-list lines are only present when they're in use. After the absolute or idle timeout parameter has activated, the temporary entries are cleared. The number of matches displayed indicates the level of activity for that statement since the last time the counters were cleared.

```
Rtr1#sho access-lists
Extended IP access list filter-in
    permit tcp any host 199.45.5.7 eq telnet (66 matches)
    permit tcp any any established (232 matches)
    permit udp any any eq rip (44 matches)
    Dynamic allow-in permit ip any any log
      permit ip host 192.168.0.14 any log (21 matches) (time left 253)
Rtr1#
00:36:10: %SEC-6-IPACCESSLOGP: list filter-in permitted tcp 192.168.0.14(1107) -
> 192.168.2.1(2001), 2 packets
Rtr1#
```

Line seven identifies the host that authenticated and specifies that 21 packets were permitted. The time left is 253 seconds. Each time another match is made, the idle timer is reset to 300, the five minutes specified in the **autocommand access-enable host timeout 5** command.

Reflexive Access Lists

As you learned in Chapter 3, the TCP Established option could be useful in limiting access to the local networks to those sessions that originated from inside the network. Any packets originating from the outside that are trying to establish a connection would be rebuffed. But what about other protocols like UDP and ICMP that aren't connection oriented? In fact, a similar problem occurs with TCP applications that send on one port and receive on another. Either way, the established feature doesn't always work in these cases.

A feature of IP-named extended access lists, called *reflexive access lists*, provides what is referred to as *reflective filtering*. Any IP packet going out of a filtered port will create a temporary access list statement inbound, which is a mirror image of the outgoing packet source and destination information. This temporary opening remains active until a TCP FIN packet is received or the idle timer expires. The idle timer timeout feature is critical for protocols that don't have a session-ending message, such as UDP and ICMP.

If host 192.168.0.10 wants to use TCP port 1045 to establish a telnet session with 192.168.1.2 in another network through its TCP 23 (telnet) port, the access list statement to allow this—if one had to be created—would look like the following:

```
permit tcp host 195.168.0.10 eq 1045 host 192.168.1.2 eq 23
```

What gets created is a temporary reflected inbound statement, which looks like the following:

```
permit tcp host 192.168.1.2 eq 23 host 195.168.0.10 eq 1045
```

Reflexive access lists can filter network traffic based on IP upper-layer protocol information by using port numbers. They create temporary ACL statements that permit returning IP traffic for sessions originating from within the network, but prevent other IP traffic from outside. They use reflexive filtering that can only be defined in *named extended IP access lists*.

Reflexive access lists can be an important part of securing your network against spoofing and certain DoS attacks. They can be a major component in a router firewall defense strategy. Reflexive access lists are relatively simple to use and provide greater control over which packets enter your network than do basic access lists.

Selecting the Interface to Use

Reflexive access lists are generally applied to a perimeter router. You can apply reflexive ACLs to either the external or the internal interface. Whenever possible, the reflexive ACL should be configured on the outside interface to block unwanted IP traffic.

Figure 5-4 shows two topologies that might be found on the perimeter of the company network. The topology on the left is a basic perimeter router. In this case, the reflexive ACL can be applied to the external interface (serial) and keep all blocked traffic out of the router.

The topology on the right shows a single-device firewall design for supporting shared servers in a DMZ network. In this case, putting the reflexive ACL on the external interface would undoubtedly interfere with legitimate traffic going to the DMZ area. Placing

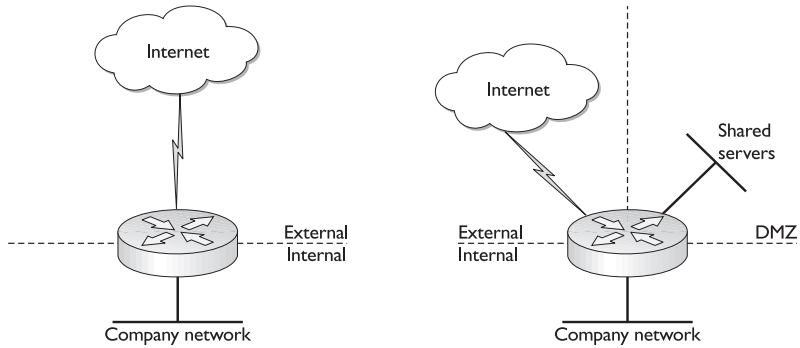

Figure 5-4 Topology impacts on interface selection

the ACL on the internal interface can protect the local network, while not interfering with traffic to and from the shared resources in the DMZ.

Configure Reflexive Access Lists

Reflexive access lists aren't applied directly to an interface but, instead, are nested into an extended-named IP access list that's applied to an interface. Because of the concept of nesting, reflexive access lists don't have the usual implicit deny-all at the end of the list.

Two primary commands need to be configured in building reflexive access lists: the **reflect** command and the **evaluate** command. In addition, the **ip reflexive-list timeout** command is for setting the idle timer.

The reflect Statement The keyword **reflect** in a **permit** statement is the key to creating the return entries in the temporary access list. The **reflect** statement in an outbound ACL creates the entries in the inbound ACL when packets matching the condition(s) are recognized.

The ACL that contains the **reflect** statement can also contain regular **permit** and **deny** statements, either before or after the **reflect** statement. If a match occurs, the packet will be processed but won't trigger a temporary entry in the reflexive access list.

The basic syntax is

Rtr(config)#ip access-list extended *name*
Rtr(config-ext-nacl)#permit *protocol source destination* reflect *name* [timeout *seconds*]

protocol	Any TCP/IP protocol that can use named extended ACLs
source and destination	Same as any ACL source and destination, including keywords **host** and **any**, wildcard masks, and port numbers
name	Name of temporary reflexive ACL to be created
timeout	Optional idle timer for this line only that overrides the absolute timeout set by the **ip reflexive-list timeout** global configuration command (default is 300 seconds)

Examples follow the discussion of the evaluate statement in the next section.

The evaluate Statement The **evaluate** statement is the final configuration entry in another named extended IP access list that partners with the **permit/reflect** statements previously discussed. The syntax is as follows:

Router(config-ext-nacl)#evaluate *name*

name	Name of temporary reflexive ACL specified in the **permit/reflect** statements

The following is a simple example where outgoing TFTP and HTTP traffic is allowed and generates return openings in a temporary ACL called newlist. The Telnet entries in both ACLs are only an example of normal ACL statements that don't generate entries in newlist.

```
Internet serial 0
  ip access-group ok-out out
  ip access-group ok-in in
!
ip access-list extended ok-out
  permit tcp any any eq telnet
  permit udp any any eq tftp reflect newlist
  permit tcp any any eq www  reflect newlist
!
ip access-lst extended ok-in
  permit tcp any any eq telnet
  evaluate newlist
```

Two named access lists, ok-out and ok-in, are applied to the external serial interface.

ok-out contains two **reflect permit** statements, plus a normal Telnet statement. Outbound TFTP and HTTP packets will create entries in a temporary access list named newlist, which reverses the source and destination designations of the outbound packets.

The ok-in evaluate newlist statement incorporates the newlist statements.

The ip Reflexive-list timeout (Seconds) Command The global mode command **ip reflexive-list timeout** ***seconds*** specifies the length of time reflexive access list statements remain in the ACL if no packets in that session are detected. Each new matching packet rests the timeout timer. If no timeout statement is specified, the default timeout value is 300 seconds. The next example shows changing the idle timeout to two minutes.

External Interface Example The following is a simple example of a reflexive ACL applied to an external interface allowing the LAN (192.168.0.0) users access to the Internet and the DNS servers.

```
ip reflexive-list timeout 120
!
Interface serial 0
  ip access-group lan-out out
  ip access-group lan-in in
!
ip access-list extended lan-out
```

```
  deny icmp any any
  permit tcp any any eq dns reflect ok-packets timeout 240
  permit tcp any any eq www reflect ok-packets
!
ip access-list extended lan-in
  deny icmp any any
  evaluate ok-packets
```

The first line changes the global timeout value to two minutes for all reflexive access lists entries on this router that don't specify their own timeout value. Line nine sets a four-minute idle timer for the DNS packets only.

Internal Interface The following is a simple example of a reflexive ACL applied to an internal interface allowing the LAN (192.168.0.0) users access to the Internet and the DNS servers. Typically, this implementation would be used if there's another network (DMZ) of shared servers. Placing the ACL on the external interface would interfere with the outside access to the servers.

```
Interface ethernet 1
  ip access-group lan-in in
  ip access-group lan-out out
!
ip access-list extended lan-in
  deny icmp any any
  permit tcp any any eq dns reflect ok-packets timeout 120
  permit tcp any any eq www reflect ok-packets timeout 180
!
ip access-list extended lan-out
  deny icmp any any
  evaluate ok-packets
```

The two **permit/reflect** entries set their own timeout, two and three minutes, respectively. Without these timeout settings, they would be governed by the global default (300 seconds).

Viewing Reflexive Access Lists

Reflexive access lists can be displayed using the **show access-list** command. No reflexive access list entries will appear if no outbound traffic has triggered the reflexive access list. After a HTTP session is initiated from within the network, the **show access-list** command could display the following:

```
Extended IP access list lan-in
  permit icmp any 195.168.1.0 0.0.0.255 echo-response
  permit tcp any host 195.168.1.20 eq 80 (7 matches)
  permit bgp any any (4 matches)
  evaluate ok-packets
Extended IP access list lan-out
  permit tcp 195.168.1.0 0.0.0.255 any eq 80 reflect ok-packets timeout 180
  permit tcp 195.168.1.0 0.0.0.255 any eq 53 reflect ok-packets
Reflexive IP access list ok-packets
  permit tcp host 201.37.14.127 eq www host 195.168.52 eq 80 (5 matches)
    (time left 140 seconds)
```

Reflexive Access Lists Caveats Some applications, like FTP, allow port numbers to be changed during a session, meaning the port numbers of returning packets aren't the same as the originating packet. Reflexive access lists can't resolve this and the return packet will be denied—even if it's part of the same session. Passive FTP is an option that will work with reflexive access lists if the FTP site supports it.

Reflexive access lists can't examine data in the packet beyond the Layer 4 (OSI model) information, such as TCP and UDP port numbers and the related IP addresses. They can't follow the application as you just saw with the FTP discussion in the preceding paragraph.

Lack of Legal IP Addresses

Organizations are facing two related challenges. First, the depletion of registered IP addresses means that getting blocks of "real" addresses larger than class C is virtually impossible. Second, trying to scale larger organizations with a hodge-podge of class C addresses will reduce routing efficiency. The following private address pools (RFC 1918) provide the only logical solution, but they can't go out on the Internet.

10.0.0.0–10.255.255.255	1 class A	16.6+ million hosts
172.16.0.0–172.31.255.255	16 class Bs	1+ million hosts
192.168.0.0–192.168.255.255	256 class Cs	65,000+ hosts

Network Address Translation (NAT) is a mechanism that allows private addresses to be translated to real addresses, so they can travel through the Internet. NAT allows an organization with unregistered "private" addresses to connect to the Internet by translating those addresses into globally registered IP addresses. Incoming traffic is translated back for delivery within the inside network.

NAT can save an organization the hassle of readdressing its network when it changes ISPs. The real addresses leased from the original ISP can continue to be used, but must be translated at the perimeter to addresses that will summarize to the new ISP. This can be a real sanity saver when an ISP fails without notice to clients.

NAT can also provide a limited level of network privacy by hiding internal IP addresses from external networks. The external hosts will see the assigned "real" address and respond to that address.

NOTE Important to know is that some applications, such as some e-mail programs, capture the internal address and store it in the data portion of the IP packet, which means that under some circumstances the internal address can be seen by the outside world. Programs like McAfee SpamKiller display the entire path back to the source, including the internal local addresses.

This chapter covers Cisco IOS NAT and how to configure it, typically on a perimeter router. You must understand NAT technology and concepts because these are revisited in the Firewall chapters (6–8 and 17–20) and IPSec chapters (9–16 and 21) and can be on any of the exams. In some form, NAT is available on personal routers (cable and DSL con-

nections), firewall devices like the PIX devices, and proxy servers working as firewalls. This section covers basic NAT operations and the following NAT implementations:

- Static NAT
- Dynamic NAT
- Dynamic NAT with overloading (PAT)

NAT Technology and Terminology

NAT is the process of altering the IP header of a packet, so the source local address of the internal host is replaced in the header by real global addresses. In some cases, the destination address might also be modified. This swapping process is performed by a NAT device, usually on the network perimeter. The NAT server then maintains a table of the translations that allows returning packets to be addressed with the correct internal address. Figure 5-5 is a simple example of NAT where the local hosts (10.0.0.0 private network) are translated by the NAT device to global addresses in the 192.168.1.0 network.

When an outside host sends a response to 192.168.1.97 or 192.168.1.98, the NAT router checks the current table of network address translations, and then replaces the destination address with the original inside source address.

NOTE Obviously, 192.168.1.0 isn't a globally routable "real" network, but we'll use only private addresses in all examples for two reasons. The public addresses belong to someone, and the current or future owners might not agree with their use. Second, inevitably, someone will build a "practice" lab like the ones in the book and a slim possibility would exist for impacting the real owners of the address.

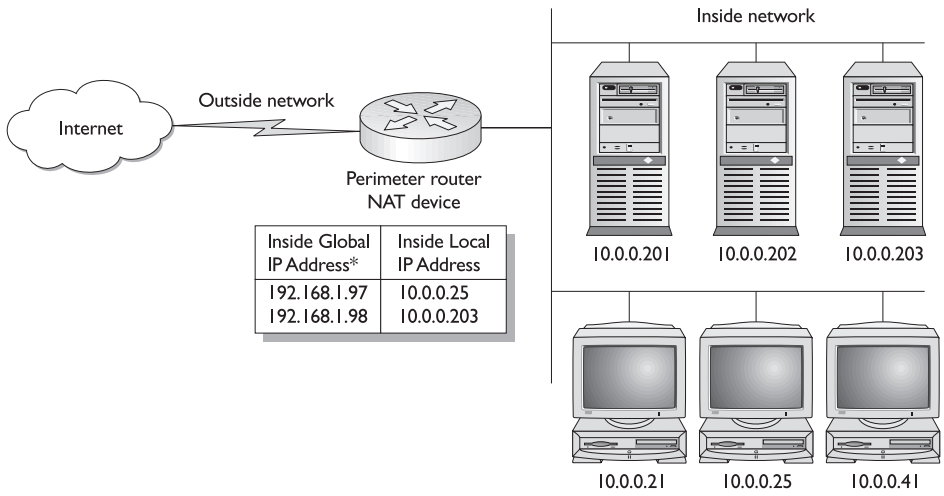

Figure 5-5 NAT translation for inside network 10.0.0.0

Cisco IOS Software uses the following terms when working with NAT. While other implementations might change the terms, the concepts remain the same:

- **Inside local address** The configured IP address assigned to a host on the inside network. Quite often, these addresses are drawn from the RFC 1918 private address pools. They could also be real addresses officially assigned to some other organization.
- **Inside global address** The inside global address is the translated address. This is the IP address the outside world sees for an inside host. Typically, these addresses are allocated from a pool of real IP addresses provided by the ISP.
- **Outside local address** The IP address of an outside host as it appears to the inside network. Because these addresses are only used on the inside network, they might not necessarily be real public addresses. They can be assigned from the RFC 1918 private address pools.
- **Outside global address** The configured IP address assigned to a host in the outside network. This address is a real address assigned from the globally routable pool.

NAT translations can occur dynamically or statically, and can be used for a variety of purposes, as described in the following sections.

Device Interfaces

A key concept in all NAT translations is identifying the internal (inside) interface(s) and the external (outside) interface(s). NAT needs to know which interfaces are connected to inside networks and which are connected to outside networks. Only packets moving between inside and outside interfaces can be translated. At least one inside interface and one outside interface for each border router must be defined. The following configuration additions would be typical for a perimeter router with a single serial connection to the Internet and a single LAN:

```
interface Serial0/0
  ip nat outside
!
interface FastEthernet0/0
  ip nat inside
```

Outside interfaces should never be included in internal route tables and must remain unknown to internal hosts. Similarly, internal addresses should never be advertised or shared with the outside world. The internal interface typically is either the default gateway for hosts that share that network or the next-hop address for default routes set on other internal routers.

Static NAT

Static translation involves permanently mapping an inside local address to an inside global address. For most organizations, this would be reserved for those network servers

shared with the outside world that, therefore, need to keep the same address. This is particularly true for servers, such as web or e-mail servers, that will have domain name translations in DNS servers. Any change in their global address results in missed traffic until the DNS server system is fully resolved to the change. This might take the form of those dreaded 404 Page Not Found Errors when searching the Web.

Except in the smallest of networks, or those with few users authorized to go out into the Internet, static NAT isn't a suitable implementation for users because it requires a one-for-one ratio of real addresses to inside addresses.

Configuring Static NAT

To enable static NAT of an inside source address; use the **ip nat inside source** command in global configuration to create a single static translation. To remove the static translation, use the no form of this command. The syntax is

Rtr1(config)#ip nat inside source static *local-ip global-ip*
Rtr1(config)#no ip nat inside source static *local-ip global-ip*

local-ip	Local IP address assigned to an inside host—often taken from RFC 1918 pool
global-ip	Globally unique IP address of an inside host as it appears to the outside world

Creating three static translations for servers in the example at the beginning of this section might look like the following:

```
Rtr1(config)#ip nat inside source static 10.0.0.201 192.168.1.105
Rtr1(config)#ip nat inside source static 10.0.0.202 192.168.1.106
Rtr1(config)#ip nat inside source static 10.0.0.203 192.168.1.106
```

Confirming NAT Translations

To see the NAT translations, use the **show ip nat translations** command. The results would look something like the following output. The translations appear immediately because they're permanent. An outside user could access these servers by using the 192.168.1.105–107 addresses, assuming no access lists prevent it.

```
Rtr1#show ip nat translations
Pro Inside global    Inside local    Outside local   Outside global
---  192.168.1.105    10.0.0.201      ---             ---
---  192.168.1.106    10.0.0.202      ---             ---
---  192.168.1.107    10.0.0.203      ---             ---
Rtr1#
```

To see any NAT activity, use the **show ip nat statistics** command. The results would look something like the following output. The Total active translations summarizes the current status, while hits and misses indicates no traffic has been attempted. If an outside host were to ping one of the servers, the hits would show the number of packets—four or five.

```
Rtr1#show ip nat statistics
Total active translations: 3 (3 static, 0 dynamic; 0 extended)
```

```
Outside interfaces:
  Serial0/0
Inside interfaces:
  FastEthernet0/0
Hits: 0  Misses: 0
Expired translations: 0
Dynamic mappings:
Rtr1#
```


Dynamic NAT

You can configure dynamic NAT in the following three steps:

- Define a pool of global IP addresses to be allocated.
- Use a standard ACL to define the local addresses eligible for translation.
- Link the pool of global addresses with the eligible local addresses.

Define a Pool of Global IP Addresses to Be Allocated

To define a pool of global IP addresses for NAT, use the **ip nat pool global** configuration command. To remove one or more addresses from the pool, use the no form of this command. The command defines a pool of addresses using start address, end address, and either netmask or prefix length. The **pool** command can define either an inside global pool or an outside local pool. The syntax is

Rtr1(config)#ip nat pool *name start-ip end-ip* {netmask *netmask* | prefix-length *prefix-len*}
Rtr1(config)#no ip nat pool *name*

name	Name of the pool
start-ip	Starting IP address that defines the range of addresses in the address pool
end-ip	Ending IP address that defines the range of addresses in the address pool
netmask *netmask*	Network mask that defines address bits identifying the network ID
prefix-length *prefix-len*	CIDR notation indicating how many bits make up the netmask

Examples:

```
Rtr1(config)#ip nat pool 2net-out 172.16.2.10 172.16.2.20 netmask 255.255.255.0
Rtr1(config)#ip nat pool 3net-out 172.16.3.10 172.16.3.99 prefix-length 24
```

Use a Standard ACL to Define the Local Addresses Eligible for Translation

A standard access list is used to define the internal addresses that can and can't be translated. In Global Configuration mode on SanJose1, make the following entries:

Examples:

```
Rtr1(config)#access-list 1 permit 192.168.1.0 0.0.0.255
Rtr1(config)#access-list 1 deny 192.168.2.5
Rtr1(config)#access-list 1 permit 192.168.2.0 0.0.0.255
```

Link the Pool of Global Addresses with the Eligible Local Addresses

To enable dynamic NAT of the inside source addresses, use the second form of the **ip nat inside source** command. To remove the static translation or remove the dynamic association to a pool, use the no form of this command. The syntax is

```
Rtr1(config)#ip nat inside source list {acl# | acl-name} {pool name | interface dialer-name}
Rtr1(config)#no ip nat inside source list {acl# | acl-name}{pool name | interface dialer-name}
```

list *acl# | acl-name*	Standard IP ACL number or name (packets permitted by the ACL are allowed to translate)
pool *name*	Name of the address pool from which global IP addresses are dynamically allocated
interface *dialer-name*	Name of the dialer interface on which the PPP/IPCP address negotiation takes place

Examples:

```
Rtr1(config)#ip nat pool 3net-out 172.16.3.10 172.16.3.99 netmask 255.255.255.0
Rtr1(config)#ip nat inside source list 1 pool 3net-out
Rtr1(config)#access-list 1 permit 192.168.2.0 0.0.0.255
```

Clearing the NAT Translation Table

Some dynamic translations stay in the translation table for 24 hours. Clearing dynamic NAT translations from the translation table before they timeout is possible by using the **clear ip nat translation** * command. This is particularly useful when testing NAT configuration.

Static NAT entries are immediately reestablished in the table.

Additional options exist for clearing only part of the table. Use the **clear ip nat translation ?** feature, if necessary.

Changing the NAT Default Inactivity Timeout Timers

TCP translations inactivity timeout defaults to 24 hours, unless an RST or FIN bit packet is seen on the stream, in which case they timeout in one minute. To change the amount of time after which NAT translations timeout, use the global configuration **ip nat translation** command. To disable the timeout, use the no form of this command.

```
Rtr1(config)#ip nat translation timeout seconds
Rtr1(config)#no ip nat translation timeout
```

timeout	Applies to dynamic translations, except for overload (PAT) translations. Default is 86,400 seconds (24 hours).
seconds	Seconds after which the specified port translation times out.

Examples:

The following example demonstrates changing the NAT default timeout value to 120 seconds on Rtr1.

```
Rtr1(confif)#ip nat translation timeout 120
```

Additional timer options are covered in the overload (PAT) section.

Dynamic NAT with Overloading (PAT)

Unless an organization has a large pool of global IP addresses, it seems the basic dynamic NAT translations could prove too limiting and cumbersome. Furthermore, having to maintain a large pool of global addresses somewhat defeats the address preservation aspects of private addresses.

The alternative is to have the NAT router create a unique identifier for every session by using a single global IP address and appending a port number, such as 1.1.1.1:1540. While over 65,000 port numbers exist, the effective limit is about 4,000. This means an organization could use and pay for only one real IP address and allow up to 4,000 hosts out onto the Internet. Even if NAT has a pool of two or more IP addresses to work with, the IOS version of NAT chooses to continue using the first IP address in the pool for subsequent translations.

This process is referred to as *dynamic overloading* or as Port Address Translation (PAT). To implement the feature requires adding only the one word **overload** to the **ip nat inside source list** command. An example would be as follows:

```
Rtr1(config)#ip nat pool 3net-out 172.16.3.10 172.16.3.99 netmask
     255.255.255.0
Rtr1(config)#ip nat inside source list 1 pool 3net-out overload
Rtr1(config)#access-list 1 permit 192.168.2.0 0.0.0.255
```

When the router receives the packet from an inside computer, it saves the computer's local IP address and port number to a NAT table. The router replaces the local IP address with a global IP address and adds the same port number. The translation table now has a mapping of the computer's nonroutable IP address and port number along with the router's IP address. The **show ip nat translation** command displays the result.

```
Rtr1#show ip nat translation
Pro Inside global     Inside local       Outside local   Outside global
icmp 172.16.2.5:1536  192.168.0.21:1536  10.0.0.5:1536   10.0.0.5:1536
tcp  172.16.2.5:1095  192.168.0.21:1095  10.0.0.19:21    10.0.0.19:21
tcp  172.16.2.5:1094  192.168.0.21:1094  10.0.0.45:23    10.0.0.45:23
Rtr1#
```

The first column (Pro), which has been blank before, now shows the protocol used.

When a packet comes from the destination computer, the router checks the destination port on the packet. The router then looks in the address translation table to see which computer on the stub domain the packet belongs to. It changes the destination address and the destination port to the one saved in the address translation table and sends it to that computer. The NAT router continues to use that same port number for the duration of the connection. The timer is reset each time the router accesses an entry in the table. If the entry isn't accessed again before the timer expires, the entry is removed from the table.

Changing the PAT Default Inactivity Timeout Timers

When port translation (PAT) is configured, a finer control exists over specific translation entries because each entry contains more context about the traffic using it. A separate entry is made for each timer. The syntax is

Rtr1(config)#ip nat translation {udp-timeout | dns-timeout | tcp-timeout | finrst-timeout} *seconds*

Rtr1(config)#no ip nat translation {udp-timeout | dns-timeout | tcp-timeout | finrst-timeout}

udp-timeout	Applies to the UDP port. Default is 300 seconds (five minutes).
dns-timeout	Applies to DNS connections. Default is 60 seconds.
tcp-timeout	Applies to the TCP port. Default is 86,400 seconds (24 hours).
finrst-timeout	To set the timeout value after a Finish or Reset TCP packet before terminating a connection. Default is 60 seconds.
seconds	Seconds after which the specified port translation times out.

Examples:

```
Rtr1(confif)#ip nat translation udp-timeout 120
Rtr1(confif)#ip nat translation dns-timeout 30
Rtr1(confif)#ip nat translation tcp-timeout 600
```

Dynamic NAT sessions can only be initiated by an internal host. Initiating a NAT translation from outside the network is impossible. To some extent, this adds a level of security to the internal network. This might also help explain why the dynamic timeout timer for overload (PAT) sessions is so short. The window of opportunity stays open just long enough to make sure legitimate replies like web pages, FTP and TFTP copies, and ICMP messages can get in.

Debugging IP NAT

To see the actual translation process and troubleshoot NAT problems, use the **debug ip nat** command and related options. As with all **debug** commands, this can have serious impacts on production routers and should be used judiciously. The **debug ip nat** output looks like this:

```
Rtr1# debug ip nat
IP NAT debugging is on
06:37:40: NAT:  s=192.168.0.10->172.16.2.5, d=172.16.1.97 [63]
06:37:40: NAT*: s=172.16.1.97, d=172.16.2.5->192.168.0.10 [63]
06:37:41: NAT*: s=192.168.0.10->172.16.2.5, d=172.16.1.97 [64]
06:37:41: NAT*: s=172.16.1.97, d=172.16.2.5->192.168.0.10 [64]
06:37:42: NAT*: s=192.168.0.10->172.16.2.5, d=172.16.1.97 [65]
06:37:42: NAT*: s=172.16.1.97, d=172.16.2.5->192.168.0.10 [65]
06:37:43: NAT*: s=192.168.0.10->172.16.2.5, d=172.16.1.97 [66]
06:37:43: NAT*: s=172.16.1.97, d=172.16.2.5->192.168.0.10 [66]
Rtr1#
06:38:43: NAT: expiring 172.16.2.5 (192.168.0.10) icmp 1536 (1536)
Rtr1#undebug all
All possible debugging has been turned off
```

The previous output shows the results of a ping from workstation to Rtr2.

You can see both translations as the pings pass both ways through the NAT router. The number at the end of the row is the same for both translations of each ping. The *s*= indicates the source, *d*= indicates the destination, and -> shows the translation.

The previous 06:38:43 entry shows the expiration of a NAT translation.

Using the **debug ip nat ?** command indicates you can refer to an access list number to define specific traffic to debug. The detailed option provides the port numbers, as well as the IP address translations. The following output shows the results of a ping from workstation to Rtr2. The second and third NAT output lines show the actual verification of IP address and port number assignment.


```
Rtr1#debug ip nat ?
  <1-99>    Access list
  detailed  NAT detailed events
  <cr>

Rtr1#debug ip nat detailed
IP NAT detailed debugging is on
07:03:50: NAT:  i: icmp (192.168.0.10, 1536) -> (172.16.1.97, 1536) [101]
07:03:50: NAT:  address not stolen for 192.168.0.10, proto 1 port 1536
07:03:50: NAT:  ipnat_allocate_port: wanted 1536 got 1536
07:03:50: NAT*: o: icmp (172.16.1.97, 1536) -> (172.16.2.5, 1536) [101]
07:03:51: NAT*: i: icmp (192.168.0.10, 1536) -> (172.16.1.97, 1536) [102]
07:03:51: NAT*: o: icmp (172.16.1.97, 1536) -> (172.16.2.5, 1536) [102]
07:03:52: NAT*: i: icmp (192.168.0.10, 1536) -> (172.16.1.97, 1536) [103]
07:03:52: NAT*: o: icmp (172.16.1.97, 1536) -> (172.16.2.5, 1536) [103]
07:03:53: NAT*: i: icmp (192.168.0.10, 1536) -> (172.16.1.97, 1536) [104]
07:03:53: NAT*: o: icmp (172.16.1.97, 1536) -> (172.16.2.5, 1536) [104]
Rtr1#
```

Rerouting Attacks

A router ICMP redirect message directs a host to use another router as its path to a particular destination because it has a better route. The rules say a router will send redirects only to hosts on its own local subnets. No user host will ever send a redirect and no redirect will travel more than one network hop. Unfortunately, attackers don't play by the rules. Some attacks are based on this.

Figure 5-6 shows a situation where using the host default gateway wouldn't reach the target destination in network 10.1.1.0. The following debug message shows Rtr1 sending a debug message to host 192.168.1.10 to use router Rtr2 (192.168.1.200) as the gateway to reach the destination 10.1.1.10.

```
Rtr1#debug ip icmp
ICMP packet debugging is on
ICMP: redirect sent to 192.168.1.10 for dest 10.1.1.10, use gw 192.168.1.200
Rtr1#
```

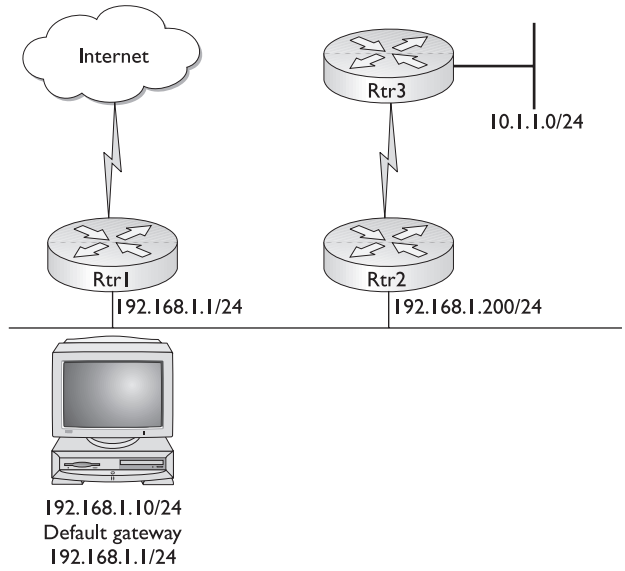

Figure 5-6 Router Rtr1 redirects traffic to Rtr2

By default, Cisco routers send ICMP redirects. You can use the interface subcommand **no ip redirects** to disable ICMP redirects. Another solution is to use an ACL to filter out any incoming ICMP redirects. The following code contains examples of each:

```
Rtr1(config)#interface ethernet 0
Rtr1(config-if)#no ip redirects
Rtr1(config-if)#^Z
Rtr1#
Rtr1#conf t
Rtr1(config)#access-list 125 deny   icmp any any redirect
Rtr1(config)#access-list 125 permit ip any any
Rtr1(config)#interface serial 0
Rtr1(config-if)#ip access-group 125 in
Rtr1(config-if)#^Z
Rtr1#
```

This filtering prevents only redirect attacks by remote attackers, but does nothing against an attacker that has internal access to the same segment as a host that's under attack.

Event Logging on Perimeter Routers

Perimeter router logs can be invaluable in troubleshooting, capacity planning, and dealing with security incidents. For security purposes, the events to log are interface status changes, changes to the system configuration, access list matches, events detected by the firewall, and intrusion detection features. System logging events might be reported to a variety of destinations, including the following:

- The system console port (**logging console** command). Because many console ports are unattended or are connected to terminals with no historical storage, this information might be unavailable to reconstruct a major event.
- Servers running the syslog daemon can send logging information to a server with the **logging** ***server-ip-address*** command, and you can control the urgency threshold for logging to the server with the **logging trap** ***urgency*** command. Even if you have a syslog server, you should probably still enable local logging. If you don't have access to a syslog server, go to Kiwi Enterprises at http://www.kiwisyslog.com /index.htm and download its free Kiwi Syslog Daemon.
- Remote sessions on VTYs and local sessions on TTYs (**logging monitor** and **terminal monitor** commands).
- Most routers can save system logging information to a local RAM buffer. This buffer is a fixed size and retains only the most recent information, and the contents are lost whenever the router is reloaded. Use the **show memory** command to make sure your router has enough free memory to support a logging buffer. Create the buffer using the **logging buffered** ***buffer-size*** configuration command.

If the router has a real-time clock or is running NTP, time-stamp log entries by adding the **service timestamps log datetime msecs** command to the configuration.

Access List Violation Logs

With traffic filtering and access control ACLs, you should consider logging packets that violate the filtering criteria. Older Cisco IOS software versions use the **log** keyword option, which captures the IP addresses and port numbers of packets matching an access list entry. Newer IOS versions use the **log-input** keyword, which adds receiving interface information and the MAC address of the host that sent it.

To manage file size and minimize performance impacts, configure logging for those critical access list entries. Don't log entries that will match a large number of packets and generate little useful information, such as the **permit any** statement.

Chapter Review

This chapter focused on the IOS features that could be used on the perimeter router as a first line of defense against security threats. The perimeter configuration of the network often includes both a perimeter router and a firewall as the second line of defense. The firewall separates the inside network from the DMZs. The dirty DMZ is protected only by the perimeter router, while the protected DMZ has the firewall and perimeter router between it and the outside.

Network security can be enhanced by disabling unused services, such as CDP, finger, and TCP and UDP small services.

Cisco IOS offers a rich selection of routing and route security tools, such as controlling directed broadcasts, blocking ICMP redirects, routing protocol authentication, and flooding control.

Controlling network access and traffic using address filtering, dynamic access lists, and reflexive access lists can all contribute to increased security.

Questions

1. True or False. In the screened subnet architecture network model, the inside network is everything from the perimeter router in to the corporate network.
   - **A.** True
   - **B.** False
2. Which one of the following is considered the trusted network?
   - **A.** Inside
   - **B.** Outside
   - **C.** Dirty DMZ
   - **D.** Protected DMZ
3. Which of the following would *not* be a function of a perimeter router in a screened subnet architecture network?
   - **A.** Providing a serial connection to the outside world
   - **B.** Providing any filtering of outside traffic
   - **C.** Providing LAN routing
   - **D.** Implementing basic security for the dirty DMZ
4. True or False. CDP facilitates a secure environment on a perimeter router.
   - **A.** True
   - **B.** False
5. Which one is *not* true about IP directed broadcast?
   - **A.** It's a datagram sent to the subnet broadcast address.
   - **B.** It's routed through the network as a unicast packet.
   - **C.** Only the router directly connected to the target subnet can positively identify it.
   - **D.** It can be blocked by a smurf defense.
6. True or False. Filtering incoming ICMP redirects on a perimeter router should never cause any problems.
   - **A.** True
   - **B.** False

7. Which two of the following reduces spoofing attacks?
   - **A.** RFC 2827 filtering
   - **B.** Weighted fair queuing
   - **C.** RFC 1918 filtering
   - **D.** Routing protocol authentication
8. Which of the following is most like the TCP established option?
   - **A.** Dynamic ACL
   - **B.** Lock and key
   - **C.** Reflexive ACL
   - **D.** Finger ACL
9. In NAT terminology, what's the IP address of a network member computer?
   - **A.** Inside local
   - **B.** Outside local
   - **C.** Inside global
   - **D.** Outside global
10. Which statement is *not* true about Network Address Translation (NAT)?
    - **A.** It's a mechanism that allows private addresses to be translated to use the Internet.
    - **B.** It can be configured both static and dynamic on the same router.
    - **C.** It provides good security by hiding internal IP addresses.
    - **D.** It reduces the cost of IP addresses.
11. True or False. Static NAT entries appear in the translation table the first time they're used.
    - **A.** True
    - **B.** False
12. Which command shows the NAT table?
    - **A.** show ip nat statistics
    - **B.** show run
    - **C.** show ip nat translations
    - **D.** show ip nat table
13. What one word changes dynamic NAT to PAT?
    - **A.** PAT
    - **B.** Overflow

C. Overload

D. Rotary

14. Which command sets the idle timeout for a dynamic (lock-and-key) access list?

A. access-list 101 dynamic temp-in timeout 30 permit ip any any

B. ip dynamic-list timeout 30

C. autocommand access-enable host timeout 30

D. ip reflexive-list timeout 30

15. Which statement is true about reflexive access lists?

A. They create temporary holes into the network security, based on a successful Telnet authentication.

B. They only work with TCP traffic.

C. They create temporary holes in the network security–based specific outbound traffic.

D. They rely on named standard access lists.

Answers

1. **B.** False. It's everything in from the inside interface of the firewall.
2. **A.** Inside
3. **C.** Providing LAN routing
4. **B.** False. It announces to any system on a directly connected segment that the router is a Cisco device, the model number, and the Cisco IOS version being run.
5. **D.** It can be blocked by a smurf defense.
6. **A.** True. They shouldn't come from outside the segment.
7. **A.** RFC 2827 filtering and **C.** RFC 1918 filtering
8. **C.** Reflexive ACL
9. **A.** Inside local
10. **C.** It provides good security by hiding internal IP addresses. It provides limited privacy.
11. **B.** False. They appear when created.
12. **C.** show ip nat translations
13. **C.** Overload
14. **C.** autocommand access-enable host timeout 30
15. **C.** They create temporary holes in the network security–based specific outbound traffic.

CHAPTER 6

IOS Firewall Feature Set—CBAC

In this chapter, you will learn how to:

- Recognize Cisco IOS Firewall security problems and solutions
- Use Context-Based Access Control (CBAC)
- Configure the Cisco IOS Firewall
- Use the two methods of IOS Firewall Administration

With the growing awareness of, and commitment to, network security as the way to secure business transactions both internally and over the Internet, more and more businesses are recognizing that security concepts and features must be integrated into the network design and infrastructure.

According to Cisco advertising, 80 percent of the Internet backbone routers run Cisco IOS software, making it the most fundamental component of today's network infrastructure. The Cisco IOS software-based security features, combined with the wide range of products supported, make it a logical solution for any organization's end-to-end Internet, intranet, and remote access network security requirements.

This chapter looks at the key components of the IOS Firewall features, and compares the functionality and limitations to the PIX Firewall devices. PIX Firewall devices are introduced in the next chapter, and then expanded on later in the advanced PIX chapters.

STUDY TIP The material in this chapter is part of the exam objectives for both the Managing Cisco Network Security (MCNS) and Cisco Secure PIX Firewall Advanced (CSPFA) exams.

Introduction to Cisco IOS Firewall

The *Cisco IOS Firewall* is a feature set option for Cisco IOS software, that is available for a wide range of Cisco routers and switches. It provides advanced firewall capabilities, as well as other security technologies such as intrusion detection and authentication.

The Cisco IOS Firewall is only one part of the growing Cisco Secure product family. It joins and integrates with other Cisco Secure products, such as Cisco Secure Access Control

Server (ACS), Cisco Secure PIX Firewall devices, Cisco Secure Intrusion Detection System (IDS), Cisco Secure Scanner, and Cisco Secure Policy Manager (CSPM), plus a variety of consulting and training services options. As key components of Cisco AVVID strategy, Cisco Secure network security solutions improve the network's capability to support mission-critical Internet applications.

Cisco IOS software firewall feature set provides an extensive array of security tools, enabling the administrator to configure a firewall providing the appropriate level of functionality to meet security policy requirements. Whether as a stand-alone unit or working in tandem with a dedicated firewall device, the IOS firewall features provide flexible and reliable options to extend the security protection throughout the network.

Router-Based Firewall Functionality

Cisco IOS Firewall is available on the Cisco 800, uBR900, 1400, 1600, 1700, 2500, 2600, 3600, 7100, 7200, and 7500, as well as RSM series routers. By supporting such a wide variety of router platforms, the firewall features will scale to meet any network's bandwidth, performance, and security requirements. The firewall features are also available on the Catalyst 5000 switch.

No absolute rules exist for choosing the right Cisco router for implementing the firewall and security features. While an organization's size and security requirements must be factored in, the following general guidelines can be used as a starting point.

- Small office/home offices (SOHO): Cisco 800, uBR900 series, 1600, and 1720 series
- Branch and extranet implementations: Cisco 2500, 2600, and 3600 series
- Central office or high-volume locations: Cisco 7100, 7200, 7500, and RSM series

Integration with Cisco IOS Software

The *Cisco IOS Firewall feature set* is an optional security solution integrated into the network through selecting the correct Cisco IOS software. The integration of the firewall features into the IOS means the organization security policy can be implemented and enforced throughout the network. Whether securing the links between departments or partner networks, or between the organization and the Internet, the IOS implementation allows for an end-to-end security solution that can grow and change with the organization.

The Cisco IOS Firewall is also completely interoperable with, and often enhances, other features, such as AAA, NAT, Cisco encryption technology (CET), and system logging, as well as standard and extended access control list features, such as Time-Based and Lock n Key.

VPN, IPSec Encryption, and QoS Support

When combined with Cisco IPSec technology, the Cisco IOS Firewall features provide integrated and enhanced virtual private network (VPN) functionality. VPNs are rapidly

evolving as the standard for providing secure data communications over public networks like the Internet.

The firewall features work with the IOS encryption, tunneling, and quality of service (QoS) features to ensure timely and reliable delivery of data and provide robust perimeter security, while providing advanced bandwidth management, intrusion detection, and service-level validation. The Firewall Authentication Proxy feature can provide user authentication and authorization for Cisco VPN client software.

Several chapters are dedicated to IPSec and to VPN implementation and features, but it's important to recognize that both the firewall and the IPSec features are separate IOS options that might be selected to match the objectives of the security policy. Figure 6-1 shows a sample of the Cisco IOS Upgrade Planner from the Cisco web site. Notice that the firewall features, identified as FW, and the IPSec features, IPSEC 56, are available in various combinations on different platforms. You can have either, both, or neither. Often memory and flash requirements will increase for each, as might the cost.

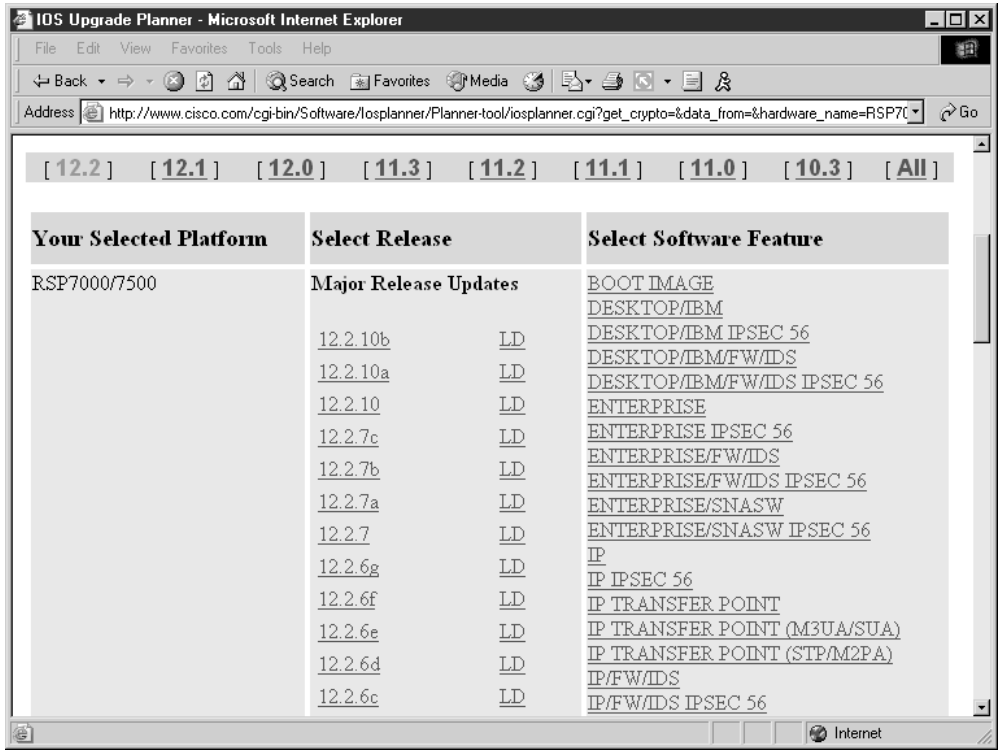

Figure 6-1 Cisco CCO IOS Upgrade Planner showing feature sets

Does the IOS Image Support Firewall and IPSec Features?

How do you tell if the IOS on a router supports the firewall or the IPSec features? Cisco and vendors who use Cisco IOS on their devices follow a naming convention for software images. The device IOS convention identifies the platform, the features of the image, and the area of memory used by the image at runtime. For the image c2600-jos56i-mz.120-4 .T.bin, the name is the portion up to the first period and everything else is version/release information. The three components of the name are separated by dashes. In the example,

- The platform is a Cisco 2600 series router.
- The *o* feature code indicates firewall features. The 56i indicates 56-bit IPSec encryption. The *s* indicates a group of features, such as NAT, ISL IBM, MMP, VPDN/L2F, VOIP, and ATM. The *j* indicates enterprise features.
- The *m* indicates that the software runs from RAM, and the *z* indicates the image is zipped.

For more information on naming convention and version numbers, go to http://www.cisco.com/warp/public/620/4.html. Naming is in Section 2.2. No CCO account is required to view this document.

Feature Summary

The Cisco IOS Firewall feature set adds the following security features to a broad selection of Cisco routers:

- Intrusion detection
- Authentication proxy
- Context-Based Access Control (CBAC)
- Dynamic port mapping
- Java blocking
- Denial of service (DoS) detection and prevention
- Simple Mail Transfer Protocol (SMTP) attack detection and prevention
- IP fragmentation attack prevention
- Configurable real-time alerts and audit trail
- Failover capabilities
- Microsoft NetShow application support

The 800, UBR904, 1600, and 2500 series of routers support all the previous firewall features, except authentication proxy and intrusion detection. There's no firewall feature set support for the 4000 series routers, but there is VPN support.

Context-Based Access Control (CBAC)

Context-Based Access Control (CBAC) is a per-application control mechanism that adds advanced traffic filtering functionality to firewalls that isn't limited, as are access lists, to examining packets at the network or transport layer. While CBAC examines both of these layers, it also examines the application-layer protocol data to monitor the state of a given TCP or UDP session. This means, as multiple channels are created or used by applications such as SQL*Net, FTP, and RPC, CBAC can respond by creating temporary openings in the firewall access lists to allow return traffic and additional data connections for specified sessions that originated from within the protected network. This application-layer awareness and capability to evolve with the traffic is beyond the capabilities of access list technologies.

Quick Access List Review

Before continuing with CBAC, it's important to be clear about how standard and extended ACLs work. By definition, standard ACLs filter only on source network addresses and are, therefore, limited to Layer 3 capabilities. Extended ACLs are able to filter on port numbers extending their reach into Layer 4. In both cases, any ACL allowing traffic to enter a network is, in fact, a hole in the firewall or perimeter security that can possibly be exploited by others.

The preceding chapter introduced reflexive ACLs as an alternative to creating permanent holes through the network security. Temporary ACL statements can be created for inbound traffic based on outbound traffic reducing risk of exploitation. Unfortunately, reflexive ACLs are limited to Layer 4 filters, like any other extended ACL. Furthermore, reflexive ACLs can't deal with changes in port designations by the outside host, such as FTP. The outbound address/port combinations for the source and destination are "mirrored" to create the inbound openings. Another limitation of reflexive ACLs is that they're limited to single channel applications.

CBAC Advantages

Understanding CBAC might be made easier if you think of it as reflexive ACLs without the limitations. CBAC adds inspection intelligence to ACL capabilities by reading the entire packet for application status information, which is stored in the state table. Like reflexive ACLS, CBAC watches outbound traffic to determine what packets to let in; but unlike reflexive ACLs, CBAC can make decisions based on how the application behaves, not only the addresses and port number it uses.

CBAC can open any additional inbound channels required for returning data that were negotiated by the outgoing data for a particular application.

When a session times out or ends, the state table and ACL entries are deleted, and the opening closes to additional traffic.

CBAC can be configured to inspect and filter the following IP sessions and application-layer protocols:

- All TCP sessions, regardless of the application-layer protocol (sometimes called *single-channel* or *generic TCP inspection*).
- All UDP sessions, regardless of the application-layer protocol (sometimes called *single-channel* or *generic UDP inspection*).
- CU-SeeMe (White Pine version only), an Internet videoconferencing program developed as freeware by Cornell University. WhitePine, Inc., sells an enhanced commercial version.
- FTP doesn't support third-party connections (three-way FTP transfer). Allows data channels with the destination ports 1024 to 65535. CBAC won't open a data channel if the FTP client-server authentication fails.
- HTTP (Java blocking).
- Microsoft NetShow.
- UNIX R-commands, such as rlogin, rexec, and rsh.
- RealAudio.
- H.323, such as NetMeeting and ProShare
- Real-Time Streaming Protocol (RTSP): CBAC supports the following RTSP data transport modes.
  - Standard Real-Time Transport Protocol (RTP) IETF standard (RFC 1889) for real-time audio and video applications, such as Cisco IP/TV and Apple QuickTime 4 software. RTP uses the RTP Control Protocol (RTCP) to manage the multimedia data stream.
  - RealNetworks Real Data Transport (RDT) Proprietary protocol developed by RealNetworks used for RealServer G2. Uses RTSP for communication control and RDT for the data connection and retransmission of lost packets.
  - Interleaved (Tunnel Mode) Uses the control channel to tunnel RTP or RDT traffic.
  - Synchronized Multimedia Integration Language (SMIL) Layout language that allows the creation of multimedia presentations made up of music, voice, image, text, video, and graphics elements. Uses multiple RTSP control and data streams between the player and the servers. Currently available only for RTSP and RDT, but SMIL is a proposed specification of the World Wide Web Consortium (W3C). RealNetworks RealServer and RealServer G2 support SMIL, while Cisco IP/TV and Apple QuickTime 4 don't.
  - RPC Sun RPC, but not DCE RPC.
  - Simple Mail Transport Protocol (SMTP) CBAC can inspect SMTP, but not Extended Simple Mail Transport Protocol (ESMTP).

- SQL*Net
- StreamWorks
- TFTP
- VDOLive

CBAC Limitations

As with all things that seem too good to be true, some limitations must be recognized and worked around:

- Only IP TCP and UDP traffic is inspected by CBAC, so ICMP traffic and any other Layer 3 protocols need to be filtered using extended ACLs.
- Any traffic where the router is the source or destination won't be inspected. CBAC will filter traffic passing through, but not traffic originating or terminating on that device.
- Because CBAC only detects and protects against attacks that travel through the firewall, it doesn't normally protect against attacks originating from within the protected network. Deploying CBAC on an intranet-based router is possible.
- CBAC can't inspect in-transit IPSec traffic. Because the IPSec traffic is encrypted, CBAC can't interpret it and, therefore, drops it. CBAC and IPSec can *only* work together at tunnel endpoint by applying IPSec to the external interface and CBAC on the internal interface.

Memory and Performance Issues

You need to consider two issues when determining which model of router to use for CBAC and how much memory to install.

- CBAC uses about 600 bytes of memory per connection established.
- A slight amount of additional CPU processing is used to inspect packets. While evaluating long access lists can negatively impact performance, CBAC mitigates this by hashing ACLs, and then evaluates the hash.

CBAC Process

The following steps describe the sequence of events for CBAC configured on an external interface connected to the Internet. The example starts with an outbound packet, which is the first of a new TCP session configured for CBAC inspection.

1. The outbound packet reaches the firewall's external interface.
2. The packet is checked against the interface's existing outbound ACL, and the packet is permitted. A denied packet would be discarded at this point and couldn't be evaluated by CBAC to allow inbound traffic. With no outbound traffic, there shouldn't be returning traffic.

3. If the packet's application is configured for CBAC inspection, CBAC inspects the packet to determine and record the state information of the connection. As a new session, a new state table entry is created for the new connection. If the application isn't configured for CBAC inspection, it would go directly to Step 5.
4. Using the state information, CBAC creates a temporary entry inserted at the beginning interface's inbound extended ACL. This temporary entry is designed to allow inbound packets that are part of the same session.
5. The outbound packet is forwarded out the interface.
6. When an inbound packet from the just-established session reaches the interface, it's evaluated against the inbound ACL and permitted by the temporary entry created in Step 4.
7. The packet permitted in Step 6 is then inspected by CBAC, and the connection's state table entry is updated as necessary. Based on this updated state information, the inbound extended ACL temporary entries can be modified to permit only packets that are valid for the current state of the connection.
8. Additional inbound and outbound packets from the connection are inspected to update the state table and the temporary inbound ACL entries as needed. The packets are forwarded through the interface.
9. When the connection terminates or times out, the related entries in the state table and the inbound ACL are deleted.

Configuring CBAC

To configure CBAC, perform the following tasks. Each task is described in the following sections.

- Set Audit Trails and Alerts
- Set Global Timeouts and Thresholds
- Define Port-to-Application Mapping (PAM)
- Define Inspection Rules
- Apply Inspection Rules and ACLs to an interface
- Test and Verify

Set Audit Trails and Alerts

Real-time alerts send syslog error messages to central management consoles upon detecting suspicious activity, allowing network managers to respond immediately to intrusions. Enhanced audit-trail features use syslog to track all transactions, recording time stamps, source host, destination host, ports used, session duration, and the total number of transmitted bytes for advanced, session-based reporting.

Cisco IOS Firewall alerts and audit-trail features are now configurable, enabling more flexible reporting and error tracking. The configurable audit-trail features support modular tracking of specific CBAC-supported applications and Java blocking. Both the real-time alerts and the audit-trail features are supported by a variety of third-party reporting tools.

Use the Global Configuration Mode command **ip inspect audit-trail** to turn on CBAC audit-trail messages. The messages are displayed on the console after each CBAC session closes. Use the no form of the command to turn off the feature. The syntax is

```
Rtr1(config)#ip inspect audit-trail
Rtr1(config)#no ip inspect audit-trail
```

This command has no arguments or keywords. By default, the audit-trail messages aren't displayed. This command was introduced in IOS 11.2 P.

The following messages are two examples of audit-trail messages. To determine which protocol was inspected, refer to the responder's port number following the responder's IP address.

```
%FW-6-SESS_AUDIT_TRAIL: tcp session initiator (192.168.1.13:33192)
     sent 22 bytes -- responder (192.168.129.11:25) sent 208 bytes
%FW-6-SESS_AUDIT_TRAIL: ftp session initiator 192.168.1.13:33194)
     sent 336 bytes -- responder (192.168.129.11:21) sent 325 bytes
```

CBAC alert messages are displayed on the console by default. Use the Global Configuration Mode command **ip inspect alert-off** to disable these messages. To reenable CBAC alert messages, use the no form of the command. The syntax is

```
Rtr1(config)#ip inspect alert-off
Rtr1(config)#no ip inspect alert-off
```

This command has no arguments or keywords. This command was introduced in IOS 12.0(5)T.

Logging While it isn't a CBAC feature, make sure logging is enabled and a syslog server is specified, so any incidents are properly logged for future reference.

```
Rtr1(config)#service timestamps log datetime
Rtr1(config)#logging host
Rtr1(config)#logging facility facility-type
Rtr1(config)#logging trap level
```

The commands,—in order,—do the following:

1. Adds the date and time to syslog and audit-trail messages.
2. Defines the host name or IP address of the host to send syslog messages to.
3. Configures the syslog facility in which error messages are sent.
4. (Optional) Limits messages logged to the syslog servers based on severity. The default is level 7 (informational).

Set Global Timeouts and Thresholds

This section looks at configuring the following global timeouts and thresholds used by CBAC.

- TCP SYN and FIN wait times
- TCP, UDP, and DNS idle timers
- TCP flooding thresholds (DoS indicators)

These are global default settings used by CBAC to determine how long to maintain state table entries and as indicators of possible DoS attacks. Each command has a default value and, therefore, needs to be set only to change the default to implement a security policy better.

TCP Session Establishment Timer Use the Global Configuration Mode command **ip inspect tcp synwait-time** to define the number of seconds the software will wait for a TCP session to reach the established state before dropping the session. The session is considered to have reached the established state after the session's first SYN bit is detected. Use the no form of this command to reset the timeout to the default of 30 seconds. The syntax is

```
Rtr1(config)#ip inspect tcp synwait-time seconds
Rtr1(config)#no ip inspect tcp synwait-time
```

This command was introduced in IOS 11.2 P. The default is 30 seconds.
The value specified for this timeout applies to all TCP sessions inspected by CBAC.

TCP Session Termination Timer Use the Global Configuration Mode command **ip inspect tcp finwait-time** to define how many seconds a TCP session will still be managed after the firewall detects a FIN-exchange. The *FIN-exchange* occurs when the TCP session is ready to close. Use the no form of the command to reset the timeout to the default of five seconds. The syntax is

```
Rtr1(config)#ip inspect tcp finwait-time seconds
Rtr1(config)#no ip inspect tcp finwait-time
```

This command was introduced in IOS 11.2 P. The default is five seconds.
The timeout set with this command is referred to as the *finwait timeout*. It applies to all CBAC-inspected TCP sessions.

TCP Session Inactivity Timer Use the Global Configuration Mode command **ip inspect tcp idle-time** to specify the TCP idle timeout—the number of seconds a TCP session will still be managed after no activity. Use the no form of the command to reset the timeout to the default of 3,600 seconds (one hour). The syntax is

```
Rtr1(config)#ip inspect tcp idle-time seconds
Rtr1(config)#no ip inspect tcp idle-time
```

This command was introduced in IOS 11.2 P. The default is 3,600 seconds (one hour).

When CBAC detects a valid TCP packet that's the first in a session for a protocol CBAC is inspecting, the software creates a new state table entry with the information. If no TCP packets for a particular session are detected for the time defined by the TCP idle timeout, the software drops that session entry from the state table and ACL.

This global value can be overridden for specific interfaces by defining a set of inspection rules with the Global Configuration Mode command **ip inspect name**. This command only applies the new timeout to any new or existing inspection rules *that don't have an explicitly defined timeout.*

UDP Session Inactivity Timer Use the Global Configuration Mode command **ip inspect udp idle-time** to specify the UDP idle timeout, the number of seconds a UDP "session" will still be managed after no activity. Use the no form of the command to reset the timeout to the default of 30 seconds. The syntax is

```
Rtr1(config)#ip inspect udp idle-time seconds
Rtr1(config)#no ip inspect udp idle-time
```

This command was introduced in IOS 11.2 P. The default is 30 seconds.

When CBAC detects a valid UDP packet that's the first in a session for a protocol CBAC is inspecting, the software creates a new state table entry with the information. Because UDP is a connectionless service, no actual sessions exist, so the software approximates sessions by examining the information in the packet and determining if the packet is similar to other UDP packets (for example, similar source/destination addresses) and if the packet was detected soon after another similar UDP packet.

Because UDP is a connectionless service, no actual sessions exist as with TCP, so CBAC approximates sessions. It does this by examining the packets and determining if they're similar (source/destination addresses and ports) to other UDP packets. If no UDP packets for a particular session are detected for the time defined by the UDP idle timeout, the software drops those session entries from the state table and ACL.

This global value can be overridden for specific interfaces when you define a set of inspection rules with the Global Configuration Mode command **ip inspect name**. This command only applies the new timeout to any new or existing inspection rules *that don't have an explicitly defined timeout.*

DNS Session Inactivity Timer Use the Global Configuration Mode command **ip inspect dns-timeout** to specify the DNS idle timeout, the length of time a DNS-name lookup session will still be managed after no activity. Use the no form of the command to reset the timeout to the default of five seconds. The syntax is

```
Rtr1(config)#ip inspect dns-timeout seconds
Rtr1(config)#no ip inspect dns-timeout
```

This command was introduced in IOS 11.2 P. The default is five seconds.

When CBAC detects a valid UDP packet for a new DNS-name lookup session for a protocol CBAC is inspecting, the software creates a new state table entry with the information.

If the software detects no packets for the DNS session for a time period defined by the DNS idle timeout, the software then drops that session entry from the state table and ACL.

The DNS idle timeout applies to all DNS-name lookup sessions inspected by CBAC and overrides the global UDP timeout. The DNS idle timeout value also enters Aggressive mode and overrides any timeouts specified for specific interfaces when you define a set of inspection rules with the **ip inspect name** command.

Maximum Incomplete Sessions High/Low Threshold An unusually high number of half-open sessions can indicate a DoS attack is occurring. For TCP, *half-open* means that the session hasn't reached the established state. For UDP, *half-open* means that the firewall has detected traffic from one direction only.

CBAC measures both the total number of existing half-open sessions and the rate of session establishment attempts. Both TCP and UDP half-open sessions are counted in the total number and rate measurements. Measurements are made once a minute.

When the number of existing half-open sessions rises above the threshold set by the **ip inspect max-incomplete high** command, the software then deletes half-open sessions until the number of existing half-open sessions drops below the threshold set by the **ip inspect max-incomplete low** command.

The global value specified for this threshold applies to all TCP and UDP connections inspected by CBAC.

Use the Global Configuration Mode command **ip inspect max-incomplete high** to define the number of existing half-open sessions that will cause the software to start deleting half-open sessions. Use the no form of the command to reset the threshold to the default of 500. The syntax is

```
Rtr1(config)#ip inspect max-incomplete high number
Rtr1(config)#no ip inspect max-incomplete high
```

This command was introduced in IOS 11.2 P. The default is 500 half-open sessions.

Use the Global Configuration Mode command **ip inspect max-incomplete low** to define the number of existing half-open sessions that will cause the software to stop deleting half-open sessions. Use the no form of the command to reset the threshold to the default of 400 half-open sessions. The syntax is

```
Rtr1(config)#ip inspect max-incomplete low number
Rtr1(config)#no ip inspect max-incomplete low
```

This command was introduced in IOS 11.2 P. The default is 400 half-open sessions.

The following example causes the CBAC to start deleting half-open sessions when the number of half-open sessions rises above 800 and to stop when the number drops below 500.

```
Rtr1(config)#ip inspect max-incomplete high 800
Rtr1(config)#ip inspect max-incomplete low 500
```

One Minute Incomplete Sessions High/Low Threshold This is an extension of the preceding threshold, accelerating the process to respond to a rapid increase in incomplete sessions. These rate thresholds are measured as the number of new session connection attempts detected in the last one-minute sample period.

The global value specified for this threshold applies to all TCP and UDP connections inspected by CBAC.

Use the Global Configuration Mode command **ip inspect one-minute high** to define the number of new unestablished sessions that will cause the software to start deleting half-open sessions. Use the no form of the command to reset the threshold to the default of 500. The syntax is

```
Rtr1(config)#ip inspect one-minute high number
Rtr1(config)#no ip inspect one-minute high
```

This command was introduced in IOS 11.2 P. The default is 500 half-open sessions.

Use the Global Configuration Mode command **ip inspect one-minute low** to define the rate of new unestablished TCP sessions that will cause the software to stop deleting half-open sessions. Use the no form of this command to reset the threshold to the default of 400. The syntax is

```
Rtr1(config)#ip inspect one-minute low number
Rtr1(config)#no ip inspect one-minute low
```

This command was introduced in IOS 11.2 P. The default is 400 half-open sessions.

The following example causes the software to start deleting half-open sessions when more than 1,000 session establishment attempts are detected in the last minute and to stop when fewer than 750 sessions are detected in the last minute:

```
Rtr1(config)#ip inspect one-minute high 1000
Rtr1(config)#ip inspect one-minute low 750
```

Maximum Incomplete Sessions Per Destination Host Threshold An unusually high number of half-open sessions with the same destination host address can indicate that a DoS attack is being launched against the host. Use the Global Configuration Mode command **ip inspect tcp max-incomplete host** to specify threshold and blocking time values for TCP host–specific DoS detection and prevention. Use the no form of the command to reset the threshold and blocking time to the default values. The syntax is

```
Rtr1(config)#ip inspect tcp max-incomplete host number block-time seconds
Rtr1(config)#no ip inspect tcp max-incomplete host
```

number	Specifies how many half-open TCP sessions with the same host destination address can exist at a time, before the software starts deleting half-open sessions to the host. Use a number from 1 to 250.
block-time	Specifies blocking of connection initiation to a host.
seconds	Specifies how long the software will continue to delete new connection requests to the host.

This command was introduced in IOS 11.2 P. The default is 50 half-open sessions and 0 minutes.

When the numbers of half-open sessions with the same destination host address rises above this threshold, CBAC will delete half-open sessions choosing a method based on the block-time seconds setting. If the timeout is

0 (the default)	CBAC will delete the oldest half-open session for the host for every new connection request to the host. This ensures the number of half-open sessions to a given host will never exceed the threshold.
Greater than 0	CBAC will delete all existing half-open sessions for the host, and then block all new connection requests to the host until the block-time expires.

The software also sends syslog messages whenever the max-incomplete host number is exceeded and when blocking of connection initiations to a host starts or ends.

The global values specified for the threshold and blocking time apply to all TCP connections inspected by CBAC.

The following example changes the max-incomplete host number to 70 half-open sessions and changes the block-time timeout to 90 seconds.

```
Rtr1(config)#ip inspect tcp max-incomplete host 70 block-time 90
```

Define Port-to-Application Mapping (PAM)

Flexible, per-application port mapping allows CBAC-supported applications to be run on nonstandard TCP and UDP ports. PAM allows network administrators to customize access control for specific applications and services to meet the distinct needs of their networks.

The **ip port-map** command associates TCP or UDP port numbers with applications or services, establishing a table of default port mapping information at the firewall. This information is used to support network environments that run services using ports that are different from the registered or well-known ports associated with a service or application. PAM also supports port mapping for specific host(s) or subnet(s) by using standard ACLs.

The port mapping information in the PAM table is one of three types:

- System defined
- User defined
- Host specific

System-Defined Port Mapping Initially, PAM creates a set of system-defined entries in the mapping table using well-known or registered port mapping information set up during the system startup. The Cisco IOS Firewall CBAC feature requires the system-defined mapping information to function properly. The system-defined mapping information can't be deleted or changed. It isn't possible to assign an application to an existing system-defined mapping, such as attempting to map HTTP services to port 25 (SMTP). The following table shows the well-known or registered port mapping information.

Application Name	Registered Port Number	Protocol Description
Cuseeme	7648	CU-SeeMe Protocol
Exec	512	Remote process execution
ftp	21	File Transfer Protocol (control port)
http	80	Hypertext Transfer Protocol
h323	1720	H.323 Protocol (such as MS NetMeeting and Intel Video Phone)
login	513	Remote login
msrpc	135	Microsoft Remote Procedure Call
netshow	1755	Microsoft NetShow
real-audio-video	7070	RealAudio and RealVideo
smtp	25	Simple Mail Transfer Protocol
sql-net	1521	SQL-NET
streamworks	1558	StreamWorks Protocol
sunrpc	111	SUN Remote Procedure Call
tftp	69	Trivial File Transfer Protocol
vdolive	7000	VDOLive Protocol

User-Defined Port Mapping Network applications that use nonstandard ports require user-defined entries in the mapping table. Use the Global Configuration Mode command **ip port-map** to create user-defined entries ports to application mapping. Use the no form of the command to delete user-defined PAM entries. The command can't be used to change system-defined port mappings. The syntax is

Rtr1(config)#ip port-map *appl-name* port *port-num* [list *acl#*]
Rtr1(config)#no ip port-map *appl-name* port *port-num* [list *acl#*]

appl-name	The name of the application with which to apply the port mapping
port	Indicates a port number maps to the application
port-num	Port number (1 to 65535)
list	The port mapping information applies to a specific host or subnet
acl#	Standard ACL number used to identify the host(s) or subnet(s)

This command was introduced in IOS 12.0(5)T. No default values.

This example shows PAM entries that define a range of nonstandard ports for HTTP services.

```
Rtr1(config)#ip port-map http port 8000
Rtr1(config)#ip port-map http port 8001
Rtr1(config)#ip port-map http port 8002
```

Host-Specific Port Mapping User-defined entries in the mapping table can include host-specific mapping, which establishes port mapping information for specific hosts or subnets. In some situations, it might be necessary to override the default port mapping information for a specific host or subnet, including a system-defined default

port mapping information. Use the list option for the **ip port-map** command to specify an ACL for a host or subnet that uses PAM.

In this example, a specific host uses port 8000 for FTP services. ACL 1 identifies the server address (192.168.0.100), while port 8000 is mapped with FTP services:

```
Rtr1(config)#access-list 1 permit 192.168.0.100
Rtr1(config)#ip port-map ftp port 8000 list 1
```

In the next example, the same port number is required by different services running on different hosts. Port 8000 is required for FTP services by host 192.168.0.100, while port 8000 is required for HTTP services by host 192.168.0.175. ACL 10 and ACL 2 identify the specific hosts, while PAM maps the ports with the services for each ACL.

```
Rtr1(config)#access-list 1 permit 192.168.0.100
Rtr1(config)#access-list 2 permit 192.168.0.175
Rtr1(config)#ip port-map ftp port 8000 list 1
Rtr1(config)#ip port-map http port 8000 list 2
```

This example shows a failed attempt to assign the RealAudio application to port 21, which is normally reserved for FTP services. Following that is the correct method to define the host using ACL 1. With this configuration, host(s) in List 1 won't recognize FTP activity on port 21.

```
Rtr1(config)#ip port-map realaudio port 21
Rtr1(config)#Command fail: the port 21 has already been defined for
         ftp by the system.
         No change can be made to the system defined port mappings.
Rtr1(config)#access-list 1 permit 192.168.0.100
Rtr1(config)#ip port-map realaudio port 21 list 1
```

Verify Port to Application Mapping Use the Privileged EXEC Mode command **show ip port-map** to display the Port to Application Mapping (PAM) information. This command displays the port mapping information at the firewall, including the system-defined and user-defined information. Include the application name to display only the entries for that application. Include the port number to display only the entries for that port. The syntax is

Rtr1#show ip port-map [*appl-name* | port *port-num*]

This command was introduced in IOS 12.0(5)T.

The following example shows the port mapping information for FTP services:

```
Rtr1#show ip port-map ftp
Default mapping: ftp             port 21                  system defined
Host specific:   ftp             port 1250   in list 1    user defined
```

Define Inspection Rules

To define a set of inspection rules, use the **ip inspect name** command with the same inspection name for each protocol you want CBAC to inspect. Each set of inspection rules must have a unique inspection name, which shouldn't exceed the 16-character limit. Typically, if inspection is configured for a protocol, return traffic entering the internal

network is only permitted if the packets are part of a valid, existing session for which state information is being maintained.

Define either one or two sets of rules per interface. A single set can be used to examine both inbound and outbound traffic, or you can define two sets: one for outbound traffic and one for inbound traffic. To define a single set of inspection rules, configure inspection for all the desired application-layer protocols and for TCP or UDP, as desired. This combination of TCP, UDP, and application-layer protocols join together to form a single set of inspection rules with a unique name.

In general, to configure CBAC inspection for a protocol, packets for that protocol should be permitted to exit the firewall by configuring the appropriate ACL if necessary. Packets for that protocol will only be allowed back in through the firewall if they belong to a valid existing session. Without outbound packets, there can be no existing session. Each protocol packet is inspected to maintain information about the session state.

Use the Global Configuration Mode command **ip inspect name** to define a set of inspection rules. Use the no form of this command to remove the inspection rule for a protocol or to remove the entire set of inspection rules. The syntax is

```
Rtr1(config)#ip inspect name inspection-name protocol [alert {on | off}] [audit-trail
   {on | off}] [timeout seconds]
Rtr1(config)#no ip inspect name inspection-name protocol (removes a protocol inspection rule)
Rtr1(config)#no ip inspect name (removes the entire set of inspection rules)
```

inspection-name	Name of inspection rules set. Up to 16 characters. To add a protocol to an existing set of rules, use the same *inspection-name*.
protocol	Application protocol name is supported by CBAC.
alert {on | off}	(Optional) Each inspected protocol can override the default settings created by the global **ip inspect alert-off** command.
audit-trail {on | off}	(Optional) Each inspected protocol can override the default settings created by the global **ip inspect audit-trail** command.
timeout *seconds*	(Optional) Each inspected protocol can override the default global TCP or UDP idle timeouts, but won't override the global DNS timeout.

This command was introduced in IOS 11.2P. Configurable alert and audit-trail, IP fragmentation checking, and NetShow protocol support were added in 12.0(5)T.

TCP and UDP Inspection Configuring TCP and UDP inspection will permit *any* TCP and UDP packets to enter the internal network through the firewall, even if the application-layer protocol isn't configured to be inspected. TCP and UDP inspection doesn't recognize application-specific commands or processes and might not permit all returning packets for that application. Any returning application packets with port numbers different than the exiting packet will be blocked.

Application-layer protocol inspection takes precedence over the TCP or UDP packet inspection. For example, with FTP inspection, all control channel information will be recorded in the state table, and all FTP traffic is permitted back through the firewall if the control channel information is valid for the state of the FTP session. Any TCP inspection in this case becomes irrelevant.

Generic TCP and UDP inspection is much like reflexive ACLs, in that the packets entering the network must exactly match an existing session. This match is based on having the same source/destination addresses and source/destination port numbers as the existing packet, except reversed. Otherwise, the inbound packets are blocked at the interface.

The following example causes the software to inspect TCP sessions and UDP sessions, and to specifically allow FTP and Real Audio traffic back through the firewall for existing sessions only. UDP traffic has the audit-trail feature on, while the FTP timeout is changed to three minutes.

```
Rtr1(config)#ip inspect name letusin tcp
Rtr1(config)#ip inspect name letusin udp audit-trail on
Rtr1(config)#ip inspect name letusin ftp timeout 180
Rtr1(config)#ip inspect name letusin realaudio-video
```

H.323 Inspection If you want CBAC inspection to work with Microsoft NetMeeting 2.0 traffic, an H.323 application-layer protocol, it's necessary also to configure inspection for TCP because NetMeeting 2.0 uses an additional TCP channel not defined in the H.323 specification. No special configuration options exist for H.323

To add Microsoft NetMeeting 2.0 inspection, it's necessary to add only the following entry because TCP inspection is on in the initial example:

```
Rtr1(config)#ip inspect name letusin h323
```

HTTP Inspection (Java Blocking) With the proliferation of Java applets available on the Internet, protecting networks from malicious applets has become a major concern to network managers. The Java blocking feature can be configured either to filter or to completely deny access to Java applets.

Java inspection enables Java applet filtering at the firewall. *Java applet filtering* distinguishes between trusted and untrusted applets by relying on a list of external sites you designate as "friendly." If an applet is from a friendly site, the firewall allows the applet through. If the applet isn't from a friendly site, the applet is blocked. Alternatively, you could permit applets from all sites, except sites specifically designated as "hostile."

Java blocking uses only numbered standard access lists to define friendly and hostile external sites. ACL "permit" statements define traffic from friendly sites, while ACL "deny statements" define traffic from hostile sites. If Java blocking is defined using an ACL number and no matching ACL exists, all Java applets will be blocked.

NOTE CBAC can't detect or block encapsulated or "zipped" Java applets. Applets in .zip or .jar format aren't blocked at the firewall. CBAC also doesn't detect or block applets loaded via FTP, gopher, or HTTP on a nonstandard port.

Use the no form of the command to remove Java blocking from the inspection set. The syntax is

Rtr1(config)#ip inspect name *inspection-name* [http [java-list *acl#*]] [alert {on | off}] [audit-trail {on | off}] [timeout *seconds*]

Rtr1(config)#no ip inspect name *inspection-name* http

http	Specifies Java applet blocking
java-list *acl#*	A "numbered standard" ACL used to define friendly sites.

The next example would add blocking of all Java applets to the inspection set.

```
Rtr1(config)#ip inspect name letusin http
```

If the next example was used instead, all Java applets would be blocked except those originating from the server 192.168.0.10.

```
Rtr1(config)#access-list 15 permit 192.168.0.10
Rtr1(config)#ip inspect name letusin http java-list 15
```

RPC Inspection *RPC inspection* allows the specification based on program numbers. To define multiple program numbers, create multiple entries for RPC inspection, each with a different program number. If a program number is specified, all traffic for that program number is permitted. Any program number that isn't specified will have all traffic for that program number blocked.

Use the no form of the command to remove RPC inspection from the inspection set. The syntax is

Rtr1(config)#ip inspect name *inspection-name* [rpc program-number *number* [wait-time *minutes*]] [alert {on | off}] [audit-trail {on | off}] [timeout *seconds*]

Rtr1(config)#no ip inspect name *inspection-name* rpc

rpc program-number *number*	The RPC program number to permit.
wait-time *minutes*	(Optional) Number of minutes to keep a small hole in the firewall to allow subsequent RPC connections from the same source address, and to the same destination address and port. Default is zero minutes.

The following example adds two RPC program numbers to the inspection set.

```
Rtr1(config)#ip inspect name letusin rpc program-number 100002
Rtr1(config)#ip inspect name letusin rpc program-number 100010
```

Fragment Inspection (DoS Attacks) CBAC inspection rules can help protect hosts against DoS attacks involving fragmented IP packets. Even if the firewall prevents an attacker from completing a connection to a protected host, the attacker might still be able to disrupt services provided by that host in two ways: by sending a large number of "noninitial" IP fragments or by sending a complete set of fragmented packets through a router that has an ACL that filters out the first fragment of the packet. Either way, the remaining fragments can tie up host resources as it tries in vain to reassemble the incomplete packets.

Using fragmentation inspection, the CBAC maintains an interfragment state for IP traffic. Noninitial fragments are discarded unless the corresponding initial fragment has already been permitted to pass through the firewall. Noninitial fragments received before the corresponding initial fragments are discarded. Because many circumstances exist that can cause out-of-order delivery of legitimate fragments, fragmentation inspection can have a severe performance implications. For this reason, fragment inspection is off by default.

Use the no form of the command to remove fragment inspection from the inspection set. The syntax is

Rtr1(config)#ip inspect name *inspection-name* fragment [max *number* timeout *seconds*]

Rtr1(config)#no ip inspect name *inspection-name* fragment

fragment	Specifies fragment inspection for the named inspection set.
max *number*	This is the maximum number of unassembled packets for which state structure information memory space is allocated by the IOS. Unassembled packets are those arriving before the initial packet for a session. Range is 50 through 10,000 entries, with 256 being the default. Because memory is allocated for these state structures, the larger the number, the more memory is reserved, possibly causing memory resources to be exhausted.
timeout *seconds* (fragmentation)	The number of seconds a packet state structure remains active. When the timeout value expires, the router drops the unassembled packet, freeing that structure for use by another packet. The default timeout value is one second. If timeout is set to more than one second, it's automatically adjusted by the IOS if the number of "free" state structures dips below certain thresholds. If the number of free states is less than 32, the timeout is divided by 2. If the number of free states is less than 16, the timeout is set to one second.

Even if the system is under heavy attack from fragmented packets, legitimate fragmented traffic, if any, will still get some fraction of the firewall's fragment state resources and, more important, regular unfragmented traffic can pass through the firewall unimpeded.

The following example adds fragment checking to that software inspection set. The firewall will allocate 100 state structures, meaning 100 initial fragments for 100 different packets sent through the router can be stored before new ones must be discarded. The timeout value for dropping unassembled packets is set to four seconds.

```
Rtr1(config)#ip inspect name letusin fragment max 100 timeout 4
```

SMTP Inspection SMTP inspection searches SMTP commands for illegal commands. Packets with illegal commands are dropped, and the SMTP session will hang and eventually time out. An illegal command is any command except for the following legal commands: DATA, EXPN, HELO, HELP, MAIL, NOOP, QUIT, RCPT, RSET, SAML, SEND, SOML, and VRFY.

Use the no form of the command to remove fragment inspection from the inspection set. The syntax is

Rtr1(config)#ip inspect name *inspection-name* smtp [alert {on | off}] [audit-trail {on | off}] [timeout *seconds*]
Rtr1(config)#no ip inspect name *inspection-name* smtp

The following example adds SMTP inspection with default audit-trail and alert setting.

```
Rtr1(config)#ip inspect name letusin smtp
```


Apply Inspection Rules and ACLs to an Interface

It's time to apply the inspection set to a router interface. If the interface connects to the external network, apply the inspection rules to outbound traffic. If the interface connects to the internal network, apply the inspection rules to inbound traffic. By applying the rules to outbound traffic, then returning inbound packets will be permitted if they belong to a valid connection with existing state information. This connection state must be initiated with an outbound packet.

Normally, you apply only one inspection rule per interface. The only exception is if you want to enable CBAC in two directions between two departments or partner networks. To configure CBAC in both directions on a single firewall interface, apply two rules, one for each direction.

Use the interface configuration **ip inspect** command to apply a set of inspection rules to an interface. Use the no form of the command to remove the set of rules from the interface. The syntax is

Rtr1(config-if)#ip inspect *inspection-name* {in | out}
Rtr1(config-if)#no ip inspect *inspection-name* {in | out}

in	Applies the inspection rules to inbound traffic (relative to the router)
out	Applies the inspection rules to outbound traffic (relative to the router)

This command was introduced in IOS 11.2.

This example applies a set of inspection rules named letusin to an external interface's outbound traffic. The inspection set is the one created in the example and isn't repeated here. Inbound IP traffic is permitted only if the traffic is part of an existing session created by the letusin inspections. All other inbound traffic will be denied unless it's the specific ICMP traffic allowed by the inbound ACL 101. The outbound ACL 150 specifies the traffic that can leave the network.

```
Rtr1(config)#access-list 101 deny udp any any
Rtr1(config)#access-list 101 deny tcp any any
Rtr1(config)#access-list 101 permit icmp any any echo-reply
Rtr1(config)#access-list 101 permit icmp any any packet-too-big
Rtr1(config)#access-list 101 permit icmp any any time-exceeded
Rtr1(config)#access-list 101 permit icmp any any traceroute
Rtr1(config)#access-list 101 permit icmp any any unreachable
Rtr1(config)#access-list 150 permit tcp any any
Rtr1(config)#access-list 150 permit udp any any
Rtr1(config)#access-list 150 permit icmp any any
Rtr1(config)#interface serial0
Rtr1(config-if)#ip inspect letusin out
Rtr1(config-if)#ip access-group 150 out
Rtr1(config-if)#ip access-group 101 in
```

Select the Interface

When implementing CBAC, deciding whether to configure CBAC on an internal or an external interface of your firewall is necessary. In firewall terminology, *internal* refers to the interface(s) where sessions must originate for the traffic to be permitted through the firewall. *External* refers to the interface(s) where sessions can't originate. Sessions originating from the external side will be blocked.

Figure 6-2 shows a simple topology common on perimeter routers or networks without a protected DMZ. CBAC would be configured on the external interface—Serial 0—to prevent unwanted traffic from entering the firewall and the internal network.

Figure 6-3 shows a common topology that includes a protected DMZ for shared servers. CBAC is configured for the internal interface—Ethernet 1—allowing external traffic free access to the DMZ. At the same time, external traffic is prevented from entering the internal network, unless it's part of a session initiated from within the internal network.

To provide firewall security between two departments or two partner networks, configure CBAC in two directions. CBAC can be configured in two directions at one or more interfaces. First, configure CBAC in one direction using the appropriate internal and external interface designations. Then, configure CBAC in the other direction with the interface designations swapped.

CBAC on an External Interface Create an outbound standard or extended IP access list and apply it to the external interface. This ACL permits all packets to be allowed to exit the network, including any packets you want inspected by CBAC.

Create an inbound extended IP access list and apply it to the external interface. This ACL denies any traffic to be inspected by CBAC. When an outbound packet triggers CBAC, a temporary opening is created in this inbound ACL to permit only traffic from a valid existing session. If the inbound ACL is configured to permit the desired traffic, then CBAC creates unnecessary openings in the firewall for packets that would be allowed anyway.

CBAC on an Internal Interface Create an inbound standard or extended IP access list and apply it to the internal interface. This ACL permits all packets to be allowed to exit the protected network, including any packets you want inspected by CBAC.

Create an outbound extended IP access list and apply it to the internal interface. This ACL denies any traffic to be inspected by CBAC. When an inbound packet triggers CBAC, a temporary opening is created in this outbound ACL to permit only traffic from a valid existing session.

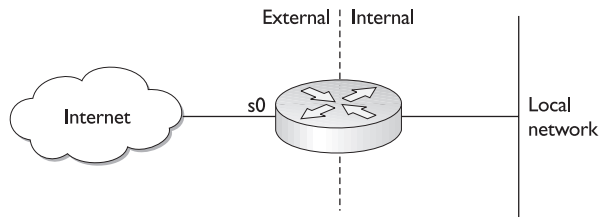

Figure 6-2 Simple perimeter router with internal and external interface

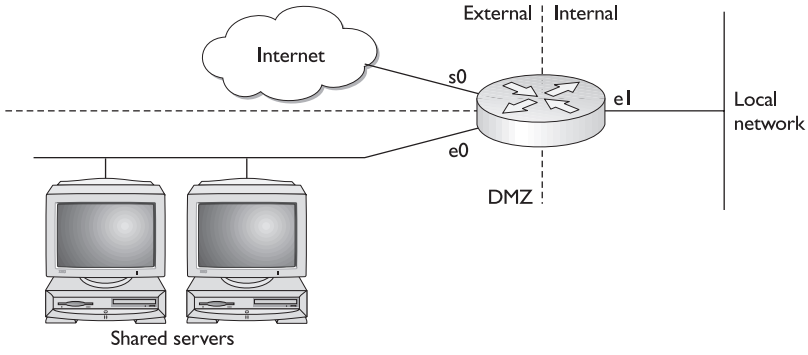

Figure 6-3 Simple firewall design with a protected DMZ

Test and Verify

Use the Privilege EXEC Mode command **show ip inspect** with its various options to view CBAC configuration and session information. The syntax is

Rtr1#show ip inspect {name *inspection-name* | config | interfaces | session [detail] | all}

config	Shows the complete CBAC inspection configuration.
interfaces	Shows interface configuration for applied inspection rules and ACLs.
session [detail]	Shows existing CBAC inspection sessions. (Optional) keyword **detail** displays more information about the sessions.
all	Shows all CBAC configuration and all sessions being tracked and inspected by CBAC.

This command was introduced in IOS 11.2 P.

The following example shows sample output for the **show ip inspect name testinspect** command, where testinspect is the inspection rule set. The output shows the protocols as inspected by CBAC and the idle timeouts for each.

```
Rtr1#show ip inspect name testinspect
Inspection Rule Configuration
 Inspection name testinspect
    tcp timeout 3600
    udp timeout 90
    ftp timeout 2400
```

The following, which is sample output for the **show ip inspect config** command, shows CBAC configuration, including global timeouts, thresholds, and inspection rules:

```
Rtr1#show ip inspect config
Session audit trail is disabled
one-minute (sampling period) thresholds are [500:750] connections
max-incomplete sessions thresholds are [500:750]
max-incomplete tcp connections per host is 50. Block-time 0 minute.
tcp synwait-time is 30 sec -- tcp finwait-time is 5 sec
tcp idle-time is 3600 sec -- udp idle-time is 90 sec
dns-timeout is 5 sec
Inspection Rule Configuration
 Inspection name testinspect
```

```
    tcp timeout 3600
    udp timeout 90
    ftp timeout 2400
```

This output demonstrates the **show ip inspect interfaces** command:

```
Rtr1#show ip inspect interfaces
Interface Configuration
 Interface Ethernet0
  Inbound inspection rule is testinspect
    tcp timeout 3600
    udp timeout 90
    ftp timeout 2400
  Outgoing inspection rule is not set
  Inbound access list is not set
  Outgoing access list is not set
```

This output demonstrates the **show ip inspect sessions** command, indicating an FTP session and showing the source and destination addresses and port numbers (separated by colons):

```
Rtr1#show ip inspect sessions
Established Sessions
 Session 25A3318 (192.168.0.17:20)=>(192.168.1.9:47091) ftp-data SIS_OPEN
 Session 25A6E1C (192.168.1.9:47098)=>(192.168.0.17:21) ftp SIS_OPEN
```

This output demonstrates the **show ip inspect sessions detail** command, including times, number of bytes sent, and which access lists are applied:

```
Rtr1#show ip inspect sessions detail
Established Sessions
 Session 25A335C (192.168.0.17:20)=>(192.168.1.9:47091) ftp-data SIS_OPEN
   Created 00:00:07, Last heard 00:00:00
   Bytes sent (initiator:responder) [0:3416064] acl created 1
   Inbound access-list 100 applied to interface Ethernet0
 Session 25A6E1C (192.168.1.9:47098)=>(192.168.0.17:21) ftp SIS_OPEN
   Created 00:01:34, Last heard 00:00:07
   Bytes sent (initiator:responder) [196:616] acl created 1
   Inbound access-list 100 applied to interface Ethernet0
```

IOS Firewall Management

You can administer the IOS Firewall features in two methods: the traditional command-line interface (CLI) and Cisco ConfigMaker.

Command Line Interface

For users who are familiar with the CLI, its primary advantage is that it can be used across virtually the entire product line. As with any complex skill, the more familiar you become with the CLI, the easier it becomes. Strengths of the CLI include

- Similarity across product lines.
- Functional online help to assist with complex tasks.

- Virtually every feature can be accessed or configured.

On the possible downside, the initial exposure to CLI can be intimidating and confusing. The additions and changes in commands with new releases can be a challenge to keep up with.

NOTE A technician I met early in my career told me the Cisco CLI was the way "to keep the fools that shouldn't be touching routers away." While I've never found that in any Cisco document, I suspect the result might be on the mark, even if the intent is not.

ConfigMaker

Cisco ConfigMaker is an easy-to-use Microsoft Windows–based software tool designed to configure smaller network Cisco routers, switches, hubs, and other network devices. Advanced features include implementing security policies and managing the Cisco IOS Firewall quickly and efficiently with GUI-based management. Cisco ConfigMaker 2.1 and later versions include a Security Wizard for step-by-step guidance for quick configuration

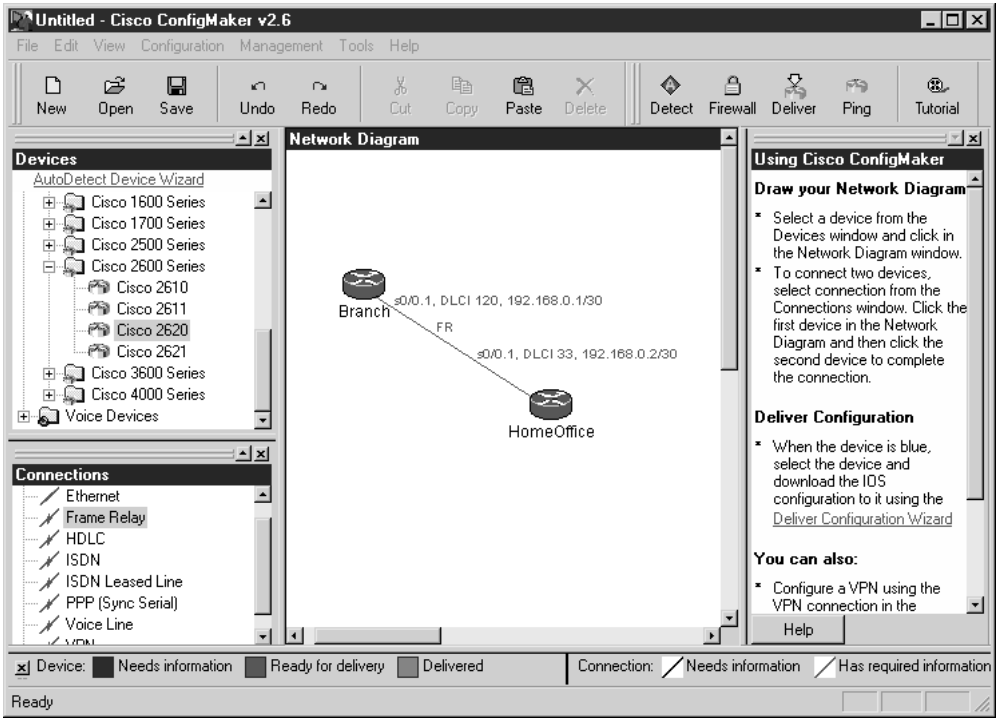

Figure 6-4 Cisco ConfigMaker tool for network design and implementation

of security policy for the Cisco IOS Firewall. They also support NAT and IPSec configuration. Figure 6-4 shows the ConfigMaker interface.

ConfigMaker prompts, wizards, and help screens guide users through the setup process. Online WAN configuration worksheets list important information the network administrator must obtain from the Internet service provider (ISP) or WAN service provider before configuring the devices. The program includes a multimedia movie tutorial.

Configurations can be downloaded or uploaded to the appropriate devices from a standalone PC or over the network.

Configuration wizards include the following:

Task wizards	Prompt user through processes and for key information.
AutoDetect Device Wizard	Can automatically detect and identify the configuration of any supported Cisco device on the network and any WAN or voice interfaces installed in a modular router.
Address Network Wizard	Quick-and-easy completion of complex network addressing procedures by using supplied address ranges and assigning the addresses to selected devices or the entire network.
Security Wizard	Enable policy-based configurations for efficient setup of the Cisco IOS Firewall feature set for Cisco 1600 and 2500 series routers.
Deliver Configuration Wizard	Once the configuration is complete, the wizard can automatically download it through the console port or over the network.

ConfigMaker supports routers up through the 4000 series, but little or no support exists for the Catalyst switches.

To download ConfigMaker, go to the Cisco web site, http://www.cisco.com/warp/public/cc/pd/nemnsw/cm/index.shtml, or do a search for ConfigMaker. No CCO account is required, but you're asked to supply some information before downloading.

STUDY TIP For those of you attempting to certify without access to devices, look at this product, which will at least challenge and assist you in working with many of the technologies covered in the exams.

Chapter Review

The Cisco IOS Firewall feature set, a part of the Cisco Secure system, is made up of the following four interrelated features:

- Cisco IOS Firewall Context-Based Access Control
- Port to Application Mapping
- Cisco IOS Firewall Intrusion Detection System
- Cisco IOS Firewall Authentication Proxy

Context-Based Access Control (CBAC) allows the firewall to take the access list type filter to a much higher level. Whereas ACLs are limited to Layer 3 and Layer 4 information

for filtering, the CBAC can incorporate knowledge of the operation of supported application protocols to make decisions. This allows for more flexibility in the number of and port addressing of communications channels through the firewall.

Other CBAC features allow for monitoring and reacting to common DoS attacks and e-mail attacks involving unauthorized SMTP commands.

Port to Application Mapping (PAM) allows the flexibility of incorporating nonstandard TCP and UDP port numbers in the secure openings through the firewall.

Questions

1. True or False. IPSec is a part of the Cisco IOS Firewall feature set.
   - **A.** True
   - **B.** False
2. True or False. The Cisco IOS Firewall feature set is implemented on all Cisco router series.
   - **A.** True
   - **B.** False
3. Which of the following IOS features is *not* part of the Firewall feature set?
   - **A.** Intrusion detection
   - **B.** Context-Based Access Control (CBAC)
   - **C.** AAA
   - **D.** Java blocking
4. True or False. CBAC can incorporate application layer information in its filtering.
   - **A.** True
   - **B.** False
5. In the following command, what does the 30 represent?
   Rtr1(config)#ip inspect tcp idle-time 30
   - **A.** Minutes
   - **B.** Packets
   - **C.** Seconds
   - **D.** Hours
6. True or False. CBAC can filter TCP, UDP, and ICMP traffic.
   - **A.** True
   - **B.** False
7. The memory required for each CBAC connection is what?
   - **A.** 600 bits
   - **B.** 600 bytes

C. 600K

D. Varies with the data

8. Which of the following is *not* a step in configuring CBAC?

A. Set audit trails and alerts.

B. Set global timeouts and thresholds.

C. Define inspection rules.

D. Remove all nonstandard Port-to-Application Mapping.

E. Apply inspection rules and ACLs.

9. Which of the following is a DoS protective measure?

A. RPC inspection

B. Fragment inspection

C. SMTP inspection

D. HTTP inspection

10. Which of the following defines the number of seconds the software will wait for a TCP session to reach the established state before dropping the session?

A. Rtr1(config)#ip inspect tcp synwait-time 20

B. Rtr1(config-if)#ip inspect tcp synwait-time 20

C. Rtr1(config)#ip inspect tcp finwait-time 20

D. Rtr1(config-if)#ip inspect tcp finwait-time 20

11. In the following command, what does the number 800 represent?
Rtr1(config)#ip inspect max-incomplete high 800

A. Seconds

B. Minutes

C. Half-open TCP sessions

D. DNS-name lookup session

12. What does the following command do?
Rtr1(config)#ip port-map realaudio port 21

A. Assigns port 21 to be used by Real Audio.

B. States a preference for port 21 to be used by Real Audio.

C. The command will fail because CBAC doesn't support Real Audio.

D. The command will fail because port 21 is reserved for FTP.

13. True or False. ConfigMaker is an alternative for configuring Firewall features.

A. True

B. False

14. Which two commands might be useful against DoS attacks?
    A. Maximum Incomplete Sessions High/Low Threshold
    B. UDP Session Inactivity Timer
    C. TCP Session Termination Timer
    D. One Minute Incomplete Sessions High/Low Threshold
15. Which statement is *not* true about CBAC?
    A. Only IP TCP and UDP traffic is inspected by CBAC.
    B. CBAC doesn't normally protect against attacks from within the protected network.
    C. CBAC and reflexive ACLs work well together.
    D. CBAC can't inspect in-transit IPSec traffic.

Answers

1. **B**. False. They're used together often, but they're separate feature sets.
2. **B**. False. It is implemented only on the Cisco 800, uBR900, 1400, 1600, 1700, 2500, 2600, 3600, 7100, 7200, and 7500 and RSM series routers.
3. **C**. AAA It is in the regular IOS feature set.
4. **A**. True
5. **C**. Seconds
6. **A**. False. It's limited to TCP and UDP traffic.
7. **B**. 600 bytes
8. **D**. Remove all nonstandard Port-to-Application Mapping.
9. **B**. Fragment inspection
10. **A**. Rtr1(config)#ip inspect tcp synwait-time 20
11. **C**. Half-open TCP session
12. **D**. The command will fail because port 21 is reserved for FTP.
13. **A**. True
14. **A**. Maximum Incomplete Sessions High/Low Threshold and **D**. One Minute Incomplete Sessions High/Low Threshold
15. **C**. CBAC and reflexive ACLs work well together

CHAPTER 7

IOS Firewall—Intrusion Detection System

In this chapter, you will learn how to:

- Use Cisco IOS Firewall IDS
- Initialize the IOS Firewall IDS
- Configure, disable, and exclude signatures
- Create and apply IDS audit rules

With the growing awareness of, and commitment to, network security as the way to secure business transactions—both internally and over the Internet—more and more businesses are recognizing that security concepts and features must be integrated into the network design and infrastructure.

According to Cisco advertising, 80 percent of the Internet backbone routers run Cisco IOS software, which makes it the most fundamental component of today's network infrastructure. The Cisco IOS software-based security features, combined with the wide range of products supported, makes it a logical solution for any organization's end-to-end Internet, intranet, and remote access network security requirements.

This chapter looks at the Cisco IOS intrusion detection system (IDS) features that are now a part of the IOS Firewall feature set on many router platforms. In the last four chapters of this book, you willlearn about the Cisco Secure IDS appliances that make up the first team in Cisco's intrusion detection strategy.

STUDY TIP The material in this chapter is only part of the exam objectives for the Managing Cisco Network Security Exam (MCNS 640-100), but reviewing it when you prepare for the IDS with Policy Manager Exam (IDS 9E0-572) will help you see how the products can be compared and used together.

Intrusion Detection System (IDS)

An intrusion detection system (IDS) device inspects all network activity passing through it and identifies suspicious patterns that might indicate a network or device attack. IDS can be categorized three ways.

Misuse detection vs. anomaly detection	With *misuse detection,* the IDS analyzes the data stream and compares it to databases of known attack signatures. The IDS can only look at previously documented, specific attacks. This is similar to the process many virus-scanning programs use, and, like those programs, the protection is only as good as the database of attack signatures available. In *anomaly detection* (or *profile detection*), the administrator defines the baseline profile of the "normal" network traffic characteristics. The IDS monitors network segments and compares their state to the baseline, looking for deviations or anomalies.
Network-based vs. host-based systems	With network-based systems (NIDS), such the Cisco IDS 4200 Series appliances, the individual packets flowing through a network are analyzed by dedicated devices. As specialty devices, NIDS have been optimized to detect suspicious packets, which, typically, are overlooked by a firewall's filtering rules. A *host-based system* is IDS software configured on key resources such as servers, routers, or switches to examine activity on that device.
Passive system vs. reactive system	When a *passive system* detects a potential security threat, it logs the pertinent information and signals an alert message. A *reactive system* would log the event and send out the alert but, at the same time, the IDS responds to block traffic from the suspected malicious source.

This capability to react to an evolving threat allows IDS implementations to provide protection beyond that of a firewall without IDS features. Furthermore, firewalls tend to be configured to look for bad traffic coming into the network, but remain oblivious to internal traffic. IDS can simultaneously protect the network from internal and external threats.

IOS Firewall Intrusion Detection System

The IOS Firewall IDS feature acts as an inline sensor, watching packets and sessions as they flow through the router, scanning each for pattern matches to any known IDS signatures. When packets in a session match a signature, the IDS system can be configured to do the following:

- Send an alarm to a Syslog server and/or to a Cisco Secure IDS Director centralized management device (formerly the NetRanger system).
- Discard the packet.
- Reset the TCP connection.

While enabling both the firewall and intrusion detection features of the CBAC security engine to support a network security policy is preferable, each of these features can be enabled independently and on different router interfaces.

Devices Supporting the IOS Firewall IDS Features

Cisco IOS software-based intrusion detection is available on the Cisco uBR900, 1720, 2600, 3600, 7100, 7200, and 7500 series routers, and the RSM for Catalyst 5000 switches.

The IDS technology has been included with the firewall feature set since version 12.1 or 12.2 of the IOS, depending on the device platform. More models might be supported later. You must choose a feature set that contains the firewall and IDS features when you order or upgrade the device IOS. For this reason, a router IOS that supports the firewall features doesn't necessarily mean that the IDS technology is included. Figure 7-1 shows a sample of the Cisco IOS Upgrade Planner from the Cisco web site for the model 1720 router. Notice that the firewall features (FW) and the intrusion features (IDS) are available in combination with various protocols and features, such as IPX. The IOS releases with firewall features without IDS are typically capped at version 12.0(4). Memory and flash requirements often increase, as does the cost of the IOS when additional features are added.

The following IOS image names are the latest 1720 release for IP with the firewall features only, while the second listing includes the IDS features. The last two entries are for a 2600 and a 7000 with RSP, respectively, with each having Enterprise and IPSec features. The *o* indicates firewall features, while the o3 indicates firewall and IDS features.

```
c1700-oy-mz.120-5.T1.bin
c1700-o3y-mz.122-12a.bin
c2600-jk8o3s-mz.122-12.bin
rsp-jk8o3sv-mz.122-10b.bin
```

[12.2] [12.1] [12.0] [11.3] [11.2] [11.1] [11.0] [10.3] [All]

Your Selected Platform	Select Release		Select Software Feature
1720	Major Release Updates		IP IP PLUS
	12.2.12a	LD	IP PLUS 40
	12.2.12	LD	IP PLUS 56
	12.2.10b	ED	IP PLUS IPSEC 56
	12.2.10a	LD	IP/ADSL
	12.2.10	LD	IP/ADSL PLUS
	12.2.7c	LD	IP/ADSL PLUS IPSEC 56
	12.2.7b	LD	IP/ADSL/FW/IDS PLUS IPSEC 56
	12.2.7a	LD	IP/ADSL/IPX/AT/IBM PLUS
	12.2.7	LD	IP/ADSL/IPX/AT/IBM/FW/IDS PLUS IPSEC 56
	12.2.6h	LD	IP/ADSL/IPX/AT/IBM/VOICE/FW/IDS PLUS IPSEC 56
	12.2.6g	LD	IP/ADSL/IPX/FW/IDS PLUS
	12.2.6f	LD	IP/ADSL/IPX/VOICE/FW/IDS PLUS
	12.2.6e	LD	IP/ADSL/VOICE PLUS
	12.2.6d	LD	IP/ADSL/VOICE PLUS IPSEC 56
	12.2.6c	LD	IP/ADSL/VOICE/FW/IDS PLUS
	12.2.6b	LD	IP/ADSL/VOICE/FW/IDS PLUS IPSEC 56 IP/FW IP/FW PLUS IPSEC 56 IP/FW/IDS IP/FW/IDS PLUS IPSEC 56

Figure 7-1 Cisco CCO IOS Upgrade Planner showing IDS feature sets

Cisco IDS Attack Signatures

The most recent Cisco IOS Firewall IDS uses 59 *attack signatures*, representing a broad cross section of intrusion-detection signatures, which identify severe breaches of security and the most common network attacks and information-gathering scans. Unlike virus protection software, IDS signatures aren't updated periodically by the system. Currently, the number of signatures only changes if a version upgrade contains any additions or deletions. The Cisco IOS Firewall IDS signatures are categorized into four types:

- Info Atomic
- Info Compound
- Attack Atomic
- Attack Compound

To understand these categories better, the signature keywords are as follows:

Info	Information-gathering activity, such as a port sweep.
Attack	Attacks attempted into the protected network, such as denial of service (DoS) attempts or the execution of illegal commands during an FTP session.
Atomic	Simple patterns, such as an attempt to access a specific port on a specific host.
Compound	Complex patterns, such as a sequence of operations distributed across multiple hosts over an arbitrary period of time.

The intrusion detection signatures included in the Cisco IOS Firewall were chosen from a broad cross section of intrusion detection signatures as representative of the most common network attacks and information-gathering scans. A small sample of the signatures is included in the following table.

Sig ID	Signature Name	Sig Type	Description
1100	IP Fragment Attack	Attack, Atomic	Triggers when any IP datagram is received with the more fragments flag set to 1 or if an offset is indicated in the offset field.
2003	ICMP Redirect	Info, Atomic	Triggers when an IP datagram is received with the protocol field in the IP header set to 1 (ICMP) and the type field in the ICMP header set to 5 (Redirect).
2154	Ping of Death Attack	Attack, Atomic	Triggers when an IP datagram is received with the protocol field in the IP header set to 1 (ICMP), the Last Fragment bit is set, and (IP offset × 8) + (IP data length) > 65535. In other words, the IP offset plus the rest of the packet is greater than the maximum size for an IP packet.
3050	Half-open SYN Attack/SYN Flood	Attack, Compound	Triggers when multiple TCP sessions were improperly initiated on any of several well-known service ports. Detection of this signature is currently limited to FTP, Telnet, HTTP, and e-mail servers (TCP ports 21, 23, 80, and 25, respectively).

For a complete listing and more information on IDS signatures, go to http://www.cisco.com and do a search for Cisco IOS Firewall IDS Signature List, and then look through the resulting choices for the same phrase in bold.

False Positives

The signatures integrated into the IOS software monitor for severe breaches of security. They are used to watch for those data flows you wouldn't normally expect to see in an operating network. A *false positive* is an erroneous report from an IDS, indicating it detected a potentially malicious pattern. The pattern appears to matches a signature but, in fact, is a valid and acceptable transmission. Any intrusion detection technology can and does report false positives. This can be looked at as erring on the side of security or caution, but it can also block necessary traffic.

The IOS-based intrusion-detection features were developed with flexibility in mind, so individual signatures could be disabled in case of false positives.

Cisco Secure IDS Director Support

The Cisco IOS Firewall intrusion detection capabilities have an enhanced reporting mechanism that permits logging to the Cisco Secure IDS Director console in addition to a Syslog server to provide a consistent view of all intrusion detection sensors throughout a network. Administrators can deploy the IOS Firewall IDS to complement their existing IDS systems. This allows IDS protection to be deployed to areas that might not support a Cisco Secure IDS Sensor. The IOS Firewall IDS signature features can be deployed alongside or independent of other Cisco IOS Firewall features.

The Cisco Secure IDS consists of three components:

- Sensor
- Director
- Post Office

Cisco *Secure IDS Sensors,* dedicated high-speed network appliances, analyze the content and context of individual packets to determine if traffic constitutes a threat. If a data stream appears unauthorized or suspicious, such as a ping sweep or a SATAN attack, the sensors can detect the policy violation in real-time, forward alarms to a Cisco Secure IDS Director management console, and remove the offender from the network.

The *Cisco Secure IDS Director* is a software-based management system that can monitor the activity of multiple Cisco Secure IDS Sensors located on local or remote network segments. Events are sent to the Director by an IDS Sensor or an IDSM that detects a security violation. The *smid daemon* on the Director interprets this event information and passes it to the *nrdirmap daemon,* which is responsible for displaying this information on the Director's maps.

Depending on the severity of an alarm, the alarm icon displays in different colors: red for severe, yellow for moderate, green otherwise. The Cisco Secure IDS Director is an application that runs on either HP or Sun Solaris UNIX workstations. The Director is covered in detail in the final chapter of this book.

The *Cisco Secure IDS Post Office Protocol* is the communication backbone that allows Cisco Secure IDS services and hosts to communicate with each other. All communication is supported by a proprietary, connection-based protocol that can switch between alternative routes to maintain point-to-point connections.

NOTE Version 2.2.2 of the Cisco Secure IDS Director replaces the name "Cisco Secure IDS Post Office Protocol" with "Communication Service." The version 2.2.2 Installation program replaces the nr.postofficed daemon.

Performance Implications

The impact on performance of the IOS intrusion detection features depends on the number of signatures enabled, the router platform, the overall traffic level on the router, and other features enabled on the router—such as encryption, CBAC, and so on. Because the router is working as a security device, no packets are allowed to bypass the security mechanisms. The IDS process in the IOS Firewall router acts as a filter in the packet path, thus searching each packet for signature matches. Because the entire packet is searched in many cases, the state information, and even the application state and awareness, must be maintained by the router.

IOS IDS vs. Cisco Secure IDS

The *Cisco Secure IDS Sensor* is a dedicated appliance that passively monitors the network and reacts to suspicious signatures indicating potential malicious activity. The IDS device can be configured to block this activity. The Cisco IOS-based IDS system is an integral component of the IOS software and, therefore, lies directly in the packet path, rather than being a separate appliance. The IOS IDS technology expands the perimeter protection capabilities offered by the IOS Firewall by being able to take appropriate actions on packets and data flows that appear to be malicious network activity or to violate the organization security policy.

Other differences include the following:

- The Cisco Secure IDS Sensor device processing speed is faster than that of the IOS feature because of the shared resources design of the integrated router-based solution.
- The Cisco Secure IDS Sensor device includes more signatures than the intrusion detection feature on the Cisco IOS Firewall.
- The Cisco Secure IDS Sensor device can reconfigure a Cisco router by dynamically adding an access control list to block intruders, but the IOS version can't do this.
- The Cisco Secure IDS Sensor device can be managed remotely by the Cisco Secure IDS Director. While the IOS version can send output to the IDS Director, it doesn't take instructions from it.

When to Choose the Cisco IOS Firewall IDS Features

Because the IOS Firewall IDS supports intrusion detection features for a wide range of Cisco router platforms, it can make a powerful addition to any network perimeter. The features can be especially useful in locations where a router is being deployed to provide additional security between network segments, such as between the organization and a partner site.

The Firewall IDS features can provide increased protection between intranet connections, such as branch-office connections to the corporate office or even providing additional security for an internal department like an R&D program. Three examples of IOS Firewall IDS supporting the security goals of all sizes of organizations include:

- Small and medium-sized businesses looking for a cost-effective way to add IDS features to their security policies for their network router(s).
- Enterprise customers looking for a cost-effective way to extend their IDS security protection and policies across all network boundaries, including branch-office, intranet, and extranet perimeters.
- Service providers that want to provide router-based managed firewall and intrusion detection services for their customers.

The IOS IDS support of the Cisco Secure IDS Director security-management system allows many routers and the Catalyst 6500 family of switches to provide additional security and visibility into the network in support of the organization's Cisco Secure IDS appliance implementation. The Cisco Secure IDS appliance features and implementation are covered in detail in the last four chapters of this book.

Cisco IOS Firewall IDS Configuration Task List

Use these four basic steps to configure the IDS features if the network is using the Cisco Secure Director:

- Initialize Cisco IOS Firewall IDS (required)
- Initialize the Post Office (required)
- Configure and apply audit rules (required)
- Verify the configuration (optional)

The second step is only applicable when working with the Director. If the network isn't using the Director, a logging feature within initializing the Post Office steps should still be used to create a history of the IDS activity.

STUDY TIP The MCNS 640-100 exam doesn't assume the Cisco Secure Director is present, and therefore doesn't (at press time) require you to know those commands. The commands are included here to complete the coverage and in anticipation that the exam objectives could change. Assume the commands in the other three steps are required.

Initializing the IOS Firewall IDS

Only two commands initialize Cisco IOS Firewall IDS on a router. This section introduces the purpose, syntax, and examples of each.

The ip audit smtp spam Command

Use the **global configuration mode** command **ip audit smtp spam** to set the threshold beyond which you suspect e-mail messages contain spam. Use the no version of this command to set the number of recipients to the default setting. The syntax is

```
Rtr1(config)#ip audit smtp spam recipients
Rtr1(config)# no ip audit smtp spam
```

recipients	Integer (1–65535) that designates the maximum number of recipients in a mail message before a spam attack is suspected. The default is 250.

This command was introduced in IOS 12.0(5)T. The default is 250.
The following example changes the maximum number of e-mail recipients to 70:

```
Rtr1(config)#ip audit smtp spam 70
```

The ip audit po max-events Command

Use the **global configuration mode** command **ip audit po max-events** to specify the maximum number of event notifications placed in the router's event queue. Additional events are dropped from the queue for sending to the CiscoSecure IDS Director. This command can be used regardless of whether the Cisco Secure IDS Director is used. Use the no version of this command to return the number to the default setting. The syntax is

```
Rtr1(config)#ip audit po max-events number_events
Rtr1(config)#no ip audit po max-events
```

number-of-events	Integer (1–65535) that designates the maximum number of events allowed in the router event queue. Increasing this number could impact memory and performance—each event in the queue requires 32KB of memory.

This command was introduced in IOS 12.0(5)T. The default is 100.
The following example changes the maximum number of events queued to 50.

```
Rtr1(config)#ip audit po max-events 50
```

Initializing the Post Office

Only three commands initialize the Post Office system on a router. This section introduces the purpose, syntax, and examples of each. If the Cisco Secure IDS Director isn't

being used, the only command that should be used is the **ip audit notify** command to specify that alarms be sent to a Syslog server to preserve a record of the event.

NOTE Each time a change is made to the Post Office configuration, the router must be reloaded.

The ip audit notify Command

Use the **global configuration mode** command **ip audit notify** to specify the method(s) of event notification. If alarms are to be sent to a Cisco Secure IDS Director, use the **nr-director** keyword in the command syntax. If alarms are to be sent to a Syslog server, use the **log** keyword in the command syntax. The two commands can be used together to log to both devices. Use the no version of this command to return the number to the default setting. The syntax is

```
Rtr1(config)#ip audit notify {nr-director | log}
Rtr1(config)#no ip audit notify {nr-director | log}
```

nr-director	Sends messages in Director format to the Cisco Secure Director or Sensor
log	Sends messages in syslog format

This command was introduced in IOS 12.0(5)T. The default is to send messages in syslog format.

The following example directs logging to both the Cisco Secure IDS Director and a Syslog server.

```
Rtr1(config)#ip audit notify nr-director
Rtr1(config)#no ip audit notify log
```

While they're not a part of IDS, the following commands must also be issued to identify the Syslog server, using the IP address or host name and direct logging input to the server.

```
Rtr1(config)#logging 192.168.1.10
Rtr1(config)#logging on
```

Without these commands, the output would be displayed on the router console and would look like the following output. Notice the output includes the IDS signature detected, plus the source and destination IP addresses.

```
01:04:33: %IDS-4-ICMP_ECHO_SIG: Sig:2004:ICMP Echo Request - from 192.168.1.10 t
o 192.168.2.1
01:04:34: %IDS-4-ICMP_ECHO_REPLY_SIG: Sig:2000:ICMP Echo Reply - from 192.168.2.
1 to 192.168.1.10
```

The previous signatures triggered are

2000 ICMP Echo Reply (Info, Atomic)	Triggers when an IP datagram is received with the IP header protocol field set to 1 (ICMP) and the type field in the ICMP header set to 0 (Echo Reply)
2004 ICMP Echo Request (Info, Atomic)	Triggers when an IP datagram is received with the IP header protocol field set to 1 (ICMP) and the type field in the ICMP header set to 8 (Echo Request)

Figure 7-2 shows what the entries might look like in a Syslog server. The sample uses a free Syslog daemon from Kiwi Enterprises.

If messages are sent to the Cisco Secure IDS Director, then it's necessary also to configure the Cisco Secure Director's Post Office transport parameters for both the router (using the **ip audit po local** command) and the Cisco Secure IDS Director (using the **ip audit po remote** command).

The ip audit po local Command

Use the **global configuration mode** command **ip audit po local** to specify the local Post Office parameters used when sending event notifications to the Cisco Secure IDS Director. Use the no form of this command to set the local Post Office parameters to their default settings. The syntax is as follows:

```
Rtr1(config)#ip audit po local hostid host-id orgid org-id
Rtr1(config)#no ip audit po local [hostid host-id orgid org-id]
```

host-id	Unique integer (1–65535) that identifies the router
org-id	Unique integer (1–65535) that identifies the organization to which both the router and the Director belong

This command was introduced in IOS 12.0(5)T. The default organization ID is 1. The default host ID is 1.

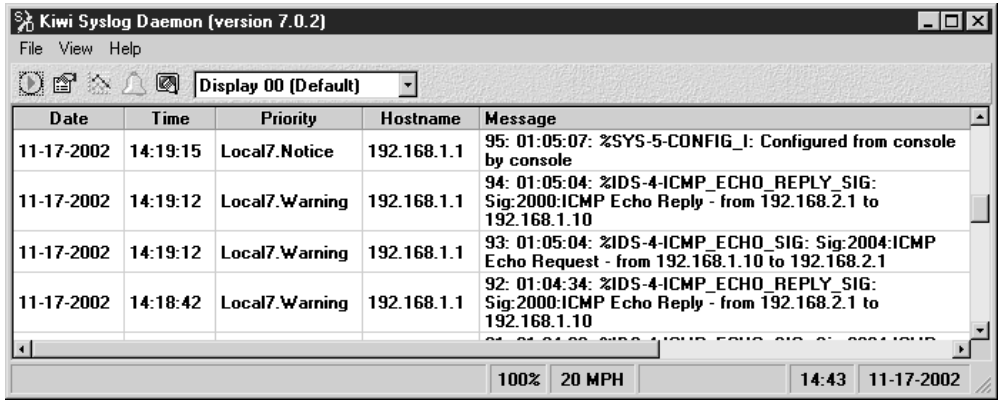

Figure 7-2 Sample Syslog output showing IDS activity

In the following example, the local host is assigned a host ID of 777 and an organization ID of 25:

```
Rtr1(config)#ip audit po local hostid 777 orgid 25
```

The ip audit po remote Command

Use the **global configuration mode** command **ip audit po remote** to specify one or more sets of Post Office parameters for Cisco Secure IDS Director(s) receiving event notifications from the router. A router can report to more than one Director by adding an **ip audit po remote** command for each Director. Use the no form of this command to remove a Director's Post Office parameters as defined by host ID, organization ID, and IP address. The syntax is

```
Rtr1(config)# ip audit po remote hostid host-id orgid org-id rmtaddress ip-add
localaddress ip-add [port port-num] [preference pref-num] [timeout seconds]
[application {director | logger}]
```

host-id	Unique integer (1–65535) that identifies the router.
org-id	Unique integer (1–65535) that identifies the organization to which both the router and the Director belong.
rmtaddress *ip-add*	The Director's IP address.
localaddress *ip-add*	The router's interface IP address.
port-num	The UDP port the Director listens for alarms (the default is 45000).
pref-num	The relative priority of the route to the Director (1 is the default). If more than one route is used to reach the same Director, then one must be a primary route (preference 1) and, the other, a secondary route (preference 2).
seconds	The integer representing the heartbeat timeout value—in seconds—the Post Office waits before it determines a connection has timed out (the default is 5).
app-type	Either *director* (Cisco Secure IDS Director) or *logger* (the default is director).

This command was introduced in IOS 12.0(5)T. The default organization ID is 1, the default host ID is 1, the default UDP port number is 45000, and the default preference is 1. The default heartbeat timeout is five seconds.

Director vs. Logger Application

Use *logger application* instead of the default *director* when sending Post Office notifications to a sensor. Sending to a logging application means no alarms are sent to a GUI. Instead, the Cisco Secure IDS alarm data is written to a flat file, which can then be processed with scripts or filters, such as perl and awk, or staged to a database. Use logger only in advanced applications where you want the alarms only to be logged and not displayed.

Multiple Routes to the Same Director

More than one route can be established to the same Director by giving each route a preference number that establishes the relative priority of routes. The router always attempts

to use the lowest numbered route, switching automatically to the next higher number when a route fails and then switching back when the route begins functioning again.

In this example, two routes for the same dual-homed (residing on two networks) IDS Director are defined.

```
Rtr1(config)#ip audit po remote hostid 777 orgid 25 rmtaddress 192.168.1.3 localaddress 192.168.1.1 preference 1

Rtr1(config)#ip audit po remote hostid 777 orgid 25 rmtaddress 192.168.6.3 localaddress 192.168.6.1 preference 2
```

The router will use the first entry to establish communication with the Director, host ID 777, and organization ID 25. The router will switch to the secondary route if the preferred route fails. When the preferred route returns to service, the router switches back and closes the secondary route.

In this example, the Director is assigned a longer heartbeat timeout value, with a default of five seconds, because of predictable network congestion. This is also designated as a logger application.

```
Rtr1(config)#ip audit po remote hostid 727 orgid 25 rmtaddress 192.168.4.3 localaddress 192.168.4.1 timeout 10
application logger
```

After you configure the router, it's necessary to add the Cisco IOS Firewall IDS router's Post Office information to Cisco Secure IDS Sensors and Directors communicating with the router. This process is covered in the last four chapters of the book.

NOTE Remember, if the Post Office features are added or changed, then it's necessary to save the configuration and reload the router.

Creating and Applying Audit Rules

The Cisco IOS IDS technology works as an inline sensor, monitoring packets as they travel between the router's interfaces. If a packet, or group of packets, in a session matches an active signature, the IOS IDS can perform any or all of the following actions based on the predefined router configuration.

- **Alarm** Sends an alarm to a Syslog server and/or a Cisco Secure IDS Director
- **Drop** Discards the packet
- **Reset** Resets the questionable TCP connection

Two basic steps are necessary to set up the packet auditing process for the Cisco IOS Firewall IDS router.

1. Create an audit rule specifying which signatures are to be applied to packet traffic and the specific action(s) to take when a match is found.
2. Apply the audit rule to a router interface, specifying a traffic direction (in or out).

As packets pass through an interface covered by the audit rule, they're monitored by a series of audit modules in the following order:

- IP module
- TCP, UDP, or ICMP modules (as appropriate)
- Application-level modules

If a pattern matching a known signature is found by any audit module, then the following action(s) occur, based on the instructions included in the router configuration. Any or all of the actions can be configured.

Action	Result
alarm	The module completes its audit. It sends an alarm to the Syslog and/or IDS Director. The packet is forwarded to the next module.
drop	The packet is dropped from the module, discarded, and not sent to the next module.
reset	If this is a TCP session, the reset flag (bit) is set On and sent to both ends of the session. The packets are forwarded to the next module.

NOTE Cisco recommends the drop and reset actions be used together.

If multiple signature matches occur as a packet is processed by a module, only the first match triggers the specified action—the packet is either discarded (drop) or moved immediately to the next audit module (alarm or reset). Additional matches in other modules can trigger additional alarms, but only one per audit module. This separates the IOS IDS implementation from the Cisco Secure IDS Sensor appliance, which identifies all signature matches for each packet.

Creating an Audit Rule

The idea is to define an audit rule, specifying which signatures should be applied to packet traffic and the action(s) to take when a match is found. The signature list can have a single signature, all available signatures, or a partial list of the signatures. The signatures can be informational and attack signatures. Signatures can be disabled in case of false positives or to meet the unique requirements of the network environment.

Using the **ip audit info** and **ip audit attack** commands is necessary to set the default actions for respective signatures. Both types of signatures can take any or all of the following actions: alarm, drop, and reset.

Define Info Audit Actions

Use the **global configuration mode** command **ip audit info** to specify the default action(s) for info-type signatures. Use the no form of this command to set the default action for info signatures. The syntax is

```
Rtr1(config)#ip audit info {action [alarm] [drop] [reset]}
Rtr1(config)#no ip audit info
```

Action	Sets an action for the info signature to take when a match occurs
Alarm	Sends an alarm to the console, IDS Director, or to a Syslog server
Drop	Drops the packet
Reset	Resets the TCP session

This command was introduced in IOS 12.0(5)T. The default action is Alarm.

In the following example, the default action for info signatures is set to all three actions:

```
Rtr1(config)#ip audit info action alarm drop reset
```

Define Attack Audit Actions

Use the global configuration mode command ip audit attack to specify the default action(s) for attack-type signatures. Use the no form of this command to set the default action for info signatures. The syntax is

```
Rtr1(config)#ip audit attack {action [alarm] [drop] [reset]}
Rtr1(config)#no ip audit info
```

action	Sets an action for the info signature to take when a match occurs
alarm	Sends an alarm to the console, IDS Director, or to a Syslog server
drop	Drops the packet
reset	Resets the TCP session

This command was introduced in IOS 12.0(5)T. The default action is alarm.

In this example, the default action for attack signatures is set to all three actions:

```
Rtr1(config)#ip audit attack action alarm drop reset
```

Create Named Audit Rules

Use the **global configuration mode** command **ip audit name** to create audit rules for info and attack signature types. Any signatures disabled with the **ip audit signature** command don't become a part of the audit rule created with the **ip audit name** command. Use the no form of this command to delete an audit rule. The syntax is

```
Rtr1(config)ip audit name audit-name {info | attack} [list standard-acl] [action [alarm] [drop] [reset]]

Rtr1(config)no ip audit name audit-name {info | attack}
```

audit-name	The name for an audit specification.
info	Specifies that the audit rule is for info signatures.
attack	Specifies that the audit rule is for attack signatures.
list	Specifies an ACL to attach to the audit rule.
standard-acl	The integer representing an access control list. Use with the list keyword.
action	Sets an action for the info signature to take when a match occurs.
alarm	Sends an alarm to the console, IDS Director, or a Syslog server.
drop	Drops the packet.
reset	Resets the TCP session.

This command was introduced in IOS 12.0(5)T. The default action is alarm.

In this example, the default action for attack signatures is set to all three actions:

```
Rtr1(config)#ip audit attack action alarm drop reset
```

The following example creates an audit rule named Audit.99 for info signatures that's configured with all three actions:

```
Rtr1(config)#ip audit name Audit.99 info action alarm drop reset
```

This example demonstrates disabling signature 1000 and then creating an info signature audit rule named Audit.33 that doesn't include that signature:

```
Rtr1(config)#ip audit signature 1001 disable
Rtr1(config)#ip audit name Audit.33 info action alarm drop reset
```

Using ACLs with Named Audit Rules

Using a Standard ACL to help define the traffic to be audited by an audit rule is possible. In the following example, an audit rule named Attack.7 is created that uses ACL 25, which is defined later in the configuration. The ACL doesn't behave the way you'd assume, particularly if you're thinking of it as if it were filtering interface traffic. Instead, when used in this context, the deny statements are indicating that the private networks—192.168.0.0 to 192.168.255.0—aren't filtered through the audit process because they're trusted hosts. All other hosts are defined by the permit any statement and are to be processed by the audit rule.

```
Rtr1(config)#ip audit name Attack.7 list 25
Rtr1(config)#access-list 25 deny 192.168.0.0 0.0.255.255
Rtr1(config)#access-list 25 permit any
```

If some individual hosts or subnets should have been included in the audit, the following example shows how this might be addressed.

```
Rtr1(config)#ip audit name Attack.7 list 25
Rtr1(config)#access-list 25 permit host 192.168.1.117
```

```
Rtr1(config)#access-list 25 permit 192.168.100.16 0.0.15
Rtr1(config)#access-list 25 permit 192.168.8.0 0.0.1.255
Rtr1(config)#access-list 25 deny 192.168.0.0 0.0.255.255
Rtr1(config)#access-list 25 permit any
```

Disabling Individual Signatures

Use the **global configuration mode** command **ip audit signature** to attach a policy to a signature. You can implement two policies: disable a signature or qualify the audit of a signature with an access list. This command is generally used to disable the auditing of a signature or to exclude specific hosts or network segments from being audited. Use the no form of this command to remove the policy. If the policy disabled a signature, then the **no** command re-enables the signature. If the policy attached an access list to the signature, the **no** command removes the access list. The syntax is

```
Rtr1(config)#ip audit signature signature-id {disable | list acl-list}
Rtr1(config)#no ip audit signature signature-id
```

signature-id	A unique integer that specifies a signature in the Director Network Security Database
disable	Disables the ACL associated with the signature
list	Specifies an ACL to associate with the signature
acl-list	The ACL configured on the router

This command was introduced in IOS 12.0(5)T. The default is that no policy is attached to a signature.

Using ACLs When Disabling Individual Signatures

You can use the **ip audit signature** command to apply ACLs to individual signatures to help filter out sources of false alarms. When attaching an ACL to a signature, then it's also necessary to create an audit rule with the **ip audit name** command and to apply that named rule to an interface with the **ip audit** command.

In this example, the 1001 signature is disabled and signature 1004 has ACL 10 attached. As in the preceding example, the ACL doesn't behave as you might initially assume. The hosts on the defined network aren't filtered through the signature because they're trusted hosts or, possibly, because they're causing false positives to occur. All other hosts are defined by the permit any statement and are to be processed by the audit rule.

```
Rtr1(config)#ip audit signature 1001 disable
Rtr1(config)#ip audit signature 1004 list 10
Rtr1(config)#access-list 10 deny 192.168.45.0 0.0.0.255
Rtr1(config)#access-list 10 permit any
```

Apply the Audit Rule to the Interface(s)

The *audit rule* is applied to an interface on the router specifying a traffic direction (in or out) in much the same way that ACLs are applied. As with ACLs, in or out is referenced

to the center of the router, not to the connected network. An *inbound rule* is auditing traffic coming into the router from the specified interface. In deciding which interfaces to use and whether to apply the audits in or out, consider the following information.

Inbound Audits

When an audit rule is applied to the in direction on an interface, packets are audited before the inbound ACL has a chance to discard them. This order allows the administrator, Syslog server, and/or IDS Director to be alerted if an attack or information-gathering activity is underway, even if the router would normally reject the activity.

Outbound Audits

When an audit rule is applied to the out direction on an interface, an *outbound rule,* packets are audited after they enter the router through another interface. In this case, an inbound ACL on the other interface might discard packets before they're audited, meaning the administrator, Syslog server, and/or IDS Director won't be alerted of an attack or information-gathering activity that's occurring. Even though the attack or information-gathering activity was thwarted, the network is unaware of it, and so, while the attacker is preparing their next assault, the administrator doesn't even know to prepare for it.

Applying the Audit Name

Use the **interface configuration mode** command **ip audit** to apply an audit specification created with the **ip audit name** command to a specific interface and for a specific direction. Use the no version of this command to disable auditing of the interface for the specified direction. The syntax is

```
Rtr1(config-if)#ip audit audit-name {in | out}
Rtr1(config-if)#no ip audit audit-name {in | out}
```

audit-name	Name of an audit specification created with the **ip audit name** command
in	Inbound traffic—toward the router
out	Outbound traffic—away from the router

This command was introduced in IOS 12.0(5)T. The default that no audit specifications are applied to an interface or direction is implied.

In the following example, the default action for attack signatures is set to all three actions, and the audit specification—Attack.7—is applied inbound on the Ethernet interface:

```
Rtr1(confisg)#interface e0
Rtr1(config-if)#ip audit Attack.7 in
```

Define the Protected Networks

After you apply the audit rules to the router interfaces, use the **global configuration mode** command **ip audit protected** to specify whether an address is on a protected network. A single address at a time or a range of addresses can be entered at one time. You can make as many entries to the protected networks list as needed. In case of a detected

attack, the corresponding event contains a flag that denotes whether the source and/or destination of the packet belong to a protected network.

Use the no form of this command to remove network addresses from the protected network list. If you specify an IP address for removal, that address is removed from the list. If you don't specify an address, then *all IP addresses* are removed from the list. The syntax is

```
Rtr1(config)#ip audit protected ip-addr [to ip-addr]
Rtr1(config)#no ip audit protected [ip-addr]
```

to	Specifies a range of IP addresses
ip-addr	IP address of a network host

This command was introduced in IOS 12.0(5)T. The default is that if no addresses are defined as protected, then all addresses are considered outside the protected network.

The following example shows three individual addresses and two ranges of addresses to be added to the protected network list. The final entry shows an address removed from the protected list.

```
Rtr1(config)#ip audit protected 192.168.5.1
Rtr1(config)#ip audit protected 192.168.5.8
Rtr1(config)#ip audit protected 192.168.5.211
Rtr1(config)#ip audit protected 192.168.4.1 to 192.168.4.254
Rtr1(config)#ip audit protected 192.168.6.1 to 192.168.7.254
Rtr1(config)#no ip audit protected 192.168.4.75
```

Verifying the IDS Configuration

Four **show** commands are used to verify and monitor IDS configuration and performance. These commands include

- **show ip audit statistics**
- **show ip audit configuration**
- **show ip audit interface**
- **show ip audit all**

The show ip audit statistics Command

Use the **show ip audit statistics EXEC** command to display the number of packets audited and the number of alarms sent, among other information. This command shows any signatures used, how many interfaces are configured for audit, and a summary of session information.

```
Rtr1#show ip audit statistics
Signature audit statistics [process switch:fast switch]
  signature 2000 packets audited: [1:569]
```

```
  signature 2004 packets audited: [2:569]
Interfaces configured for audit 2
Session creations since subsystem startup or last reset 3
Current session counts (estab/half-open/terminating) [1:0:0]
Maxever session counts (estab/half-open/terminating) [1:1:0]
Last session created 00:00:10
Last statistic reset never
Rtr1#
```


The clear ip audit statistics Command.

To reset the audit statistics for packets analyzed and alarms sent, use the **clear ip audit statistics EXEC** command. This command could be handy to eliminate historical data when trying to see current activity. If the data is being logged, clearing these counters shouldn't present any downside. The syntax is

```
Rtr1#clear ip audit statistics
```

The show ip audit configuration Command

Use the **show ip audit configuration EXEC** command to display additional configuration information, including default values that might not be displayed using the **show run** command. You can tell this configuration is logging to a Syslog server only and not doing anything more than logging events for info signatures (alarm only). It's performing all three actions on attack signatures:

```
Rtr1#show ip audit configuration
Event notification through syslog is enabled
Event notification through Net Director is disabled
Default action(s) for info signatures is alarm
Default action(s) for attack signatures is alarm
Default threshold of recipients for spam signature is 25
PostOffice:HostID:0 OrgID:0 Msg dropped:0
          :Curr Event Buf Size:0  Configured:100
Post Office is not enabled - No connections are active
Audit Rule Configuration
 Audit name Audit-1
    info actions alarm
    attack actions alarm drop reset
Rtr1#
```

The show ip audit interface Command

Use the **show ip audit interface EXEC** command to display the interface configuration. This command shows the auditing being done on each interface.

```
Rtr1#show ip audit interface
Interface Configuration
 Interface FastEthernet0
  Inbound IDS audit rule is Audit-1
    info actions alarm
    attack actions alarm drop reset
  Outgoing IDS audit rule is not set
```

```
 Interface Serial1
  Inbound IDS audit rule is Audit-1
    info actions alarm
    attack actions alarm drop reset
  Outgoing IDS audit rule is not set
Rtr1#
```

The show ip audit all Command

Use the **catch-all** command—**show ip audit all**—to include the output from the previous commands:

```
Rtr1#show ip audit all
Event notification through syslog is enabled
Event notification through Net Director is disabled
Default action(s) for info signatures is alarm
Default action(s) for attack signatures is alarm
Default threshold of recipients for spam signature is 25
PostOffice:HostID:0 OrgID:0 Msg dropped:0
          :Curr Event Buf Size:0  Configured:100
Post Office is not enabled - No connections are active
Audit Rule Configuration
 Audit name Audit-1
    info actions alarm
    attack actions alarm drop reset
Interface Configuration
 Interface FastEthernet0
  Inbound IDS audit rule is Audit-1
    info actions alarm
    attack actions alarm drop reset
  Outgoing IDS audit rule is not set
 Interface Serial1
  Inbound IDS audit rule is Audit-1
    info actions alarm
    attack actions alarm drop reset
  Outgoing IDS audit rule is not set
Rtr1#
```

Chapter Review

The Cisco IOS Firewall Intrusion Detection System has been a feature of a growing list of Cisco router platforms running the firewall feature set since version 12.0(5). The IDS features are currently available on the Cisco uBR900, 1720, 2600, 3600, 7100, 7200, and 7500 series routers, as well as the RSM for Catalyst 5000 switches.

The IOS-based IDS features extend the Cisco Secure IDS appliance and host-based software features to include the router-based firewalls. The IDS features can be especially useful in locations where a router is being deployed to provide additional security between network segments, such as between the organization and a partner site. The key advantage of an IDS-enabled device is the capability to take preconfigured steps to thwart an attack, rather than simply report it.

The most recent Cisco IOS Firewall IDS uses 59 attack signatures representing a broad cross section of intrusion detection signatures that identify severe breaches of security, as well as the most common network attacks and information-gathering scans.

Four basic steps configure the IDS features if the network is using the Cisco Secure Director (NetRanger). Those steps include initializing Cisco IOS Firewall IDS (required), initializing the Post Office (required), configuring and applying audit rules (required), and verifying the configuration (optional).

The MCNS exam objectives only include the first, third, and fourth steps because the IDS Director is covered in the IDS exam. In the configuring audit rules step, it's possible to disable certain signatures because they're unneeded or create false positive responses. You can also exempt some hosts from the auditing process reflecting their trusted status or to avoid false positive responses.

Questions

1. Which of the following is *not* a method of categorizing IDS systems covered in this chapter?
   - **A.** Misuse detection vs. anomaly detection
   - **B.** Network-based vs. host-based systems
   - **C.** Open vs. proprietary
   - **D.** Passive system vs. reactive system
2. True or False. The Cisco IOS Firewall feature set is implemented on all Cisco router series.
   - **A.** True
   - **B.** False
3. Which of the following IOS features is always found with the IDS features on the new IOS?
   - **A.** VPN features
   - **B.** Firewall feature set
   - **C.** AAA
   - **D.** Java blocking
4. Which of the following is *not* an action the IOS IDS can be configured to do?
   - **A.** Drop
   - **B.** Shut down
   - **C.** Alarm
   - **D.** Reset
5. Cisco IDS is based on matching traffic patterns to which of the following?
   - **A.** Virus profiles
   - **B.** Autographs
   - **C.** Signatures
   - **D.** Traffic baselines

6. Which of the four basic steps to configure the IDS features isn't required?
   - **A.** Initializing the Post Office
   - **B.** Initializing Cisco IOS Firewall IDS
   - **C.** Verifying the configuration
   - **D.** Configuring and applying audit rules
7. Which two of the following are the commands to initialize Cisco IOS IDS on a router?
   - **A.** Rtr1(config)#ip audit notify {nr-director | log}
   - **B.** Rtr1(config)#ip audit smtp spam *recipients*
   - **C.** Rtr1(config)#logging on
   - **D.** Rtr1(config)#ip audit po max-events *number_events*
8. Which of the following signature keywords means "information-gathering activity, such as a port sweep"?
   - **A.** Attack
   - **B.** Atomic
   - **C.** Info
   - **D.** Compound
   - **E.** Apply Inspection Rules and ACLs
9. Which of the following signature keywords means "simple patterns, such as an attempt to access a specific port on a specific host"?
   - **A.** Attack
   - **B.** Atomic
   - **C.** Info
   - **D.** Compound
10. According to the text, how many IDS signatures are supported in the IOS version of IDS?
    - **A.** 39
    - **B.** 59
    - **C.** 79
    - **D.** The number changes with daily updates.
11. With the **ip audit info** and **ip audit attack** commands, how many actions can be applied?
    - **A.** 1
    - **B.** 2

C. 3

D. Depends on where it's applied

12. Which command will disable the IDS signature 1001?

A. Rtr1(config-if)#ip audit signature 1001 disable

B. Rtr1(config)#ip audit signature 1001 disable

C. Rtr1(config-if)#ip audit signature disable 1001

D. Rtr1(config)#ip audit signature disable 1001

13. In the command **Rtr1(config)#ip audit name Attack.7 list 25**, what does the 25 represent?

A. The IDS signature

B. The first 25 matching packets

C. An ACL number

D. The hold-time for attack packets

14. What does the **ip audit** ***audit-name*** command do?

A. Creates an IDS audit specification

B. Is used to disable an IDS signature

C. Applies an IDS audit specification to an interface

D. No such command

15. Which is *not* a valid IDS **show** command?

A. show ip audit transactions

B. show ip audit statistics

C. show ip audit configuration

D. show ip audit interface

Answers

1. **C.** Open vs. proprietary is *not* a category used.
2. **B.** False. This is implemented only on the Cisco 800, uBR900, 1400, 1600, 1700, 2500, 2600, 3600, 7100, 7200, and 7500, and RSM series routers
3. **B.** Firewall feature set is always with IDS in the IOS.
4. **B.** Shut down
5. **C.** Signatures
6. **C.** Verifying the configuration

7. **B. and D.** **Rtr1(config)#ip audit smtp spam** *recipients* and **Rtr1(config)#ip audit po max-events** *number_events*
8. **C.** **Info**
9. **B.** **Atomic**
10. **B.** 59
11. **C.** 3: Alarm, Drop, and Reset
12. **C.** Rtr1(config)#ip audit signature 1001 disable
13. **C.** An ACL number
14. **C.** Applies an IDS audit specification to an interface
15. **A.** show ip audit transactions is *not* a real command

CHAPTER 8

IOS Firewall—Authentication Proxy

In this chapter, you will learn how to:

- Use authentication proxy concepts and processes
- Understand AAA server configuration
- Make use of AAA router configuration
- Apply authentication proxy configuration on the router
- Verify authentication proxy configuration

One of the concerns of every network administrator with a connection to the outside world is how to make sure only approved users can access resources. A secondary need is to make sure those approved users are limited to the resources and services specified in the organization security policy. In this chapter, you will learn about the Cisco IOS Firewall authentication proxy features and the steps necessary to configure them.

STUDY TIP The material in this chapter is part of the new exam objectives for only the Managing Cisco Network Security (MCNS). This was also included in a lighter form in the objectives for the old Cisco Secure PIX Firewall Advanced (CSPFA) exam.

Cisco IOS Firewall Authentication Proxy

The Cisco IOS Firewall authentication proxy feature allows network administrators to implement security policies on a per-user basis through personalized ACLs. Without firewall authentication proxy, user identity and any authorized access was associated with a user's IP address. Any single security policy had to be applied to an entire user group or subnet. Now, users can be identified and authorized on the basis of their per-user policy, and any access privileges can be customized based on their individual access profiles.

With the authentication proxy feature, users can log in to the network or access the Internet via HTTP. Their specific access profiles are automatically retrieved and applied

from a Cisco Secure Access Control Server (ACS), or another RADIUS or TACACS+ authentication server. The user profiles are active only while there's active traffic from the authenticated user.

The authentication proxy is supported on the latest IOS versions (12.2) of the SOHO 70, 800, uBR900, 1720, 2600, 3600, 7100, 7200, and 7500 series routers. Earlier versions won't support the feature on the smaller units (SOHO 70 to 1720s). Authentication proxy is compatible with other Cisco IOS security features, such as NAT, CBAC, IPSec encryption, and VPN client software.

How the Authentication Proxy Works

Unlike many Cisco IOS Firewall features that operate transparently to the user, the authentication proxy feature requires some user interaction on the client host. When a user, using a web browser, initiates an HTTP session through a firewall configured to support the authentication proxy, the process is triggered. The first thing the authentication proxy checks is to see if the user has already been authenticated. If so, the connection is completed without further intervention. But, if no valid authentication entry exists, the authentication proxy responds by providing a screen that prompts the user for a user name and a password. Figure 8-1 shows the message that greets the user.

In the sample, the Rtr1 text is the host name defined for the firewall router. The rest of the text might vary slightly with the version of Cisco Secure ACS and the operating system (OS) platform.

The users must successfully authenticate by supplying a valid user name and password combination recognized by the defined authentication server. Figure 8-2 shows a successful attempt response.

If the authentication attempt failed, the authentication proxy would display a message stating Authentication Failed! and then prompt the user for retries. After five failed attempts to authenticate, the user would wait two minutes, and then would have to initiate another HTTP session to trigger authentication proxy.

User Profiles and Dynamic ACL Entries

With a successful authentication, the user's authorization profile stored on the AAA server is downloaded to the firewall to create dynamic access control entries (ACEs). These temporary ACEs are added to the inbound ACL for an input interface and to the

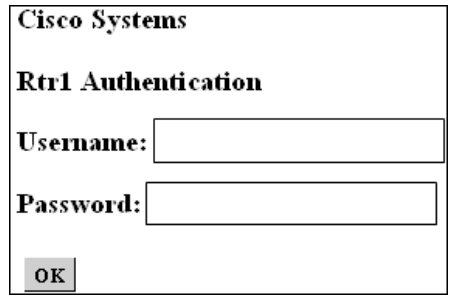

Figure 8-1 Authentication proxy login screen

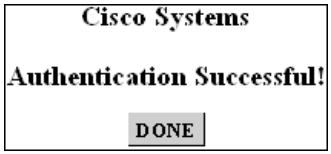

Figure 8-2 Authentication successful attempt response message

outbound ACL of an output interface, if an output ACL exists at the interface. These temporary ACEs configure the firewall to allow only the authenticated users access to the network that was specifically permitted by the authorization profile. For example, the user could only initiate a Telnet session through the firewall if either Telnet or all of TCP is specifically permitted in that user's profile. The following example is a TACACS+ user profile stored on a Windows AAA server.

```
default authorization = permit
key = test
user = newuser1 {
login = cleartext test
service = auth-proxy
{
priv-lvl=15
proxyacl#1="permit tcp any any eq 69"
proxyacl#2="permit icmp any host 192.168.7.2"
proxyacl#3="permit tcp any any eq ftp"
proxyacl#4="permit tcp any any eq ftp-data"
proxyacl#5="permit tcp any any eq telnet"
proxyacl#6="permit tcp any any eq smtp"
}
}
```

The various proxyacl#*n* lines will be downloaded as ACEs in the firewall router to allow the defined protocol or activity. Notice no deny statements exist because all traffic should have already been denied by a standard ACL. These entries are creating the exceptions. All source addresses are set to any because they'll be replaced by the IP address of the authenticating user.

The following output shows input ACL (160) following the user on host 192.168.1.10 successfully authenticating. The first six permit statements represent the downloaded user profile enabling the specified features for the authenticated user. The last four lines are the "real" ACL entries.

```
Rtr1#show ip access-lists
Extended IP access list 160
permit tcp host 192.168.1.10 any eq 69
permit icmp host 192.168.1.10 host 192.168.4.2
permit tcp host 192.168.1.10 any eq telnet
permit tcp host 192.168.1.10 any eq ftp
permit tcp host 192.168.1.10 any eq ftp-data
permit tcp host 192.168.1.10 any eq smtp
deny tcp any any eq telnet
deny udp any any
permit tcp any any (76 matches)
permit ip any any
```

Idle Timer

With a successful authentication, the authentication proxy sets up an inactivity (idle) timer for the created user profile. As long as any authorized user traffic goes through the firewall, the idle timer is continuously reset and the user's host doesn't trigger the authentication proxy. The **ip auth-proxy auth-cache-time** ***minutes*** command is used (covered in the section "Authentication Proxy Configuration on the Router") to set the idle timer. The default is 60 minutes.

If the idle timer expires, the authentication proxy then removes the user's profile information and any related dynamic ACEs in the access lists. Any subsequent traffic from the client host is blocked until they initiate another HTTP session to trigger the authentication proxy.

Secure Authentication

The authentication proxy uses a JavaScript applet to help establish a secure authentication session with the user's web browser. This secure authentication prevents a client from mistakenly submitting a user name/password combination to a network web server other than the authentication proxy router. For this to work, user browsers must allow JavaScripts prior to initiating the HTTP connection. With the JavaScript feature enabled on the browser, secure authentication is done automatically and the user sees the authentication message displayed earlier in Figure 8-2.

NOTE There is a work-around for networks that won't allow the JavaScript feature. While this technique works, it's neither automatic nor particularly reassuring to the user. The Cisco online documentation has the steps to follow if the JavaScript feature can't be enabled.

Applying the Authentication Proxy

Figure 8-3 shows a simple authentication proxy implementation where a single connection exists to the local, protected network and a serial connection exists to the Internet via an Internet service provider (ISP). Access control lists are applied inbound on both interfaces to block all incoming traffic. The lone exception is this: an opening must be created that allows the AAA server to communicate with the firewall router.

One strategy is to apply the authentication proxy in the inbound direction for any router interface for which per-user authentication and authorization is required. By applying the authentication proxy inbound, all user connection requests are intercepted before they're processed by any IOS or firewall features. If the user authentication attempt fails, the connection request is dropped.

The authentication proxy feature supports using a standard access list to specify a particular host or group of hosts whose initial HTTP traffic triggers the proxy. How the authentication proxy will be applied should be driven by the organization security policy.

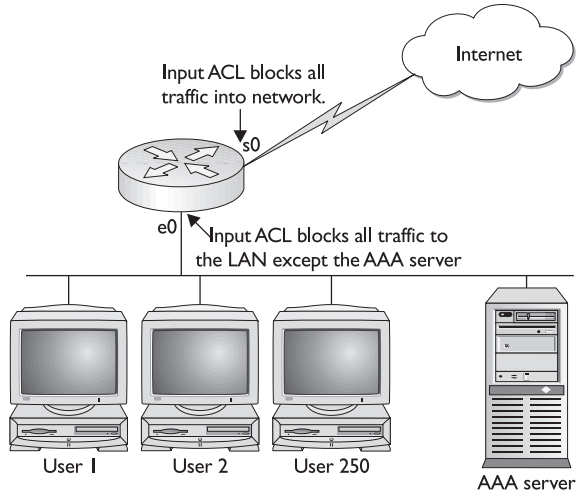

Figure 8-3 Simple network implementation for authentication proxy

Comparison with the Lock-and-Key Feature

Chapter 5 covered another IOS feature—lock-and-key—which seems to provide a similar service to authentication proxy. Lock-and-key also uses authentication and dynamic access lists to provide user access through a firewall. The following table compares the authentication proxy and lock-and-key features.

Authentication Proxy	Lock-and-Key
Triggers on HTTP connection requests.	Triggers on Telnet connection requests.
TACACS+ or RADIUS authentication and authorization.	TACACS+, RADIUS, or local authentication.
Access list entries are created from profile information retrieved from the AAA server only.	Access lists are configured on the router only.
Access privileges are granted on a per-user and host IP address basis.	Access privileges are granted on the basis of the user's host IP address only.
ACLs can have multiple entries defined by the user profiles on the AAA server.	Access lists are limited to one entry for each host IP address.
Allows DHCP-based host IP addresses, meaning users can log in from any host location and obtain authentication and authorization.	Associates a fixed IP address with a specific user. Users must log in from the host with that IP address.

Compatibility with Other Features

The authentication proxy works transparently with the Cisco IOS Firewall intrusion detection system and Cisco IOS IPSec encryption features. This section looks at issues related to authentication proxy and features, such as NAT, CBAC, VPN client software, AAA, and one-time password (OTP) features.

CBAC Compatibility

Authentication proxy *does not create ACEs to support returning data traffic or data channel traffic*, so you must either create static ACLs to allow the return traffic or configure the CBAC inspection rules in the firewall configuration. Remember, static ACLs increase the network vulnerability because they're permanent "holes" in the firewall, whereas a CBAC "hole" only exists as long as it's needed.

Configuring CBAC with any authentication proxy implementation is the most reliable and secure method to ensure return traffic for authorized user connections through the firewall.

Network Address Translation (NAT) Compatibility

Authentication proxy can be successfully implemented on a network using network address translation (NAT) if the CBAC features are configured. Figure 8-4 shows a firewall router performing NAT translation on traffic heading out into the Internet.

With the NAT service, the client local IP address is translated to a "real" global address for any packets traveling out into the Internet. This translation is stored in a NAT translation table on the router, so returning traffic can be retranslated back to the appropriate local address. For example, traffic from host 192.168.1.2 might travel out through the Internet using 172.16.1.1:2070.

But, earlier, when 192.168.1.2 used authentication proxy, a set of temporary ACEs were established using 192.168.1.2 as the authenticated address, not 172.16.1.1. Depending on where those ACEs were applied, our user might not have ACEs to allow the traffic to pass through the router.

Assuming authentication proxy is applied inbound on the LAN interface (e0) only, temporary ACEs are created inbound on that interface to allow that user the use of whatever protocols and features were defined in the downloadable profile. Because NAT will

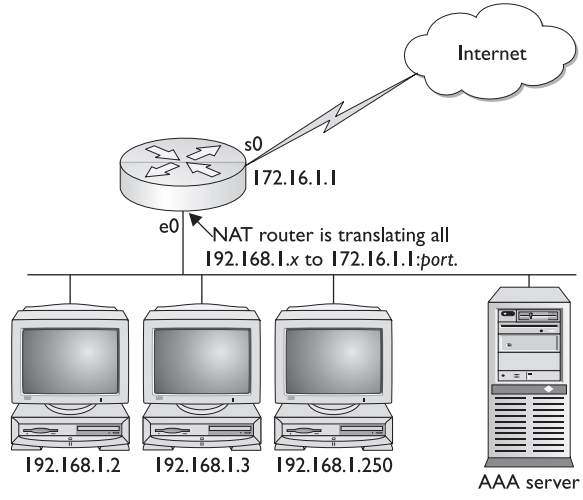

Figure 8-4 Firewall router performing NAT translations

occur outbound after the traffic has passed through the LAN interface, everything should be okay so far. The NAT translations on returning traffic shouldn't be an issue because the returning traffic is unfiltered outbound on the LAN interface. Even if the traffic was filtered, the addresses have been retranslated back to the local addresses before heading out into the LAN.

But, what if the ACEs were applied to the Serial 0 interface? Now the traffic reaching the interface doesn't match the original dynamic ACEs created by the authentication proxy. Configuring CBAC when using authentication proxy with NAT services will take care of the problem by using the translated addresses to create matching ACEs. CBAC makes sure the NAT translated address for the session is associated with the original host address.

When in doubt, use CBAC with NAT and the authentication proxy feature.

VPN Client Compatibility

Authentication proxy can add an extra layer of security and access control for VPN client sessions. A VPN client initiating an HTTP connection through the firewall router is treated like any other user. The authentication proxy checks for a prior client authentication. If a valid authentication is found, authorized traffic is permitted. If not, the HTTP request triggers the authentication proxy and the user is prompted for a user name and password.

Once the VPN user authenticates successfully, the authentication proxy retrieves the user profile from the AAA server. The source address in the user profile entries is then replaced with the IP address of the authenticated VPN client from the decrypted packet.

Compatibility with AAA Accounting Features

Authentication proxy supports the AAA Accounting feature by being able to generate "start" and "stop" accounting records with enough detail to be used successfully for billing and security auditing purposes. This also allows the network administrator to monitor the actions hosts that have used the authentication proxy service.

If the accounting Start option has been enabled, a start record will be generated every time an authentication proxy cache entry and associated dynamic access control list entry are created. Any subsequent traffic from the authenticated host will be recorded when the ACEs created by the authentication proxy receive the packets. The accounting feature saves data about this event in a data structure stored with the data of other users on a defined server.

When the authentication proxy cache expires and is deleted along with the related ACEs, any additional data, such as elapsed time, is added to the accounting information and a stop record is sent to the server.

Operation with One-Time Passwords (OTS)

Authentication proxy supports the use of one-time passwords (OTPs) because AAA supports them. Remember, OTPs typically require an additional authentication server and will require additional configuration of the AAA server. When using an OTP, the user enters the user name and OTP in the HTML login page as usual.

One minor difference is this: the user must enter the correct token password within the first three attempts. If a user has made three consecutive incorrect entries, they must enter two valid token passwords in succession before authentication is granted by the AAA server.

Security Vulnerability Issues

Any time holes are punched through a firewall defense, an increased vulnerability occurs to those who would capitalize on those vulnerabilities. Because authentication proxy should be configured with CBAC, it only makes sense to implement the CBAC attack defenses introduced in Chapter 6.

Denial of Service (DoS) Attacks

Any time a firewall experiences a high level of connection requests requiring authentication, legitimate network users might experience delays in making needed connections or the connection could be rejected and the user would need to try reestablishing the connection. For each request, the authentication proxy prompts the users for login credentials. A high number of open requests could indicate that the router is the subject of a denial of service (DoS) attack.

As a protective mechanism, the authentication proxy monitors incoming HTTP requests, limits the number of open requests, and then drops any additional requests until the number of open requests has fallen below 40.

Source Address Spoofing

Spoofing a source address involves a hacker substituting a known or suspected trusted address for their address both to gain access and conceal their identity. When an HTTP session triggers a successful authentication proxy session, a dynamic opening is created in the firewall by reconfiguring an interface with temporary user-access privileges. As long as this opening exists, it's always technically possible for another user to detect the source IP address information and spoof the address to gain access behind the firewall.

Spoofing is a problem inherent to all access list implementations. Important to realize is that the authentication proxy feature doesn't cause this address spoofing problem, but it does have some level of vulnerability. The authentication proxy doesn't specifically address this problem.

Before Configuring Authentication Proxy

Authentication proxy (auth-proxy) is available only on Cisco IOS Software containing the firewall feature set since version 12.0.5.T.auth-proxy can be used to authenticate inbound or outbound users, or both, by using a web browser to pass through the firewall and authenticate on a TACACS+ or RADIUS server. Traditional access lists are in place to block all traffic. After authentication, the AAA server passes temporary access list entries to the firewall router to allow predefined types of traffic.

Authentication proxy has certain requirements that must be met for effective use of this technology in a network. This section looks briefly at some of the system requirements, as well as the skill sets required of a network administrator tasked with implementing and supporting Cisco secure authentication proxy.

- Client hosts must be running one the following software browsers. These client browsers must have the JavaScript capabilities enabled for secure authentication.
  - Microsoft Internet Explorer 3.0 or later
  - Netscape Navigator 3.0 or later
- Because the authentication proxy activates only on HTTP connections, HTTP services must be running on the standard (well-known) port 80 on the firewall router.
- The authentication proxy feature and related access list entries only apply to traffic passing through the router. Any administrative traffic with the router as the destination is authenticated by the standard authentication methods provided by IOS software.
- Authentication proxy doesn't support concurrent use, meaning it can't be used for multiple users to log in from the same host device at the same time. The authentication and authorization apply only to the first user to submit a valid user name and password. Any others are ignored until the first user ends their session.
- Load balancing between multiple AAA servers isn't supported. Requests only go to any additional AAA server if the first one configured times out.
- Because authentication proxy can use standard access lists, how access lists are used to filter traffic before configuring the authentication proxy is important. For more information on how to configure and use access lists, see Chapter 2 and Appendix A.
- The authentication proxy is a feature of the Cisco Authentication, Authorization, and Accounting (AAA) strategy. It's important to understand how to configure AAA user authentication, authorization, and accounting before you attempt to configure the authentication proxy. For more information on how to configure AAA, see Chapters 3 and 4.
- To create a completely secure and successful implementation of the authentication proxy feature of the Cisco IOS Firewall, you must configure CBAC on the firewall. CBAC features are typically required to allow filtering of protocols permitted by the downloadable user profiles. For more information on configuring and using the CBAC features, see Chapter 6.

Cisco suggests the following tips for implementing the auth-proxy features to reduce the impact on the existing network and the variables involved. This makes it easier to see if the process is being implemented as defined in the security policy.

- Confirm that traffic is flowing properly through the firewall before configuring auth-proxy. Remember, other firewall features exist, such as ACLs and CBAC, that can restrict traffic flow.
- To reduce disruption of the network during testing, modify any existing access list or add an access list to deny access to only one test client.

- Confirm that only the one test client can't get through the firewall and that all other hosts can pass through.
- Add the **exec-timeout 0 0** command under the console port (line con 0) and/or virtual type terminals (line vty 0 4) to prevent your sessions from timing out while you're busy on other devices or reading online documentation. The first 0 is minutes, the second 0 is seconds: exec-timeout 20 30 would set the idle timer to 20 1/2 minutes. The 0 0 combination turns off the idle timer.

Authentication Proxy Configuration Task List

The authentication proxy feature requires an AAA server running Cisco Secure Access Control Server (ACS) to be present on the network. This is necessary to configure the AAA server to enable the features.

Next, the router running the firewall feature set, typically the perimeter router, must be configured by performing the following tasks:

- Configuring AAA support
- Configuring the HTTP server feature
- Configuring the authentication proxy
- Verifying the authentication proxy

Technically, only the first three tasks and configuring the AAA server are required. Skipping the optional verifying is done only at risk to the resources you're charged to protect. The next section looks first at configuring the AAA server, and then at the firewall router.

AAA Server Configuration

Cisco Secure ACS must be configured on a Windows 2000/NT or UNIX server to support either TACACS+ (Cisco preferred) or RADIUS authentication. The next lines demonstrate installing Cisco Secure ACS v2.6 for Windows 2000/NT, which is the latest version at press time. The example assumes a TACACS+ implementation. This information is included as representative of the process, but it's important to make sure you have the latest instructions matching the platform and version of the Cisco Secure ACS software.

To get the latest instructions, go to the Cisco web site at www.cisco.com and search for IOS auth-proxy, and then scan for a document referring to Implementing Authentication Proxy. While searching for Implementing Authentication Proxy could be more direct, the former also includes Command Reference listings and examples for working with other firewall features you might find useful.

The configuration steps are as follows:

1. Open a web browser.

2. In the address or location bar in the web browser, type the applicable URL, using the IP of a remote server, or either http://localhost:2002 or http://127.0.0.1:2002 if you're working directly on the server.
3. If the Cisco Secure ACS for Windows 2000/NT login page appears, follow these steps:
   - Type a valid Cisco Secure ACS administrator name in the User Name box.
   - Type a valid administrator password in the Password box.
   - Click the Login button.
   - The Cisco Secure ACS for Windows 2000/NT initial page appears (Figure 8-5).
4. Click the Interface Configuration option on the left side of the opening screen. The three-column format of the screen includes links for help on the right-hand side, if necessary.
5. From the Interface Configuration screen, select TACACS+ (Cisco IOS). Note, this is also where you could choose a RADIUS (IETF) implementation. Figure 8-6 shows the Interface Configuration screen.

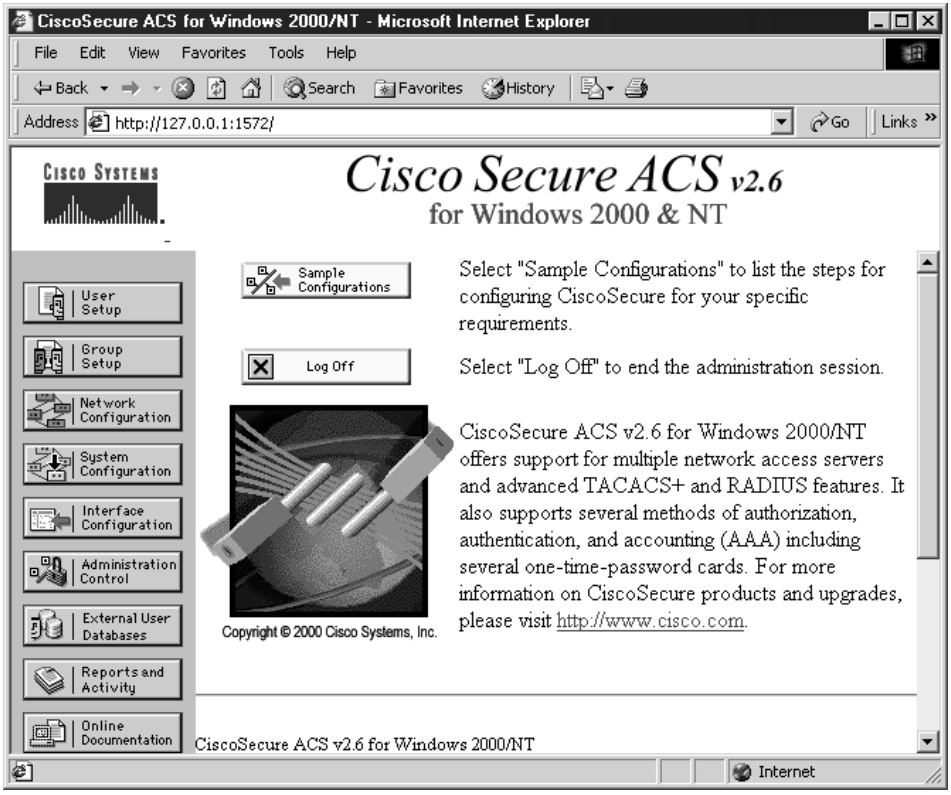

Figure 8-5 Cisco Secure ACS for Windows opening screen

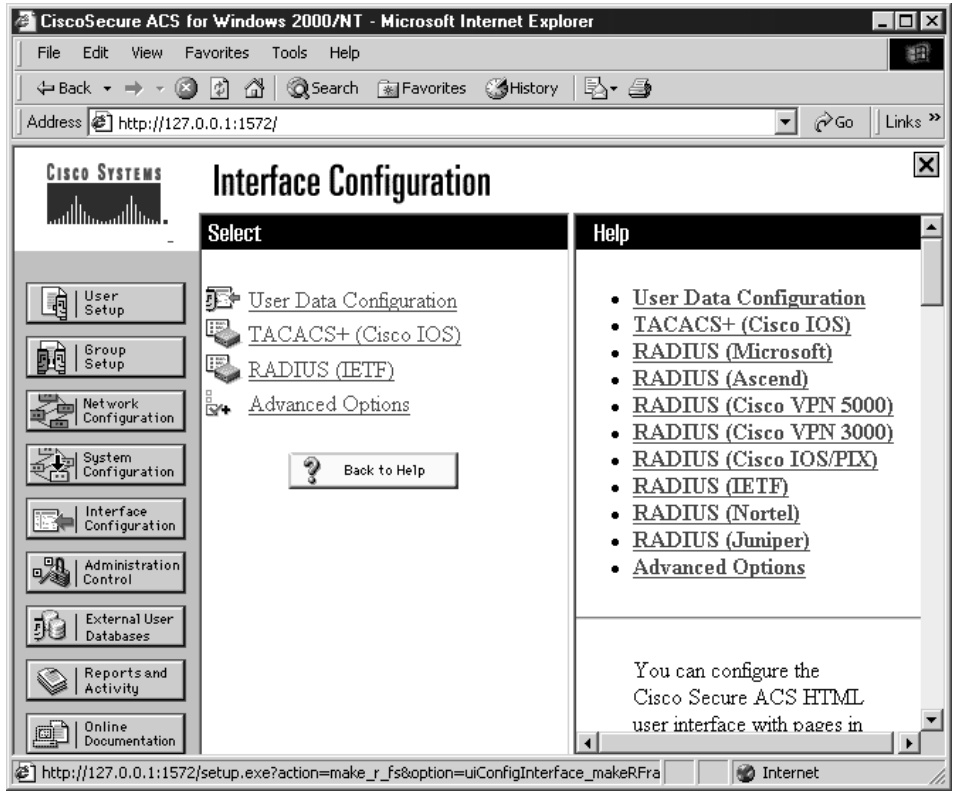

Figure 8-6 Interface Configuration screen showing the TACACS+ and RADIUS options

6. On the resulting screen, scroll down in the center column until you can see the .New Services section (Figure 8-7). Select (check) the Group option for the first blank row, and then type **auth-proxy** in the Service column. Leave the Protocol column blank. Do *not* click the Submit button yet.
7. Further down the same screen, in the Advanced Configuration Options (Figure 8-8), select the Display window for each service selected in which you can enter customized **TACACS+** attributes. Look over the other choices, including the option to create a time-of-day schedule for services. Click the Submit button.
8. Return to the button bar on the left side of the screen and click the Group Settings option. When the next screen appears, choose Edit Settings. You'll be editing the default group settings, but you can use the drop-down list to select a particular group.

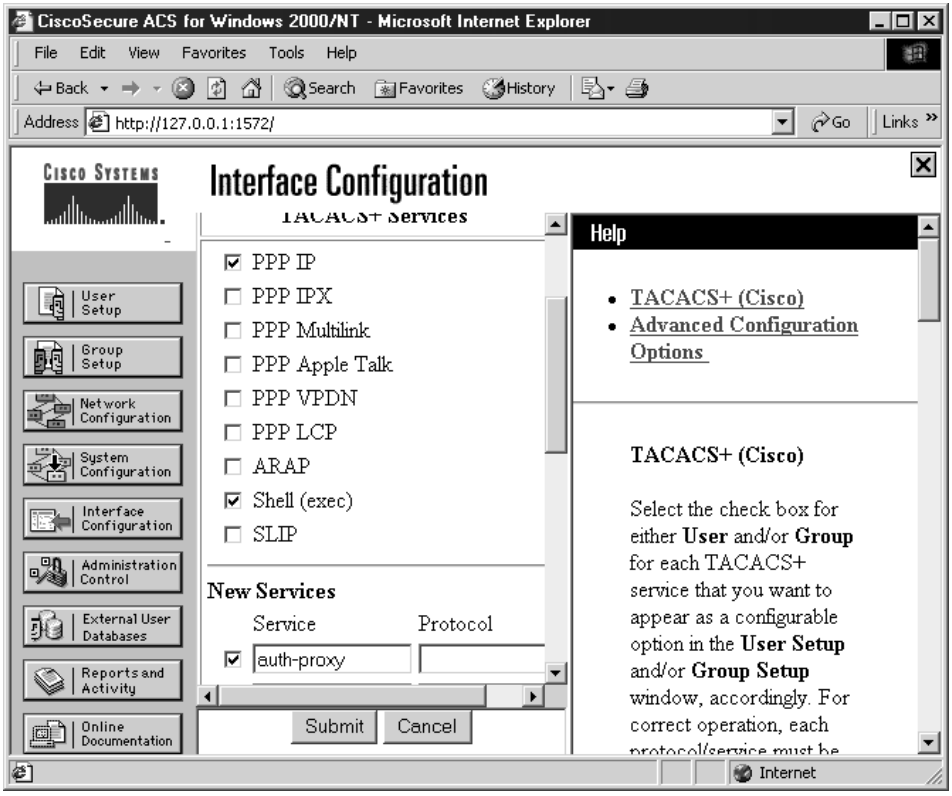

Figure 8-7 TACACS+ configuration page showing the New Services section

9. Find and select the Auth-Proxy option. This is several sections down in the center panel.
10. The authentication proxy features require a per-user downloadable user profile configuration on the AAA server. When the user uses the auth-proxy feature to access the network, the appropriate profile is downloaded and becomes a series of temporary ACL entries on the firewall router. The following is an example of what a user profile looks like on a TACACS server.

```
default authorization = permit
key = test
user = newuser1 {
login = cleartext test
service = auth-proxy
{
priv-lvl=15
proxyacl#1="permit tcp any any eq 69"
```

```
proxyacl#2="permit icmp any host 192.168.7.2"
proxyacl#3="permit tcp any any eq ftp"
proxyacl#4="permit tcp any any eq ftp-data"
proxyacl#5="permit tcp any any eq telnet"
proxyacl#6="permit tcp any any eq smtp"
}
}
```

11. This profile can be created at this point using a type of access control list. Remember the following:
    - Use the **proxyacl#*n*** attribute when configuring the access lists in the profile
    - for both RADIUS and TACACS+ attribute-value (AV) pairs.
    - Note, the **proxyacl#*n*** values are numbered sequentially
    - The privilege level must be set to 15 for all users.
    - Each user profile access list entry can contain only the **permit** keyword.
    - Use the **any** keyword for the source address in each entry. The source address of the host making the authentication proxy request will replace the **any** keyword when the user profile is downloaded to the firewall.

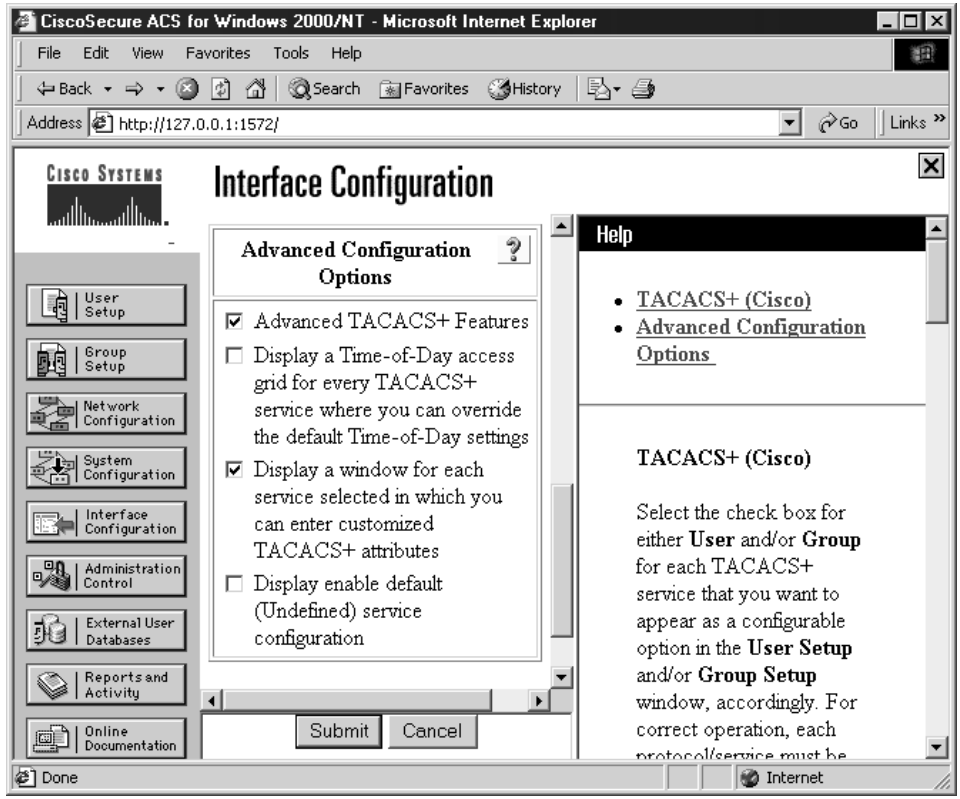

Figure 8-8 Advanced TACACS+ configuration options

In the box following the auth-proxy selection, type the following lines. These lines enable new users to create ICMP, TCP, and UDP sessions.

```
priv-lvl=15
proxyacl#1=permit icmp any any
proxyacl#2=permit tcp any any
proxyacl#3=permit udp any any
```

12. Figure 8-9 shows the completed entries and the various activation buttons. Click the Submit + Restart button when the entry is complete.

NOTE The techniques and exact commands for creating these user downloadable profiles vary with the OS platform and the version of the software. Always check the online documentation before proceeding.

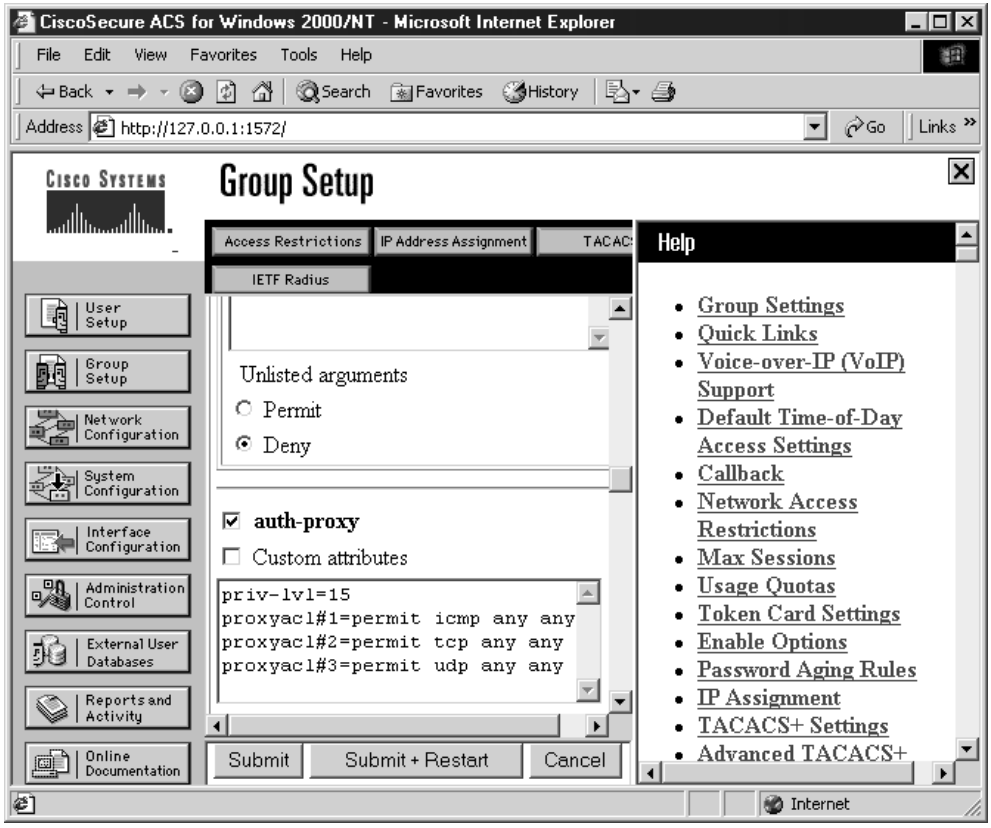

Figure 8-9 Downloadable profile entries and activation buttons

AAA Router Configuration

The following eight tasks are required to enable AAA authentication proxy services. These tasks must be configured in Global Configuration mode before authentication proxy can be configured. Fortunately, the first four tasks should be familiar from Chapter 3.

- Enable AAA.
- Define the security server (TACACS+ or RADIUS).
- Define a server key (TACACS+ or RADIUS).
- Define the list of login authentication methods.
- Enable authentication proxy for AAA.
- Activate authentication proxy accounting.
- Create an ACL entry to allow return traffic from the AAA server to the firewall.
- Enable the HTTP server feature to work with AAA.

Enable AAA

To enable the AAA Access Control model on the router or NAS, in Global Configuration mode, type the **aaa new-model** command. Use the no form of the command to disable the AAA. The syntax is

```
Rtr1(config)#aaa new-model
Rtr1(config)#no aaa new-model
```

No variables exist. The words simply turn the feature on and off. Once entered, the CON, VTY, AUX, and TTY lines require a user name and a password for access. Any previous password configuration for those lines is automatically removed.

NOTE Technically, the **aaa new-model** command not only implements AAA, but also prohibits any of the earlier **TACACS** and **extended TACACS** commands.

Define the Security Server

Defining one or more security servers that provide the authentication services is necessary. The server or servers can be TACACS+, RADIUS, or both. The method lists specified in the **authentication, authorization,** and **accounting** commands determine which type of server should be used for each service and, if both are specified, in which order. If multiple servers of either type are specified, the Cisco IOS software searches for hosts in the order in which they're specified.

Both protocols support Timeout Timer options for defining how long the device will wait for a reply from the specified host before moving on to any remaining server. Both also support Key options, which, if configured on both the server and the firewall router,

provide a level of device authentication. The key, if defined, travels in the request packet and is treated as a password or authentication token by the server. If the key doesn't match the one configured on the server, the request is denied.

The ***protocol*-server host** command is used to define the security server. While they're quite similar, the RADIUS version has more options.

Define the Security Server—TACACS+

If a TACACS+ server host is used for authentication, use the **tacacs-server host global configuration** command. Use the no form of this command to delete the specified name or address. The syntax is

```
Rtr1(config)#tacacs-server host {hostname | ip address} [single-connection] [port port#]
    [timeout seconds] [key string]
Rtr1(config)#no tacacs-server host hostname
```

hostname	Name of the TACACS+ server.
ip-addresss	IP address of the TACACS+ server.
single-connection	(Optional) Specify that the router maintain a single open connection for confirmation from the AAA/TACACS+ server. This command contains no autodetect feature and fails if the host isn't running a Cisco Secure daemon.
port *port#*	(Optional) Overrides the default port number (49) with any number from 1 to 65535.
timeout *seconds*	(Optional) Overrides the global timeout value set with the global **tacacs-server timeout** command for this server only.
key *string*	(Optional) Specifies an authentication and encryption key. This must match the key used by the TACACS+ daemon. Overrides the key set by the **global** command **tacacs-server key** for this server only.

Because some of the parameters of the **tacacs-server host** command override global settings made by the **tacacs-server timeout** and **tacacs-server key** commands, this command can be used to enhance security on a network by uniquely configuring individual TACACS+ connections.

Define TACACS+ Server Key Option The global authentication encryption key is set with the **global configuration** command **tacacs-server key.** This key value must match the key value configured on the TACACS+ server, if one is used. Use the no form of this command to disable the key. The syntax is

```
Rtr1(config)#tacacs-server key string
Rtr1(config)#no tacacs-server key [string]
```

string	Any leading spaces are ignored, *but spaces within and at the end of the key are not.* Don't enclose the key in quotation marks unless they're part of the key.

The following example specifies a global timeout of seven seconds, a global key of cisco-key, and then three TACACS+ servers. The first one—tac-serv1—is the one all requests are sent to first. If those requests aren't answered before the timeout timer expires (seven seconds), the next two servers are tried in order. The global timeout and key settings only apply to the first two servers because the third one has overriding options defined.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server timeout 7
Rtr1(config)#tacacs-server key cisco-key
Rtr1(config)#tacacs-server host tac-serv1
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server host 192.168.6.4 port 1500 timeout 3 key cisco9
```

Define the Security Server—RADIUS

To specify a RADIUS server host, use the **radius-server host global configuration** command. Use the no form of this command to delete the specified RADIUS host. The syntax is

Rtr1(config)#radius-server host {*hostname* | *ip-address*} [auth-port *port-num*] [acct-port *port-num*] [timeout *seconds*] [retransmit *retries*] [key *string*] [alias{*hostname* | *ip-address*}]
Rtr1(config)#no radius-server host {*hostname* | *ip-address*} [auth-port *port-num*] [acct-port *port-num*] [timeout *seconds*] [retransmit *retries*] [key *string*]

hostname	DNS name of the RADIUS server.
ip-address	IP address of the RADIUS server.
auth-port *port-num*	(Optional) To specify a UDP destination port for authentication requests. The host isn't used for authentication if set to 0.
acct-port *port-num*	(Optional) To specify a UDP destination port for accounting requests. The host isn't used for accounting if set to 0.

The following example defines port 12 as the destination port for authentication requests and port 16 as the destination port for accounting requests on a RADIUS host named 192.168.1.4. Because entering a line resets all the port numbers, you must specify a host and configure both the accounting and authentication ports on a single line.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4 auth-port 12 acct-port 16
```

To use separate servers for accounting and authentication, use the zero port value, as appropriate. The following example specifies RADIUS server192.168.1.4 be used for accounting, but not for authentication, and RADIUS server host1 be used for authentication, but not for accounting:

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server host 192.168.1.4 auth-port 0
Rtr1(config)#radius-server host host1.domain.com acct-port 0
```

Define RADIUS Server Key Option The global authentication encryption key is set with the **global configuration** command **radius-server key**. This key value must match the key value configured on the RADIUS server, if one is used. Use the no form of this command to disable the key. The syntax is

Rtr1(config)#radius-server key *string*
Rtr1(config)#no radius-server key [*string*]

string	Any leading spaces are ignored; *but spaces within and at the end of the key are not.* Don't enclose the key in quotation marks unless they're part of the key.

The following example specifies a global timeout of seven seconds, a global key of cisco-key, and then three RADIUS servers. The first—rad-serv1—is the one all requests are sent to first. If those requests aren't answered before the timeout timer expires (seven seconds), the next two servers are tried in order. The global timeout and key settings only apply to the first two servers because the third one has overriding options defined.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server timeout 7
Rtr1(config)#radius-server key cisco-key
Rtr1(config)#radius-server host rad-serv1
Rtr1(config)#radius-server host 192.168.1.4
Rtr1(config)#radius-server host 192.168.6.4 timeout 3 key cisco9
```

Define AAA Group Server (Optional)

The **global configuration mode** command **aaa group server** provides a way to group existing defined server hosts. By grouping a set of servers, you can use the group name with other **AAA** commands to select that subset of server hosts to use them for a particular service. A server group is comprised of server hosts of a particular type. Currently, the IOS supports RADIUS and TACACS+ server hosts. The **aaa group server** command lists the IP addresses of the selected server hosts from the global protocol-server host list. To remove a server group from the configuration list, enter the no form of this command. The syntax is

Rtr1(config)#aaa group server {tacacs+ | radius} *group-name*
Rtr1(config)#no aaa group server tacacs+ | radius *group-name*

group-name	Character string used to name the group of servers

The command was introduced in IOS version 12.0(5)T. There are no default values.

The command can only be entered after the **aaa new-model** command turns on the AAA features. The entry is followed by one or more server server-name statements identifying the servers to be included in the group. The next two sample outputs demonstrate this; be sure to note the prompt for the server entries.

The following example uses the three TACACS+ servers created in the earlier example and selects only two them to be included in the server group.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server timeout 7
Rtr1(config)#tacacs-server key cisco-key
Rtr1(config)#tacacs-server host tac-serv1
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server host 192.168.6.4 port 1500 timeout 3 key cisco9
Rtr1(config)#aaa group server tacacs+ tac-1
Rtr1(config-sg-tacacs+)#server tac-serv1
Rtr1(config-sg-tacacs+)#server 192.168.6.4
```

The following RADIUS example shows the configuration of an AAA group server named rad-1, which includes two of the three defined servers:

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server timeout 7
Rtr1(config)#radius-server key cisco-key
Rtr1(config)#radius-server host rad-serv1
Rtr1(config)#radius-server host 192.168.1.4
Rtr1(config)#radius-server host 192.168.6.4 timeout 3 key cisco9
Rtr1(config)#aaa group server radius rad-1
Rtr1(config-sg-radius)#server rad-serv1 auth-port 1800 acct-port 1801
Rtr1(config-sg-radius)#server 192.168.6.4 auth-port 1802 acct-port 1803
```

The server statement allows the auth-port and acct-port to be defined. If not specified, the default value of auth-port is 1645 and the default value of acct-port is 1646. The output for next paragraph demonstrates this.

With either protocol, if you attempt to create the server group without first defining the servers with ***protocol*-server host** commands, you'll get the following error messages because the IOS attempts to verify the server presence. Notice the entry with

```
Rtr1(config)#aaa new-model
Rtr1(config)#aaa group server radius rad-1
Rtr1(config-sg-radius)#server rad-serv1
Translating "rad-serv1"...domain server (255.255.255.255)
                          ^
% Invalid input detected at '^' marker.

Rtr1(config-sg-radius)#server 192.168.6.4
Rtr1(config-sg-radius)#
00:02:11: %RADIUS-4-NOSERV: Warning: Server 192.168.6.4:1645,1646 is
not defined.
Rtr1(config-sg-radius)#^z
Rtr1#show run
Building configuration...
!
hostname Rtr1
!
aaa new-model
aaa group server radius rad-1
 server 192.168.6.4 auth-port 1645 acct-port 1646
!
```

Notice that the entry with a named server—rad-serv1—returned an invalid input error and the command was rejected. The entry that used the server IP address—192.168.6.4—returned a system warning message indicating that the specified address can't be found, but the entry was accepted. The **show run** command in the same output confirms this. The result is identical for TACACS+. So, if you're going to define servers by name, it's critical for your **ip host** table to be created or your **ip name-server** command(s) to be issued before creating your server groups.

Define Login Authentication Methods List

To use AAA authentication for device login, use the **aaa authentication login global configuration** command. Use the no form of this command to disable AAA authentication. The syntax is

Rtr1(config)#aaa authentication login {default | *list-name*} *method1* [*method2. . .]*
Rtr1(config)#no aaa authentication login {default | list-name} method1 [*method2. . .*]

default	Uses the listed authentication methods as the default list when a user logs in.
list-name	Character string used to name the list of authentication methods that can be specified to use when a user logs in.
method	At least one of the keywords described in the next table.

The *method#* arguments identify the list of methods the authentication algorithm tries, in the stated sequence. Method keywords are described in the following table.

Keyword	Description
group tacacs+	Use the list of all TACACS+ servers to authenticate services.
group radius	Use the list of all RADIUS servers to authenticate services.
group *group-name*	Use a subset of RADIUS or TACACS+ servers for authentication, as defined by the server group *group-name* in the preceding section.
Local	Use the local user name database for authentication.
local-case	Use case-sensitive local user name database for authentication.
Enable	Use the enable password for authentication.
Line	Use the line password for authentication.
krb5	Use Kerberos 5 for authentication.
krb5-telnet	Use Kerberos 5 Telnet authentication protocol when using Telnet to connect to the router.
None	Use no authentication—no security.

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

Examples of this command are included in the next section.

Enable Authorization Proxy (auth-proxy) for AAA

To configure AAA authorization to all proxy authentications, use the global configuration mode command **aaa authorization** to enable authorization and to create a method list, defining authorization methods that can be used when a user accesses the specified function.

In Chapter 3, you learned a method list is simply a list describing the authorization methods to be queried (such as RADIUS or TACACS+), in sequence. Method lists enable one or more security protocols to be used for authorization, thus ensuring a backup system if the initial method fails.

Use the auth-proxy option of the **global configuration** command **aaa authorization** to set parameters for proxy authentications. Use the no form of this command to disable AAA authentication proxy. The general syntax of the command is

```
Rtr1(config)#aaa authorization {authorization-type} {default | list-name} method1 [method2. . .]
Rtr1(config)#no aaa authorization {authorization-type}
```

The AAA authorization types include the following:

```
Rtr1(config)#aaa authorization {network | exec | commands level| reverse-access |
    configuration | config-commands | auth-proxy | ipmobile}
    {default | list-name} method1 [method2. . .]
```

The authorization proxy feature uses the **auth-proxy** keyword. Unlike other AAA authorization options, this one requires the keyword **default**, not allowing the creation of a named list. Also, this feature only allows three method options: group tacacas+, group radius, and group *group-name*. The actual syntax and choices include

```
Rtr1(config)#aaa authorization auth-proxy default method1 [method2. . .]
```

The proxy authentication feature only allows three method options: group tacacs+, group radius, and group *group-name*.

group tacacs+	Uses the list of all TACACS+ servers defined with the **tacacs-server host** command to provide authorization services. The servers are tried in sequential order from the top, first entered.
group radius	Uses the list of all RADIUS servers defined with the **radius-server host** command to provide authorization services. The servers are tried in sequential order from the top, first entered.
group *group-name*	Uses the list of specified TACACS+ or RADIUS servers defined with the **aaa group server** command to provide authorization services. The servers are tried in sequential order within the group from the top, first entered.

The following example uses TACACS+ servers for both login authentication and proxy authorization.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server timeout 7
Rtr1(config)#tacacs-server key cisco-key
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#aaa authentication login default group tacacs+ local secret
Rtr1(config)#aaa authorization auth-proxy default group tacacs+
```

The following example uses RADIUS servers for both login authentication and proxy authorization.

```
Rtr1(config)#aaa new-model
Rtr1(config)#radius-server timeout 7
Rtr1(config)#radius-server key cisco-key
Rtr1(config)#radius-server host 192.168.1.4
Rtr1(config)#aaa authentication login default group radius local secret
Rtr1(config)#aaa authorization auth-proxy default group radius
```

The following example uses the three TACACS+ servers used earlier and selects only two of them to be included in the server group.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server timeout 7
Rtr1(config)#tacacs-server key cisco-key
Rtr1(config)#tacacs-server host tac-serv1
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#tacacs-server host 192.168.6.4 port 1500 timeout 3 key cisco9
Rtr1(config)#aaa group server tacacs+ tac-1
Rtr1(config-sg-tacacs+)#server tac-serv1
Rtr1(config-sg-tacacs+)#server 192.168.6.4
Rtr1(config)#aaa authentication login default group tac-1 local secret
Rtr1(config)#aaa authorization auth-proxy default group tac-1
```

Activate Authentication Proxy Accounting

The **aaa accounting** command enables accounting and creating named method lists to define specific accounting methods on a per-line or per-interface basis. Use the **global configuration mode** command **aaa accounting auth-proxy** to activate the security server that will monitor the accounting information. The router reports these attributes as accounting records, which are then stored in an accounting log on the security server. Use the no form of this command to disable accounting. The basic syntax is

Rtr1(config)#aaa accounting auth-proxy default {***accounting-method***} *method1* [*method2. . .*]

The accounting methods include the following four choices:

start-stop	Sends a start accounting notice at the beginning of a process and a stop accounting notice at the end of a process. The start accounting record is sent in the background. The requested user process begins regardless of whether the start accounting notice was received by the accounting server.
wait-start	Sends both a start and a stop accounting notice to the accounting server. But, if the **wait-start** keyword is used, the requested user service does *not* begin until the start accounting notice is acknowledged.

stop-only	Sends a stop accounting notice at the end of the requested user process.
None	Disables accounting services on this line or interface.

For minimal accounting, use the **stop-only** keyword. For more detailed accounting, include the **start-stop** keyword. For even more accounting control, include the **wait-start** keyword, which ensures the start notice is received by the RADIUS or TACACS+ server before granting the user's process request.

Unlike other AAA authorization options, this one requires the keyword **default**, not allowing the creation of a named list. Also, this feature only allows three method options: group tacacas+, group radius, and group *group-name*.

group tacacs+	Uses the list of all TACACS+ servers defined with the **tacacs-server host** command to provide authorization services. The servers are tried in sequential order from the top, first entered.
group radius	Uses the list of all RADIUS servers defined with the **radius-server host** command to provide authorization services. The servers are tried in sequential order from the top, first entered.
group *group-name*	Uses the list of specified TACACS+ or RADIUS servers defined with the **aaa group server** command to provide authorization services. The servers are tried in sequential order within the group from the top, first entered.

The following example uses TACACS+ servers for both login authentication and proxy authorization.

```
Rtr1(config)#aaa new-model
Rtr1(config)#tacacs-server timeout 7
Rtr1(config)#tacacs-server key cisco-key
Rtr1(config)#tacacs-server host 192.168.1.4
Rtr1(config)#aaa authentication login default group tacacs+ local secret
Rtr1(config)#aaa authorization auth-proxy default group tacacs+
Rtr1(config)#aaa accounting auth-proxy default start-stop group tacacs+
```

ACL Entry for Return Traffic from the AAA Server

Typically, it's necessary to create an ACL entry to allow the AAA server TACACS+ or RADIUS return traffic to get to the firewall. If CBAC has already been configured, an input ACL should already be implemented on an interface. Because two ACLs can't be on the same interface monitoring traffic that's traveling in one direction, it's necessary to add the appropriate entries to that ACL.

If there's no existing ACL, consider the following

- The *source* address is the IP address of the AAA server.
- The *destination* is the IP address of the router interface nearest to the AAA server.
- Do you want to permit ICMP traffic?
- Block all other traffic.
- Apply the ACL inbound on the interface connected to the AAA server.

The syntax of the specify entry required for the return traffic looks like the following:

```
Rtr1(config)#access-list acl# permit tcp host source eq tacacs host dest
```

In the following example, ACL 105 blocks all inbound traffic on interface Ethernet0/0, except for traffic from the AAA server, assuming that traffic from the AAA server would come in through Ethernet0/0. The permit ip any any statement allows any ICMP traffic through.

```
interface Ethernet0/0
  ip address 192.168.1.1 255.255.255.0
  ip access-group 105 in
!
access-list 105 permit tcp host 192.168.1.20 eq tacacs host 192.168.1.1
access-list 105 deny   tcp any any
access-list 105 deny   udp any any
access-list 105 permit ip any any
```


Configuring the HTTP Server

To use authentication proxy, the HTTP server on the firewall must be configured and the HTTP server authentication method set to use AAA. You can accomplish this in three basic steps. The first two are required, while the third offers some optional capabilities.

1. Enable the HTTP server on the router. The authentication proxy process uses the HTTP server features to communicate with the client for user authentication. The syntax is

   ```
   Rtr1(config)#ip http server
   ```

2. Set the HTTP server authentication method to AAA. The following output shows the authentication options. This isn't a method list, so only one can be chosen. The final line demonstrates the AAA implementation. The syntax is

   ```
   Rtr1(config)#ip http authentication ?
     aaa     Use AAA access control methods
     enable  Use enable passwords
     local   Use local username and passwords
     tacacs  Use tacacs to authorize user
   Rtr(config)#ip http authentication aaa
   ```

3. (Optional) Specify an access list for the HTTP server. This standard access list can identify which IP address(es) can access the HTTP session. If used, the *acl#* is the same as the one applied to the interface in the **ip auth-proxy name** ***auth-proxy-name*** **http** command in the next section. The syntax is

   ```
   Rtr1(config)#ip http access-class acl#
   ```

The following example enables the HTTP server on the router using the AAA authentication. The ACL 75 denies all host connections to the HTTP server.

```
Rtr1(config)#ip http server
Rtr1(config)#ip http authentication aaa
Rtr1(config)#ip http access-class 75
Rtr1(config)#access-list 75 deny any
```

Authentication Proxy Configuration on the Router

Configuring authentication proxy on a firewall router involves using three variations of the **ip auth-proxy** command in the Global Configuration mode. Then, the same **ip auth-proxy** command is used in the Interface Configuration mode to apply the rule to the interface. The variations include the following commands:

- **ip auth-proxy auth-cache-time**
- **ip auth-proxy auth-proxy-banner**
- **ip auth-proxy name**
- **ip auth-proxy** (interface configuration)

Remember, the **aaa authorization auth-proxy** command, covered earlier, is used together with the **ip auth-proxy name** command covered in this section. The former enables the AAA auth-proxy feature, while the latter defines the specific rules to apply. Together, these two commands create the authorization policy to be used by the firewall.

The ip auth-proxy auth-cache-time Command

Use the **global configuration ip auth-proxy** command with the auth-cache-time option to set the authentication proxy global idle timeout value. This *cache idle timer* monitors the length of time (in minutes) that an authentication cache entry, along with its associated dynamic user access control list entry, is managed after a period of inactivity. When that period of inactivity expires, the authentication entry and the associated dynamic access list entries are deleted.

If CBAC is also configured on the router, the auth-cache-time timeout value must be set higher than the idle timeout for any context-based access control protocols. If not, when the authentication proxy timer expires and removes the user profile and any associated dynamic user ACLs, idle connections could be monitored by CBAC. Deleting these ACL entries can cause the idle connections to hang.

The reverse isn't a problem when the CBAC idle timeout value is shorter. CBAC always resets the idle connections whenever the CBAC idle timeout expires, which would then be before the authentication proxy removes the user profile. Use the no form of the command to restore the default. The syntax is

```
Rtr1(config)#ip auth-proxy auth-cache-time min
Rtr1(config)#no ip auth-proxy auth-cache-time
```

auth-cache-time *min*	Specifies the minutes of inactivity an authentication cache entry and the associated dynamic user ACL entry will exist before being deleted. Acceptable values: 1 to 2,147,483,647 minutes.

This command was introduced in IOS 12.0(5)T. The default value is 60 minutes.

The following example sets the authorization cache idle timeout to 20 minutes.

```
Rtr1(config)#ip auth-proxy auth-cache-time 20
```

The ip auth-proxy auth-proxy-banner Command

To display a banner, such as the router name, in the authentication proxy login page, use the **ip auth-proxy auth-proxy-banner** command in Global Configuration mode. Use the ***banner-text*** argument to specify a customer banner other than the default. To disable the banner display, use the no form of this command. The syntax is

Rtr1(config)#ip auth-proxy auth-proxy-banner [*banner-text*]
Rtr1(config)#no ip auth-proxy auth-proxy-banner [*banner-text*]

banner-text	(Optional) Specifies a text string to replace the default banner, which is the name of the router. This command is similar to the standard IOS message of the day (MOTD) command. The text string should be written in the following format: *C banner-text C* , where *C* is a delimiting character. This delimiting character can be any character that won't be used in your message.

This command was introduced in IOS 12.0(5)T. By default, the command is disabled. If the command is issued without the banner-text option, the router host name will appear.

The first example adds the router name to the display in the authentication proxy login page. The second example personalizes the message using the # as the delimiter. The show run output shows the result regardless of the delimiter selected.

```
Rtr1(config)#ip auth-proxy auth-proxy-banner
Rtr1(config)#ip auth-proxy auth-proxy-banner # Jerri's Perimeter Router #
Rtr1#show run
!
ip auth-proxy auth-proxy-banner ^C Jerri's Perimeter Router ^C
ip auth-proxy auth-cache-time 20
!
```

The ip auth-proxy name Command

The **ip auth-proxy name** command is used to create a named authentication proxy rule. The command includes several options that increase the administrator's ability to control access. The rule is applied to an interface on a router using the **ip auth-proxy** command. To create an authentication proxy rule, use the **global configuration mode** command **ip auth-proxy name**.

Use the no form of this command with a rule name to remove the specific authentication proxy rules. If no rule name is specified, the no form of this command removes

all the authentication rules on the router and disables the proxy on all interfaces. The syntax is

```
Rtr1(config)#ip auth-proxy name auth-proxy-name http [list {acl# | acl-name}]
[auth-cache-time min]
Rtr1(config)#no ip auth-proxy [name auth-proxy-name]
```

auth-proxy-name	Associates a name, up to 16 alphanumeric characters long, with an authentication proxy rule.
http	The protocol that triggers the authentication proxy—only HTTP is supported.
list {*acl#* | *acl-name*}	(Optional) Uses a standard (1–99), extended (100–199), or named ACL with the authentication proxy rule to control which hosts can use the authentication proxy. Traffic not matching the ACL is blocked. If no list is specified, all HTTP traffic connections arriving at the interface are subject to authentication.
auth-cache-time *min*	(Optional) Overrides the global default authentication proxy cache timer for the named authentication proxy rule. Values can be 1 to 2,147,483,647. The default value is set by the **ip auth-proxy auth-cache-time** command.

This command was introduced in IOS 12.0(5)T. Support for named and extended access lists was added with the IOS 12.2 release. Until then, only standard ACLs could be used.

The first example creates the authentication proxy rule **net_users**. Because no ACL is specified in the command, all connection sessions using HTTP traffic will trigger the authentication window and process.

```
Rtr1(config)#ip auth-proxy name net_users http
```

The next example creates the authentication proxy rule **sales_users**. The **auth-cache-time 15** sets the idle timer for this rule only to 15 minutes. The "*list 77*" reference and related ACL defines the hosts, the 192.168.3.0 network that is allowed to authenticate. All other traffic is blocked.

```
Rtr1(config)#ip auth-proxy name sales_users http auth-cache-time 15 list 77
Rtr1(config)#access-list 77 permit 192.168.1.0 0.0.0.255
```

In this last example, the first command disables only the **sales_users** rule, while the last line disables the authentication proxy on all interfaces and removes any rules from the router configuration, as well as any related ACL entries.

```
Rtr1(config)#no ip auth-proxy name sales_users
Rtr1(config)#no ip auth-proxy
```

The auth-proxy Interface Configuration

Use the **interface configuration** command **ip auth-proxy** to apply the named authentication proxy rule to a firewall interface. Connection-initiating HTTP traffic passing into the interface is intercepted for authentication if no existing authentication cache entry exists.

If no ACL is defined as part of the rule, traffic from all hosts is given the opportunity to authenticate. If an ACL is referenced in the rule, only those IP addresses of hosts that match the ACL is allowed to attempt to authenticate. Traffic from all other addresses is discarded.

Use the no form of this command with a rule name to disable that specific authentication proxy rule on the interface. If a rule isn't specified, the no form of this command disables the authentication proxy on the interface.

Rtr1(config)#ip auth-proxy *auth-proxy-name*
Rtr1(config)#no ip auth-proxy [*auth-proxy-name*]

auth-proxy-name	The name of the authentication proxy rule to apply to the interface. The authentication proxy rule is defined using the **ip auth-proxy name *auth-proxy-name*** command.

This command was introduced in IOS 12.0(5)T. No default behavior or values.

This example applies the authentication proxy rule sales_users to interface Ethernet0.

```
Rtr1(config)#interface e0
Rtr1(config-if)#ip address 192.168.4.2 255.255.255.0
Rtr1(config-if)#ip access-group 100 in
Rtr1(config-if)#ip auth-proxy sales_users
Rtr1(config-if)#ip nat inside
Rtr1(config-if)#no shutdown
```

Verify Authentication Proxy Configuration

To check the current authentication proxy configuration, use the privileged **EXEC mode** command **show ip auth-proxy configuration**.

In this example, the global authentication proxy idle timeout is 20 minutes, the named authentication proxy rule is "sales_users," and the idle timeout value for this rule is 15 minutes. The display shows no host list is specified, meaning all connections initiating HTTP traffic at the interface are subject to the authentication proxy rule.

```
Rtr1#show ip auth-proxy configuration
Authentication global cache time is 20 minutes
Authentication Proxy Rule Configuration
 Auth-proxy name sales_users
    http list 77 auth-cache-time 15 minutes
```

The auth-proxy Cache

When the authentication proxy is in use, dynamic access lists grow and shrink as temporary authentication entries are added and deleted. After users initiate HTTP connections through the router, use the privileged **EXEC mode** command **show ip auth-proxy cache** to display the list of authentication entries.

The authentication proxy cache lists the host IP address, the source port number, the timeout value for the authentication proxy, and the state of the connection. If the authentication proxy state is HTTP_ESTAB, the user authentication was successful.

```
Rtr1#show ip auth-proxy cache
Authentication Proxy Cache
 Client IP 192.168.1.10 Port 31219, timeout 15, state HTTP_ESTAB
```

Clearing the auth-proxy Cache

To clear authentication cache entries manually from the firewall before they time out, use the **clear ip auth-proxy cache** command in privileged EXEC mode. Use the Asterisk option to delete all authentication cache entries. Enter a specific IP address to delete an entry for a single host. The syntax is

Rtr1#clear ip auth-proxy cache {* | *host-ip-address*}

Displaying Dynamic ACL Entries

When the authentication proxy is in use, dynamic access list entries are added and removed as temporary authentication sessions are added and deleted. When no open sessions exist, there won't be any dynamic entries. To display any dynamic access list entries, use the **show ip access-lists** command in privileged EXEC mode. The number of matches displayed in parentheses indicates the number of times the access list entry was used.

Both the idle timeout parameter and the **clear ip auth-proxy cache** command from the last section can cause no dynamic entries to occur in the display. The syntax to display any access lists configured on the firewall, including dynamic ACL entries, is

Rtr1# show ip access-lists

This following shows the ACL entries prior to any authentication proxy sessions.

```
Rtr1#show ip access-lists
Extended IP access list 160
 deny tcp any any eq telnet
 deny udp any any
 permit tcp any any (41 matches)
 permit ip any any
```

This next output shows the same ACL following user authentication. The first six lines represent the downloaded user profile enabling the specified features for the authenticated user. The last four lines are the same lines as in the previous example after more traffic.

```
Rtr1#show ip access-lists
Extended IP access list 160
permit tcp host 192.168.1.10 any eq 69
permit icmp host 192.168.1.10 host 192.168.4.2
permit tcp host 192.168.1.10 any eq telnet
permit tcp host 192.168.1.10 any eq ftp
permit tcp host 192.168.1.10 any eq ftp-data
permit tcp host 192.168.1.10 any eq smtp
deny tcp any any eq telnet
deny udp any any
permit tcp any any (76 matches)
permit ip any any
```


For more information, go to the Cisco web site http://www.cisco.com and search for authentication proxy. No CCO account is required.

The debug Commands

While they aren't specifically exam objectives, the following **debug** commands can be useful in troubleshooting or simply watching authentication proxy processes. As with all **debug** commands, remember, debug can have a detrimental effect on production routers.

Rtr1#debug ip auth-proxy {function-trace}	Displays the authentication proxy functions
Rtr1#debug ip auth-proxy {http}	Displays auth-proxy–related HTTP events

The following **debug** commands are AAA tools but, because authentication proxy relies on AAA, they can be useful in troubleshooting.

Rtr1#debug {tacacs | radius}	Displays TACACS+ or RADIUS information.
Rtr1#debug aaa authentication	Displays information on AAA/TACACS+ authentication. Shows the methods of authentication and the results.
Rtr1#debug aaa authorization	Displays information on AAA/TACACS+ authorization. Hows the methods of authorization and the results.

CBAC Configuration

As you learned earlier in this chapter, authentication proxy *does not create ACEs to support returning data traffic or data channel traffic.* So, it's necessary either to create static ACLs to allow the return traffic or to configure the CBAC inspection rules in the firewall configuration. Because static ACLs increase the network vulnerability (they're permanent "holes" in the firewall) and the CBAC "hole" only exists as long as needed, the solution is quite simple.

Configuring CBAC with any authentication proxy implementation is the most reliable and secure method to ensure return traffic for authorized user connections through the firewall. Chapter 6 covered CBAC in detail. The following exercise shows an example of adding CBAC to the configuration after the authentication proxy is defined.

Chapter Review

The authentication proxy is user authentication and authorization technology, which is a part of Cisco IOS Firewall feature set. The feature is supported on a growing list of platforms using the latest IOS versions (12.2), including the SOHO 70, 800, uBR900, 1720, 2600, 3600, 7100, 7200, and 7500 series routers. Earlier versions won't support the feature on the smaller units (SOHO 70 to 1720s). Authentication proxy is compatible with other Cisco IOS security features, such as NAT, CBAC, IPSec encryption, and VPN client software.

The Cisco IOS Firewall authentication proxy feature allows network administrators to implement security policies on a per-user basis through personalized ACLs. Without firewall authentication proxy, user identity and any authorized access was associated with a user's IP address. Any single security policy had to be applied to an entire user group or subnet. Now, users can be identified and authorized on the basis of their per-user policy, and any access privileges can be customized, based on their individual access profiles.

With the authentication proxy feature, users can log in to the network or access the Internet via HTTP, and their specific access profiles are automatically retrieved and applied from a Cisco Secure ACS, or other RADIUS or TACACS+ authentication server. The user profiles and the resulting temporary ACL entries are active only while active traffic exists from the authenticated user. By default, the temporary openings close after 60 minutes of inactivity.

The authentication proxy feature requires that an AAA server running Cisco Secure Access Control Server (ACS) be present on the network. Configuring the AAA server to enable the features is necessary.

Next, the router running the firewall feature set, typically the perimeter router, must be configured by performing the following tasks:

- Configuring AAA support (required)
- Configuring the HTTP server feature (required)
- Configuring the Authentication Proxy (required)
- Verifying the Authentication Proxy (optional, but valuable)

Skipping the optional verifying is done only at risk to the resources you're charged to protect.

Questions

1. What protocol does the authentication proxy use to trigger an authentication session?

   A. Telnet

   B. HTTPS

   C. TFTP

   D. HTTP

2. Which two protocol authentication servers can be used with authentication proxy?
   - **A.** TACACS
   - **B.** Kerberos
   - **C.** RADIUS
   - **D.** TACACS+
3. Authentication proxy allows how many attempts to enter a valid user name and password?
   - **A.** 1
   - **B.** 3
   - **C.** 5
   - **D.** unlimited
4. True or False. User profile entries stored on the AAA server are made up of **permit** and **deny** statements used to create temporary ACL entries on the firewall router.
   - **A.** True
   - **B.** False
5. In the temporary ACL entry **permit icmp host 192.168.1.10 host 192.168.4.2**, which address probably represents the authenticated user?
   - **A.** 192.168.1.10
   - **B.** 192.168.4.2
   - **C.** It can be either one.
   - **D.** There's no way to know.
6. True or False. Authentication proxy is supported on all router platforms since v12.2.
   - **A.** True
   - **B.** False
7. In the **ip auth-proxy auth-cache-time** ***units*** command, what are the idle timer units?
   - **A.** Bits
   - **B.** Seconds
   - **C.** Minutes
   - **D.** Hours
8. Which technology does authentication proxy use to provide secure authentication?
   - **A.** HTTPS
   - **B.** DirectX

C. JavaScript

D. All of the above

9. According to the text, which of the following is the reason to configure CBAC with authentication proxy?

A. CBAC's attack prevention features.

B. CBAC is so easy to configure.

C. Authentication proxy doesn't create ACEs to support returning data traffic.

D. CBAC has its own authentication functions.

10. The IOS Firewall authentication proxy feature works with which technology?

A. Dial-in connection

B. Console connections

C. HTTP sessions

D. Telnet sessions

11. What additional IOS feature allows the authentication proxy to work with NAT services?

A. AAA Accounting

B. CBAC

C. VPN client

D. One-time passwords

12. Which command is *not* a step in setting up an IOS Firewall authentication proxy?

A. Configuring the HTTP server

B. Configuring the authentication proxy

C. Configuring AAA

D. Configuring CBAC

13. How many open sessions does authentication proxy support before refusing additional sessions?

A. 24

B. 40

C. 100

D. 500

14. When creating the inbound ACL on the firewall for authentication proxy, all traffic is typically blocked except which one of the following that's absolutely required?

A. Outbound traffic to the AAA server

B. Local user Telnet traffic

C. Return traffic from the AAA server

D. Local user HTTP traffic

15. When configuring the authentication proxy features, all commands are variations of which of the following?

A. Rtr1#ip auth-proxy auth

B. Rtr1(config)#ip auth-proxy auth

C. Rtr1(config-if)#ip auth-proxy auth

D. Rtr1(config-ap)#ip auth-proxy auth

Answers

1. **D.** HTTP
2. **C. and D.** RADIUS and TACACS+
3. **C.** Five tries
4. **B.** False. User profiles contain only **permit** statements.
5. **A.** 192.168.1.10
6. **A.** False. IOS versions (12.2) support authentication proxy on the SOHO 70, 800, uBR900, 1720, 2600, 3600, 7100, 7200, and 7500 series routers.
7. **C.** Minutes (default 60)
8. **C.** JavaScript
9. **C.** Authentication proxy does *not* create ACEs to support returning data traffic.
10. **C.** HTTP sessions
11. **B.** CBAC
12. **D.** Configuring CBAC
13. **B.** 40
14. **C.** Return traffic from the AAA server
15. **B.** Rtr1(config)#ip auth-proxy auth

PART III

Virtual Private Networks (VPNs)

- **Chapter 9** Cisco IOS IPSec Introduction
- **Chapter 10** Cisco IOS IPSec for Pre-Shared Keys
- **Chapter 11** Cisco IOS IPSec Certificate Authority Support
- **Chapter 12** Cisco IOS IPSec Remote Access Using Cisco Easy VPN
- **Chapter 13** Cisco VPN Hardware Overview
- **Chapter 14** Cisco VPN 3000 Remote Access Networks
- **Chapter 15** Configuring Cisco VPN 3002 Remote Clients
- **Chapter 16** Cisco VPN 3000 LAN-to-LAN Networks

CHAPTER 9

Cisco IOS IPSec Introduction

In this chapter, you will learn to:

- Recognize Virtual Private Networks
- Apply tunneling protocols
- Work with IPSec
- Use Cisco IOS IPSec technologies
- Learn about Cisco IOS Cryptosystem components
- Understand security associations
- Know the five steps for IPSec
- Make use of IPSec support in Cisco Systems products

As businesses grow and create branch locations or encourage employees to work at least part-time from home, the need to maintain remote connectivity to the organization's information systems becomes a bigger concern. At the same time, many businesses are incorporating "just in time" scheduling into many of their processes, which requires vendors, suppliers, and even customers to have immediate access to production or purchasing information.

At one time, network security could be managed quite well simply by not allowing any outside access to the company resources. If information was needed, employees either reported to their desk to retrieve the information or telephoned someone at the office to get the needed data. Both of these solutions would be considered unacceptably slow today for many firms.

Earliest remote connections were often made with dedicated (leased) lines that provided quite secure connections for a price. Typically, that price went up according to distance. A branch office across town might require only a few hundred dollars per month for connectivity, whereas the same size branch across the country could cost many times as much. The two factors that most directly impacted the cost were bandwidth required and distance covered.

WAN technologies like analog modem, ISDN, T1/E1, T3/E3, frame relay, and cell relay (ATM) all offer reliable, secure connectivity with varying degrees of performance. The problem is they can be expensive. Lack of universal access to some of these technologies in many parts of the country or world means companies often have to install and

support multiple technologies. Even for companies that routinely use dedicated lines, the problem still exists of what to do with branch locations, vendors, suppliers, or traveling employees located halfway around the world.

The advent of the Internet naturally caught the imagination of many organizations—large and small—as a possible way to extend their networks for a relatively small cost. All they had to do was give up was security and often performance. The goal became to develop a way of providing a secure connection within an inherently non-secure environment, a virtual private network. Many attempts have been made to create these VPNs and many proprietary solutions are competing for this growing business. In Chapters 10 through 16, the focus is on Cisco's implementation of IPSec, an industry standard for providing private connections over public networks. Figure 9-1 shows some of the many types of connections that an organization might try to incorporate into its VPN strategy.

Virtual Private Networks

A virtual private network (VPN) is a secure connection within a public (non-secure) network. This security is typically implemented by some form of user and/or device authentication and encryption of the data stream creating a secure, private tunnel between

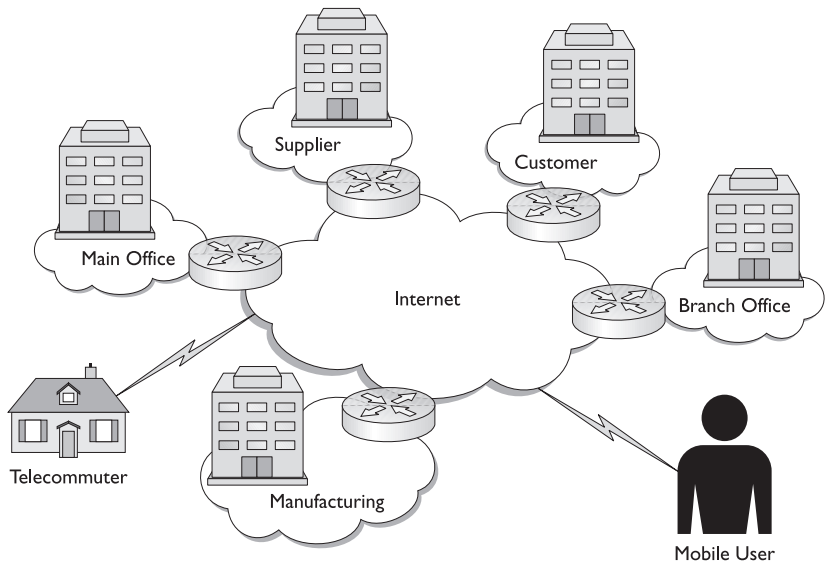

Figure 9-1 VPN opportunities to extend the corporate LAN services

a remote endpoint and a gateway. The sensitive nature of some communications requires the help of VPNs to provide the following three services:

Confidentiality	Ensures that only the intended recipient can read the transmitted data while, at the same time, thwarting efforts by other parties that might intercept it. Confidentiality is provided by encryption algorithms, such as DES or 3DES.
Authentication	Verification of the identity of a person or process that sent the data. Authentication is provided by mechanisms, such as exchanging digital certificates.
Integrity	Ensures the data received is exactly what was transmitted from the source without alterations or additions. Integrity is provided by hashing algorithms, such as MD5 or SHA.

While many VPN strategies and technologies exist, let's start by defining some basic terminology that can help describe VPNs. VPNs come in two basic types:

- Remote access
- Site-to-site

Each of these two VPN types can be further broken down into either of the following categories:

- Intranet access
- Extranet access

Remote–Access

Remote access involves connecting individual users to a LAN to provide secure, encrypted network access for telecommuters, traveling employees, and one-person offices of consultants, contractors, brokers, vendors, and so forth.

Early remote-access VPN often involved contracting (outsourcing) with enterprise service providers (ESPs), such as CompuServe. The ESP provided the network infrastructure and client software for the user computers. Many of these systems used dial-up (modem) services called virtual private dial-up networks (VPDN). ESPs were often too expensive for small to medium-sized organizations or they weren't interested in supporting that size client.

Today, many companies provide their own VPN connections through the Internet, allowing access to remote users running VPN client software through their Internet service providers (ISPs). The rapid expansion of cable modems and DSL has made it possible for telecommuters and other fixed location users to replace slower modem and ISDN services with fast connections, at a fraction of the cost of dedicated lines. Fast Internet connections offered in many hotels and the new wireless access facilities in many public places, such as airports and convention centers, means traveling employees can also use fast secure remote VPN connections. Examples in Figure 9-2 show two common types of remote-access VPN.

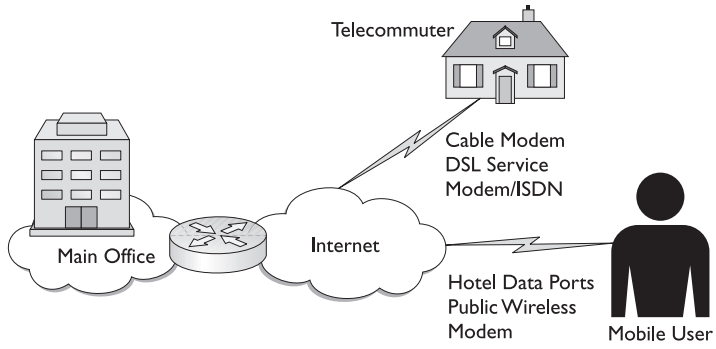

Figure 9-2 Remote-access type VPN connections

The three basic types of VPN connections include

- Access VPN
- Intranet VPN
- Extranet VPN

Access VPN

Access VPNs allow the organization's remote users to access company intranet or extranet over a shared infrastructure. Access VPNs provide secure connectivity for mobile users, telecommuters, and small branch offices using analog modems, ISDN, DSL, cable modems, and wireless IP.

Intranet and extranet VPNs are covered in the next paragraphs.

Site-to-Site

Site-to-site, or LAN-to-LAN, VPNs involve a secure connection between two end devices such as routers, firewalls, or VPN hardware devices. The hosts on each LAN connected to those end devices can access the other LAN via the secure connection based on the organization security policy and the placement of shared resources. Common examples of site-to-site VPN implementation could include connecting branch offices, vendor sites, dealer sites, or customer offices to the corporate network. Figure 9-3 shows the types of connections that might be VPN candidates.

Intranet VPN

Intranet VPN would involve allowing the remote user or connected site to have access to the company internal network and resources. A typical example might be a branch office connecting to the corporate network allowing all branch employees access to e-mail and other corporate resources. Individual telecommuters and traveling employees would be candidates to use a VPN to connect to the company intranet.

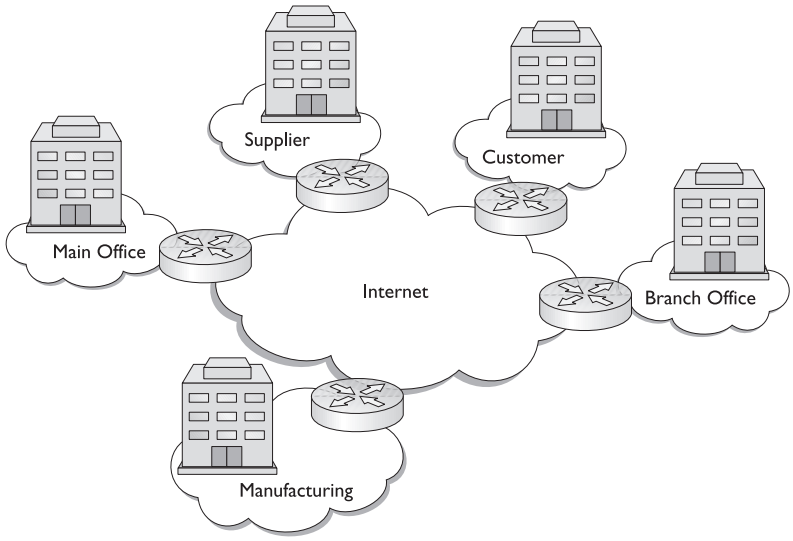

Figure 9-3
An example of site-to-site VPN connection

Extranet VPN

Extranet VPN might be one or more special networks established to share resources with vendors, suppliers, customers, consultants, business partners, and other nonemployee groups. The extranet creates a shared environment for collaborative efforts. An example might be a company web server network that allows dealers to check inventory, place orders, and track deliveries. This limits access and exposure to only those resources needed by the shared, while protecting the others.

Extranet access could be either site-to-site or a remote-access connection to a nonemployee, such as a consultant or a broker.

The extranet could be a DMZ on the network that requires some level of authentication to access and is, therefore, unavailable to the general public. If anyone could access it, this would be an Internet. Figure 9-4 shows VPN connection types supported by Cisco technologies.

Layer 2 VPNs

In an effort to provide secure connections through the public networks, service providers developed their own proprietary technologies or combinations of technologies. Some of these early VPN implementations were Layer 2 technologies, while others were Layer 3.

Using the broad definition of a VPN as "a private network built upon or within a public network," public X.25, Frame Relay, and ATM networks' use of virtual circuits within their public networks can be considered VPNs. The shared nature of these technologies allowed them to be less expensive than comparable dedicated circuits. These types of VPNs are generically referred to as Layer 2 VPNs. Unfortunately, these services don't address universal access because they are unavailable in many parts of the country and world.

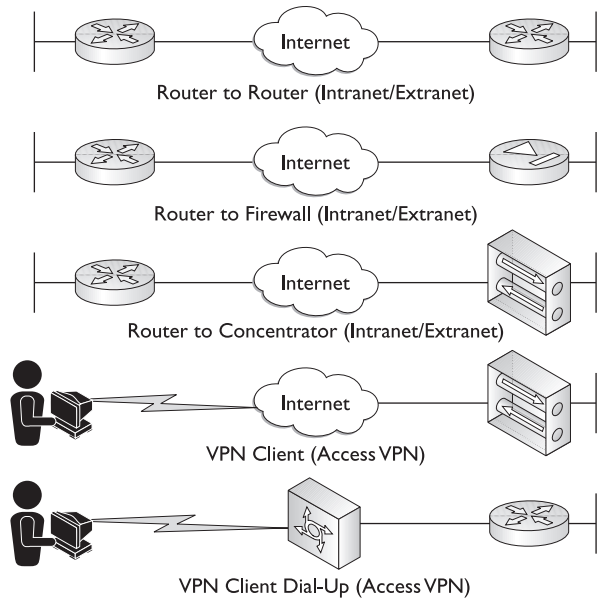

Figure 9-4 VPN connection types supported by Cisco devices

Layer 3 VPNs

The dominant emerging form of VPNs are those networks constructed across shared IP backbones, called *IP VPNs.* Because the Internet is the largest and most widely accessible of the public networks, this is where the greatest research and development is aimed. The early Layer 3 VPN implementations were provided by private companies that developed security implementations on top of the published TCP/IP and Internet standards. Cisco encryption technology (CET) was an early Cisco proprietary Layer 3 VPN technology.

The biggest drawback to these early efforts was the lack of interoperability between different manufacturers. While a strong case can be made for a single vendor end-to-end solution, the reality is this: all-too-common business mergers and acquisitions often mean companies are forced to merge different vendor VPN strategies and technologies.

A second issue is that not all vendors make products for all implementations within the network. For example, a vendor might have a strong VPN line for connecting branch locations and could even have client software for individual remote users, but they might lack a solution for small multiuser connections using cable modems or DSL service.

IPSec as an Emerging Leader

IP Security (IPSec) is a standards-based suite of protocols developed by the Internet Engineering Task Force (IETF) to provide secure exchange of packets at the IP layer (Layer 3). IPSec is rapidly becoming the most widely deployed VPN implementation. Cisco has adopted IPSec for its VPN products.

The single biggest problem with using the Internet—or any TCP/IP network—for private communication is the lack of security. The underlying protocols simply weren't designed with security as a high priority. While it's easy to place blame using hindsight, a fair analysis would recognize that not only was the technology brand new, but also that

no one involved could have visualized the masses from every corner of the Earth individually accessing the resulting network.

The evolution of the World Wide Web as a more or less unregulated playground for every interest and activity has lead to a growing number of miscreants bent on causing problems. Add to this the technologically incompetent and those who see the Internet as a tool for political and religious warfare, and you can understand why the neighborhood has become an unfriendly place.

The security concerns in using the Internet for conducting private communications can fall into the following categories:

- Loss of privacy
- Loss of data integrity
- Identity spoofing
- Denial-of-service

The ultimate goal in developing an IPSec standard is to address these threats without the need for expensive host hardware or application modifications and changes.

IPSec and the Future

Remember, IPSec is a new and still-evolving pool of standards and protocols. The IETF working group for IPSec has dozens of draft proposals they're working on to extend the capabilities and interoperability of IPSec with other common network technologies, such as NAT/Firewall traversal and MIB standards.

The mandate is for IPSec to be in IP Version 6 when it's finally implemented. Once this occurs, we'll all be using IPSec in its latest form. For more information on the IPSec working group and its current activities go to the following site: http://www.ietf.cnri.reston.va.us/html.charters/ipsec-charter.html.

NOTE While this book is current at press time, probably no other features warrant constant monitoring of Cisco documentation and releases covering IOS and device upgrades more than IPSec and VPNs. Each new release seems to bring expanded support for technologies like wireless, firewalls, and even basic IPSec protocols supported.

Other VPN Implementations

Data security can be performed at many levels, including having the user applications encrypt the data. This was a feature supported on many user applications, such as Microsoft Excel and Word. In each case, the users supplied a password when they saved a document and, next, the data was scrambled using the password as an encryption key. The recipient then needed to supply exactly the same password to open the document. While somewhat effective, this method relied on the user to implement it, forward the password in a secure fashion to the recipient, and never forget the password for the life of the document. Many corporate disasters occurred when an employee left the company, for whatever reason, and critical data was unavailable.

In an effort to remove user involvement from the process, technologies were developed that operated at the application layer, but remained invisible to the end user. Some of these applications provided only partial solutions to the problem. Secure Sockets Layer (SSL) is an example of application encryption for web browsers that protects the confidentiality of data sent from supported applications, but it doesn't protect data sent from other applications. Each host system on the network and all applications must be SSL installed and configured to work efficiently. While this is better than having the user manage encryption, too many opportunities still exist for the system to fail. SSL typically requires increased administrator involvement when computers or applications are added or changed.

Why Use VPNs?

VPN benefits to the organization can include significant cost savings because the service provider network supplies the brunt of the hardware and support for the WAN connections. The savings can be substantial, especially when comparing the price of bandwidth available from DSL and cable connections to the Internet versus the cost of dedicated or frame relay links.

Particularly in supporting remote users, the Internet provides a low-cost alternative for enabling access to the corporate network, instead of maintaining large modem banks and paying costly phone bills. With just a local phone call to an ISP in many places around the world, a user can have access to the corporate network. In some cases, VPNs enable organizations to save between 30 and 70 percent over competing remote-access solutions.

An organization might see improved geographical coverage, especially in communities or regions where dedicated links or frame relay links are unavailable or prohibitively expensive. The rapid deployment of cable modems, DSL service, and even some public wireless networks can allow the telecommuter opportunities that would be impractical at modem speeds. High-speed data ports in hotels and motels that can be secured through VPN connections mean traveling employees can often do more tasks on the road than simply checking their e-mail. In many cases, VPNs mean simplified WAN operations and increased networking reliability.

VPN Analogy

A network with remote sites can be thought of as a group of islands served by a public ferry system. While this internetwork of ferries allows residents to move from place to place at a reasonable cost, there's no privacy and any scheduling, routing, and rules are up to the ferry provider.

An alternative would be to build bridges among the islands, but the cost would be high and some islands might have too few residents to make the cost worthwhile. These dedicated links, as in WANs, would provide privacy and more freedom of movement. Unfortunately, like WANs, where distance determines the cost of the bridge, it's possible to create situations where connecting a distant small site with even minimal service could cost many times that of servicing a much larger and closer group.

Another alternative could be personal watercraft of some type. These boats could provide private transportation, routing, and scheduling at the direction of their owner. Even if the people on the ferry see a boat, they have no idea what's going on inside, or even its source or destination. Relative to building bridges, adding more boats as needed would be inexpensive.

With all the variety in boat types—from canoes and kayaks to luxury yachts—at least it's possible that some marinas might not support all types, thereby limiting the boat owner's ability to make connections. If a single standard was adopted by all marinas, such as power boats between 15 and 35 feet, then the boat owner would be free to choose from many vendors as long as they met the standard. The marinas, as service providers, could still support other options in addition to the standard.

Tunneling Protocols

Most VPNs use the concept of tunneling to create a private network that extends across the Internet. Conceptually, it's as if a secure tunnel has been built between two end devices (routers, firewall, or VPN device). Data can be directed into one end of the tunnel and it travels securely to the other end. These end devices, or tunnel interfaces, are typically the perimeter router firewalls for the LANs being connected.

Technically, no tunnel exists and the process doesn't resemble a tunnel, but the term "tunneling" somewhat describes the end result of traffic being able to pass through a non-secure environment without concerns about eavesdropping, data hijacking, or data manipulation. *Tunneling* is a process of encapsulating an entire data packet as the payload within a second packet, which is understood by the network and both end points. Depending on the protocols used, the new payload—the original packet—can be encrypted. Figure 9-5 is a common graphical representation of Layer 3 tunneling technology.

The tunneling process requires three different protocols:

- **Carrier protocol** The network protocol used to transport the final encapsulation
- **Encapsulating protocol** The protocol used to provide the new packet around the original data packet. Examples: IPSec, GRE, L2F, L2TP, PPTP
- **Passenger protocol** The original data packet that's been encapsulated. Examples: IP, IPX, NetBEUI

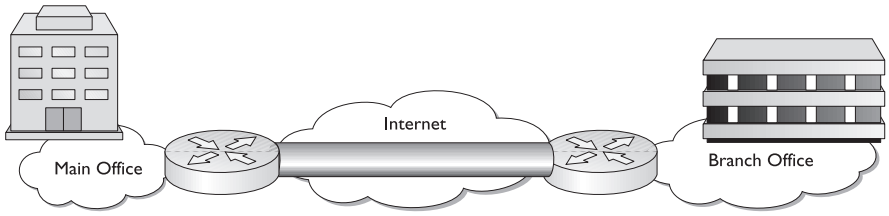

Figure 9-5 Layer 3 VPN tunneling representation

Through tunneling techniques, you can pass non-IP packets or private IP addressed packets through a public IP network. You can even route NetBEUI—the famous non-routable protocol—once it's been encapsulated for tunneling through a VPN. What happens is the new data frame, or packet, is, in fact, a legal packet with proper addressing to travel through the network. Hidden safely within the payload portion of this new frame is the original packet, which needs the assistance and/or protection.

L2F, L2TP, and PPTP are all three Layer 2 tunneling protocols that support Access VPN solutions by tunneling PPP.

Layer Two Forwarding (L2F) Protocol

L2F is a tunneling protocol developed by Cisco Systems, which is similar to PPTP developed by Microsoft. Both protocols enable organizations to set up VPNs that use the Internet backbone for transporting traffic. L2F is supported by other vendors, such as Shiva and Nortel.

Recently, Microsoft and Cisco agreed to merge their respective protocols into a single, standard protocol called Layer Two Tunneling Protocol (L2TP) an IETF and industry-standard Layer 2 tunneling solution. Microsoft supports L2TP in Windows 2000/XP client software for client-initiated VPN tunnels.

Layer 2 Tunneling Protocol (L2TP)

L2TP is an emerging IETF standard and one of the key building blocks for VPNs in the dial access space. L2TP combines the best features of Cisco's Layer 2 Forwarding (L2F) and Microsoft's Point-to-Point Tunneling Protocol (PPTP), enabling mobile workforces to connect to their corporate intranets or extranets wherever and whenever they require.

L2TP is a standard way to build Access VPNs that simulate private networks using a shared infrastructure, such as the Internet. These Access VPNs offer access for mobile users, telecommuters, and small offices through dial, ISDN, xDSL, and cable.

Benefits of L2TP include per-user authentication, dynamic address allocation from an address pool or by using DHCP server, plus RADIUS and AAA support.

Generic Routing Encapsulation (GRE)

GRE is an early Layer 3 tunneling technology that has existed for years. Cisco has supported this tunneling technology since Cisco IOS software version 9.21. IPSec is the new IETF standard for Layer 3 encryption and encrypted tunnels supported by the Cisco IOS software since version 11.3(3)T.

How IPSec Works

IPSec is a complex method of exchanging data that involves many component technologies and numerous encryption-method options. This chapter covers the major protocols and processes that make up IPSec. Configuring these protocols and processes is covered

in Chapters 10 through 16. Recognizing that IPSec operation can be broken down into the following five main steps might be useful.

Step 1	Traffic deemed "interesting" initiates an IPSec session. Access lists determine which data traffic will be protected by the IPSec technology. Other traffic will travel through the public network without IPSec protection.
Step 2	Called *IKE Phase One* A secure channel is established between two peers. Each peer determines whether the other peer can be authenticated and whether the two peers can agree on Internet Key Exchange (IKE) security rules for exchanging data.
Step 3	Called *IKE Phase Two* IPSec session is established between the two peers. Stricter IPSec security rules and protocols are established between the peers.
Step 4	Data transfer Data is transferred between IPSec peers using the IPSec-defined security rules and protocols.
Step 5	*IPSec tunnel termination* IPSec session ends through deletion or by timing out.

The complexity of VPNs in general and IPSec in particular can get a little intimidating, but, remember, at the highest level, this is just like many other communications sessions. Some data requires special attention, a session is opened, the data is exchanged, and the session is torn down. Even a simple telephone call to someone special can be an analogy.

- Step 1 Something important occurs that can't wait until the next time you speak to that person.
- Step 2 The telephone call is placed and, through the ringing and answering, you can determine you're talking to the person you expected.
- Step 3 Because the subject matter is private and important, you might ask if the person is alone, so they can speak freely.
- Step 4 Once an acceptable level of privacy is assured, the information can be shared.
- Step 5 When the information has been exchanged, both parties hang up.

These steps are revisited again in the section "IKE SAs versus IPSec SAs."

Cisco IOS IPSec Technologies

IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec peers, such as VPN software clients, VPN 3002 hardware clients, Cisco routers, PIX Firewalls, and the VPN 3000 Series Concentrators. Cisco IPSec VPNs provide the following network services:

- **Data integrity** The IPSec receiver can authenticate packets sent by an IPSec sender to make sure the data hasn't been modified during transmission.

- **Data origin authentication** The IPSec receiver can authenticate the source of any IPSec packets received. This service is dependent on the data integrity service.
- **Anti-replay** The IPSec receiver can detect and reject replayed packets.
- **Data confidentiality** IPSec can encrypt data packets before transmitting them across a network in a form that can only be unencrypted by the destination target.

IPSec Security Overview

IPSec is a framework of several open standards that provides data confidentiality, data integrity, and data authentication in exchanges between participating VPN peers at Layer 3. IPSec technologies can be used to authenticate and protect data flows between IPSec peers. The two main IPSec security protocols are the following:

- Authentication Header (AH)
- Encapsulating Security Payload (ESP)

In addition, IPSec uses other existing encryption standards to complete the protocol suite. These encryption standards are covered in the section "Other IPSec Encryption Standards."

Authentication Header (AH)

The Authentication Header (AH) security protocol provides only data authentication and integrity for IP packets forwarded between two systems. It does *not* provide encryption services for data confidentiality.

The AH authentication is provided by both ends of the tunnel performing a one-way hash calculation on the packet using a shared key value. The function of a message digest hash is never to be decrypted, but to provide a check value for the receiving party to verify data hasn't been modified. While performing a hash calculation isn't difficult, trying to modify the packet and end up with exactly the same resulting message digest is difficult. The hashing algorithm was designed to be nonpredictable and, therefore, hard to duplicate.

The fact that creating the two one-way hash calculations involves using a shared secret key known to the two systems means authenticity is guaranteed.

AH can also help to implement antireplay protection by requiring the receiving host to set the replay bit in the header to indicate the packet was seen.

AH Authentication and Integrity The AH authentication and integrity function is applied to the entire IP packet, except for those IP header fields that must change in transit; such as the Time To Live (TTL) field, which is decremented by the routers along the path. Figure 9-6 shows the AH process diagrammatically. The AH process works as follows:

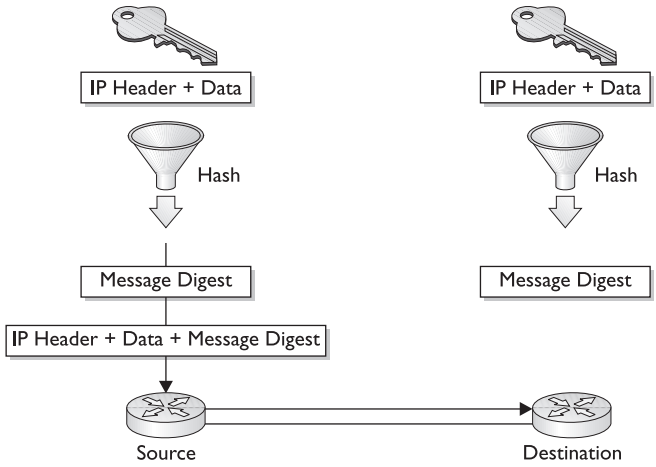

Figure 9-6 AH authentication and integrity process

Step 1	The IPSec source device performs a hash function on the IP header and data payload using a key value known also by the destination device.
Step 2	The resulting message digest is appended to the original packet as a part of creating a new packet for transit.
Step 3	The new packet is transmitted to the IPSec peer device.
Step 4	The receiving device separates the original IP header and data payload from the message digest and it performs a hash function on the IP header and data using the key value shared with the source. If the resulting message digest doesn't exactly match the one received from the source, it's rejected by the destination device.

Even a single bit change or substitution in the transmitted packet creates a different hash result and can cause the packet to be discarded. Even if the packet were captured in transit and the IP packet information was in Cleartext, the hacker doesn't know the shared key.

Encapsulating Security Payload (ESP)

The *ESP security protocol* provides confidentiality via encryption, data origin authentication, data integrity, optional antireplay protection, and limited traffic flow confidentiality by defeating traffic flow analysis. Authentication and integrity can be provided via the same algorithms used by AH. Confidentiality can be implemented independent of the other services.

The ESP confidentiality is accomplished by performing encryption at the IP packet layer. ESP supports a variety of symmetric encryption algorithms, but the default for IPSec is 56-bit DES. This particular cipher must be implemented to conform to the IPSec standard and to ensure interoperability with other vendor IPSec products. Cisco products support DES plus 3DES for even stronger encryption.

DES Encryption Algorithm Data Encryption Standard (DES) is a popular symmetric-key encryption method that uses a 56-bit key to ensure secure, high-performance encryption. The first public encryption standard, DES is based on an algorithm developed by IBM. DES is used to encrypt and decrypt packet data turning cleartext into ciphertext via an encryption algorithm. The receiving device uses a decryption algorithm and the same shared key value to restore the cleartext. Figure 9-7 shows the encryption process.

Specialized "DES cracker" machines, while uncommon, can recover a DES key after only a few hours, so Cisco recommends 3DES as the main encryption algorithm for VPN.

Triple DES Algorithm (3DES) Cisco products implementing IPSec can use the Triple DES (3DES) algorithm as a much stronger encryption method. *3DES* is a variation of the 56-bit DES that breaks the data up into 64-bit blocks, and then processes each block three times, each time with an independent 56-bit key. This process effectively doubles encryption strength over 56-bit DES.

Both DES and 3DES offer adequate performance for production network applications. Now that DES/3DES encryption is available in ASIC hardware in products, such as the VPN 3002 Hardware Client Device and VPN 3000 Series Concentrators, you can add encryption to a VPN with little impact on overall system performance.

Advanced Encryption Standard (AES) AES encryption technique was recently approved as a Federal Information Processing Standard (FIPS)-approved cryptographic algorithm (FIPS PUB 197). *AES* is based on the Rijndael (pronounced *Rhine Dahl* or *Rain Doll*) algorithm, which defines how to use 128-, 192-, or 256-bit keys to encrypt 128-, 192-, or 256-bit source blocks (all nine combinations of key length and block length are possible). AES offers greater flexibility than even 3DES because it supports multiple key sizes and multiple encoding passes.

Release 3.6 of Cisco VPN products introduce support for AES (128 and 256 bit), providing a stronger encryption standard option and improved remote access performance for both software and hardware clients. Cisco is working with the IETF IPSec Working Group to push for a new specification outlining how AES will work within the IPSec framework.

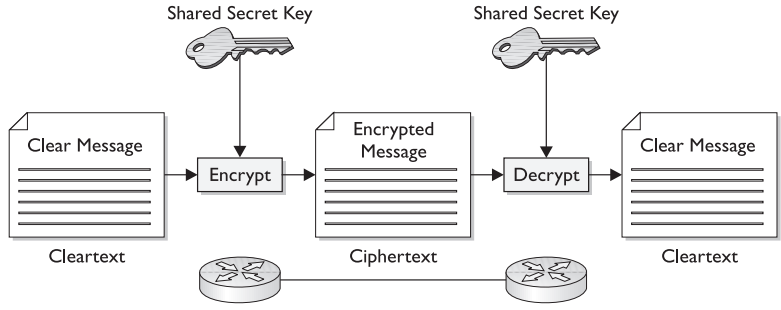

Figure 9-7 DES encryption process

Choosing Between AH and ESP

Deciding to use AH or ESP in a particular situation might seem confusing, but remember these two basics: any time more services are required, you can expect additional resources will be consumed. The same can be expected of throughput so, if a service isn't needed, don't use it. Follow both of these rules:

- If you need to make sure that data from an authenticated source is transferred with integrity, but confidentiality isn't necessary, then use the AH protocol. This saves the additional overhead and bandwidth associated with ESP encryption.
- If data confidentiality (privacy through encryption) is required, then ESP must be used. ESP can provide authentication and integrity for the packets via the same algorithms used by AH.

Other IPSec Encryption Standards

IPSec employs several different technologies to provide a complete system of confidentiality, integrity, and authenticity. These technologies are as follows:

- **Message Digest 5 (MD5)**–A hash algorithm is used to authenticate packet data.
- **Secure Hash Algorithm-1 (SHA-1)**–A hash algorithm is used to authenticate packet data.
- **Diffie-Hellman (DH)**–A key exchange standard allows two parties to establish a shared secret key to be used by encryption algorithms.
- **Rivest, Shamir, and Adelman Signatures (RSA)**–A public-key cryptographic system used for authentication.
- **Internet Key Exchange (IKE)**–A hybrid protocol that provides setup utility services for IPSec, including authentication of the IPSec peers, negotiation of IKE and IPSec security associations (SAs), and establishment of keys for encryption algorithms used by IPSec.
- **Certificate authority (CA)**–Cisco router and PIX Firewall support of CAs allows the IPSec-protected network to scale by providing the equivalent of a digital ID card for each device.

NOTE Internet Security Association Key Management Protocol (ISAKMP) is synonymous with IKE in Cisco router or PIX Firewall configuration commands. The keyword ISAKMP is always used instead of IKE.

Transport and Tunnel Mode

In configuring IPSec, one of the early decisions that must be made is whether the session is a Tunnel or a Transport mode connection. This distinction impacts other configuration decisions that have to be made.

Transport Mode

Transport mode is used between two end-host devices or between a remote host and a gateway device, where the gateway is the actual destination device. An example of a gateway device being the target destination would involve an encrypted Telnet session to configure a router or a PIX Firewall. In either case, this is basically a one-device to one-device connection. Figure 9-8 shows two possible examples of a Transport mode connection. In VPN 1, administrator Nancy must be able to access the perimeter router from home to check the status and make any configuration changes. In VPN 2, Nancy needs to access a server to make user or group account changes. In each case, a host-to-host connection exists.

NOTE Both examples are offered with the warning that these practices might be banned by a security policy. Allowing a VPN to pass through any perimeter router and/or firewall to get directly into the protected LAN is an especially risky proposition.

Transport Mode Encryption In Transport mode, if encryption is performed, only the upper-layer IP protocol fields (IP packet payload) are encrypted, leaving the IP header untouched. The IP header must be left unencrypted, so the packet can be routed through the network. Any device recording a packet in transit would be unable to read the data, but could easily determine the source and destination information.

Tunnel Mode

Tunnel mode is the most common mode involving a VPN connection between two gateway devices or a connection between an end-station and a gateway device. The Tunnel mode connection would be common between a branch office and the main office providing multiple hosts on the branch LAN to multiple shared resources on the main network. An example of the end-host to gateway VPN tunnel would be a traveling employee or telecommuter connecting to the company network to access shared resources, such as e-mail, files, printing, and so forth.

In Figure 9-9, each of the connections to the main office would normally be a Tunnel mode connection between the two routers. The mobile user or telecommuter would typically have a VPN tunnel connection from their workstation to the Main Office router.

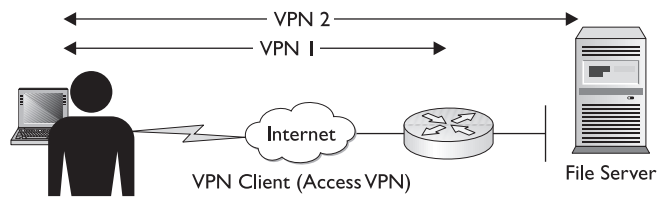

Figure 9-8 VPN Transport mode connections

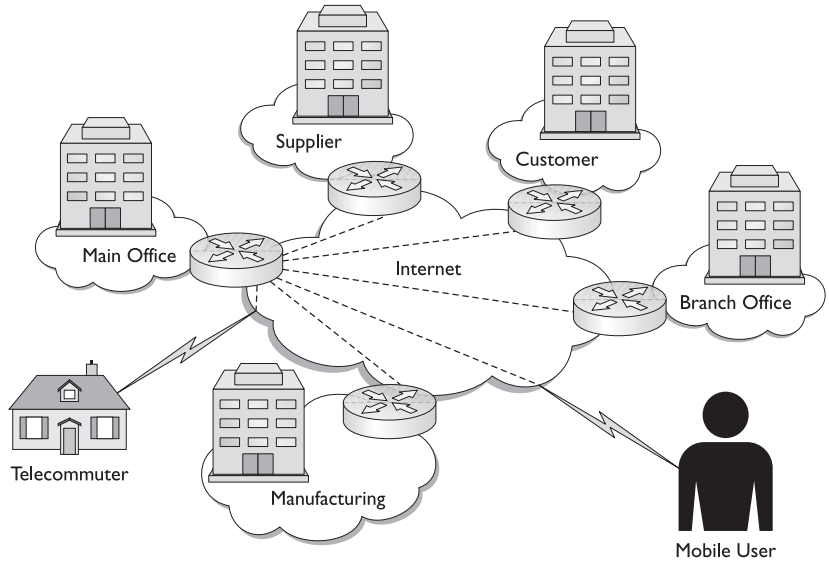

Figure 9-9 Typical Tunnel mode VPN connections

Tunnel Mode Encryption The more secure Tunnel mode encrypts both the IP header and the payload. This is possible because while the packet is in transit through the tunnel, it's fully encapsulated in a packet that uses the tunnel endpoints as the source and destination address. Any device recording a packet in transit would be unable to read any part of the original packet and could only determine the end points of the tunnel.

Tunnel Mode Benefits Tunnel mode allows a router or VPN hardware host device to act as an IPSec proxy, which means the device performs encryption services for the hosts. The Tunnel mode endpoint device is used to protect datagrams that originate from or are destined to non-IPSec host systems, making the process invisible to end users. Another great advantage is the source and the destination addresses are invisible while encrypted.

AH Transport and Tunnel Mode

Figure 9-10 compares AH Transport mode versus AH Tunnel mode. In *AH Transport mode*, AH authentication and integrity services protect the original IP packet by hashing the header fields that don't change in transit through the network, plus the data payload. Fields like TTL aren't included.

The AH header, which is the result of the hashing algorithm, is inserted after the IP header and before the data payload (higher layer fields) in the original packet. Because no encryption is involved, the IP header destination address is readable by any Layer 3 device encountered in transit and any fields, such as TTL, are still available for required changes. One advantage of Transport mode is it only adds a few bytes to each packet.

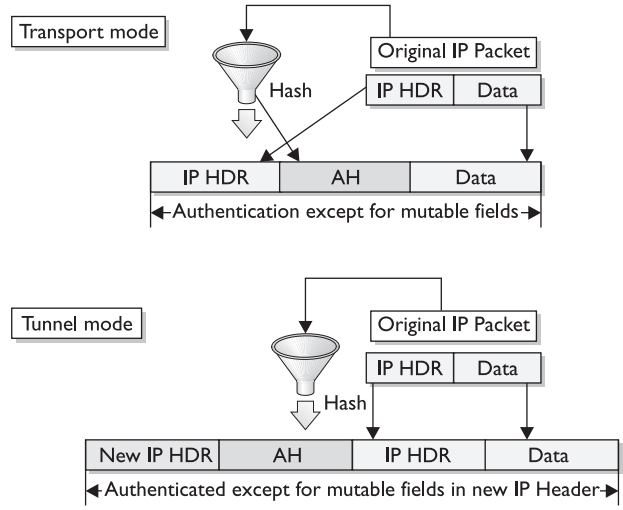

Figure 9-10 AH Transport mode versus AH Tunnel mode

In *AH Tunnel mode,* the entire original IP header and data becomes the "payload" for the new packet: a new IP header reflecting the end points of the VPN tunnel is added and the hash is performed on the resulting packet. The new IP header is protected exactly the same as the IP header in Transport mode. The fields that change in transit are excluded. The hash result again becomes the AH header following the new IP header.

Remember, both AH methods provide sender authentication and integrity verification that nothing in the data has been changed, but neither mode prevents the information from being read by an eavesdropper or a packet capture device.

A second, more serious concern is this: AH is incompatible with NAT. This is an issue if NAT occurs after the AH packet has been built because NAT changes the source IP address. This can cause the hash value stored in the AH header not to match the one generated at the destination peer and cause the packets to be rejected. This would always be a problem with AH Transport mode because the host computer is the source and, therefore, NAT would have to occur in transit. With planning, Tunnel mode would work if NAT were performed before the AH packet was built, for example, if NAT was performed at the firewall and IPSec at the perimeter router. Figure 9-11 shows how AH and NAT can be implemented to work together.

ESP Transport and Tunnel Mode

Figure 9-12 compares ESP Transport mode versus ESP Tunnel mode. In *ESP Transport mode,* the IP payload is encrypted, while the original IP header is left intact. An ESP header is inserted after the IP header and before the upper-layer protocol header (original Layer 3 data), while an ESP trailer is appended after the original data. The inserted ESP header contains a Security parameter index (SPI) value to identify the VPN security association, a sequence number field, and authentication data to verify packet authenticity.

In Transport mode, only the original data and ESP trailer fields are encrypted. The IP header isn't encrypted because some fields will be required by a Layer 3 device encountered in transit. The ESP header field isn't encrypted because it's needed by the destination peer to decipher the payload.

Figure 9-11 AH implementation to work with NAT

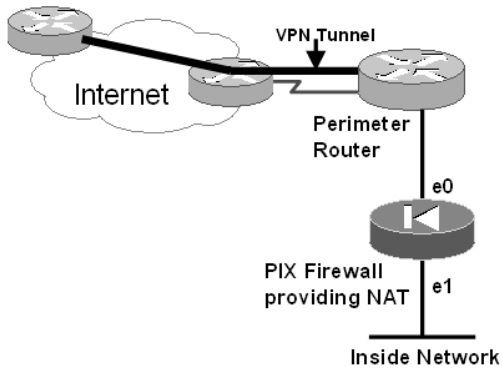

If ESP authentication is used, the upper-layer protocols (original data payload) are hashed with the ESP header and trailer, and then appended to the packet as the ESP Authentication field. Cisco IOS software and the PIX Firewall refer to this authentication service as ESP HMAC. ESP Transport mode doesn't authenticate any portion of the IP header itself.

Figure 9-12 ESP Transport mode versus AH Tunnel mode

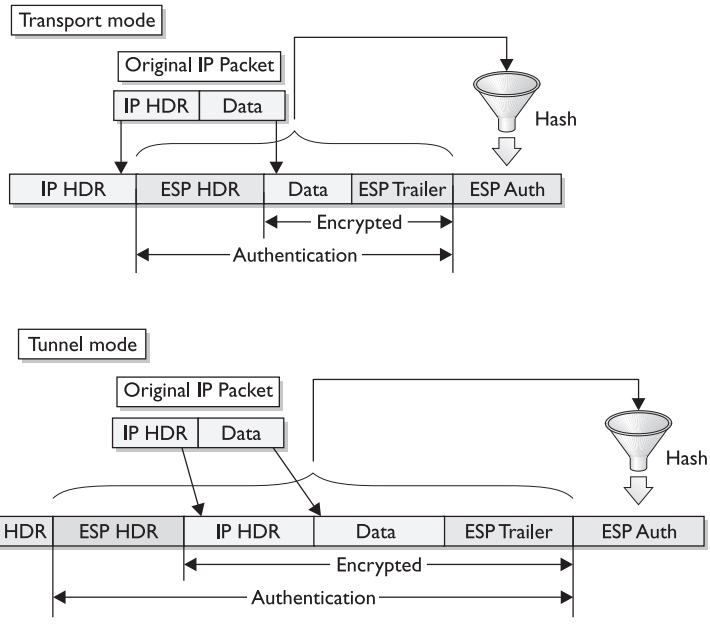

In ESP Tunnel mode, as in AH, a new IP header reflecting the end points of the VPN tunnel is added. Both encryption and authentication incorporate the entire original IP header.

When both authentication and encryption are selected, encryption is performed before authentication. This order facilitates rapid detection and rejection of replayed or bogus packets by the receiving peer, potentially reducing the impact of denial-of-service (DoS) attacks.

Choosing AH versus ESP

Both AH and ESP support both MD5 and SHA-1 hashing algorithms for authentication. The main difference between ESP and AH authentication is this: ESP doesn't protect any IP header fields in Transport mode. Both ESP and AH authenticate all IP header fields in Tunnel mode.

IPSec Transforms and Transform Sets

One set of decisions that must be made early in the IPSec implementation is whether to use AH or ESP, or MD5 or SHA-1 hashing, combined with Transport or Tunnel mode. These decisions create what are called transforms.

An *IPSec transform* defines a single IPSec security protocol—AH or ESP—with its associated security algorithms and mode. You can see the possible transform choices in Figure 9-13.

Two possible transform choices might include the following:

- AH protocol with HMAC-MD5 authentication algorithm in Transport mode, where authentication and performance are important, but encryption security isn't required.
- ESP protocol, 3DES encryption algorithm, and HMAC-SHA-1 authentication algorithm in Tunnel mode, where both data confidentiality and authentication are critical. Performance will be traded away for improved security.

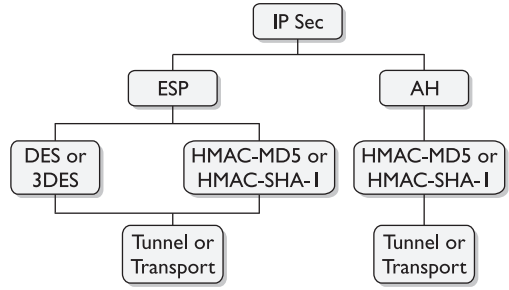

Figure 9-13 IPSec transform options

The actual transforms supported might vary by device type. VPN devices support a slightly different group than the PIX Firewall. The following are transforms supported by the IOS-based devices.

AH Transforms—Choose up to one

Transform	Description
ah-md5-hmac	AH with the MD5 (HMAC variant) authentication algorithm
ah-sha-hmac	AH with the SHA (HMAC variant) authentication algorithm
ah-rfc1828	Older version of the AH protocol (RFC 1828)

ESP Encryption Transforms—Choose up to one

Transform	Description
esp-des	ESP with the 56-bit DES encryption algorithm
esp-3des	ESP with the 168-bit DES (Triple DES) encryption algorithm
esp-rfc1829	Older version of the ESP protocol (RFC 1829). Doesn't support using ESP authentication transform
Esp-null	ESP without cipher. Can be used with esp-md5-hmac or esp-sha-hmac for ESP authentication without encryption. Shouldn't be used in production network because of the lack of security

ESP Authentication Transform—Choose up to one, only if you also selected the esp-des or esp-3des transform (not esp-rfc1829)

Transform	Description
esp-md5-hmac	ESP with the MD5 (HMAC variant) authentication algorithm
esp-sha-hmac	ESP with the SHA (HMAC variant) authentication algorithm

Transform Sets

A *transform set* is a combination of up to three individual IPSec transforms designed to implement a specific security policy for secure data transmission. The transform sets represent the choices available during IPSec security negotiation between two IPSec peers. The peers must agree to use a particular transform set for protecting a particular data flow or the exchange can't occur. Transform sets are limited to no more than one AH transform, plus no more than two ESP transforms: one for encryption and one for authentication.

Some possible examples of acceptable transform combinations include the following:

- ah-md5-hmac AH protocol with MD5 authentication
- esp-des ESP protocol with DES encryption

- esp-3des and esp-md5-hmac ESP protocol with DES encryption, plus ESP MD5 authentication
- ah-sha-hmac and esp-des and esp-sha-hmac AH protocol with SHA-1 authentication, ESP DES encryption, plus ESP SHA-1 authentication
- ah-rfc1828 and esp-rfc1829 Legacy AH protocol with ESP encryption

When configuring transform sets, the parser prevents you from entering invalid combinations. Transform sets are discussed in greater detail in Chapters 10 and 11 when configuring IPSec is covered.

Cisco IOS Cryptosystem Components

You should be familiar with two terms used in network security. While these terms are quite closely related, they aren't quite as interchangeable as they might first seem.

Encryption	The process of translating data into a secret code, which can only be read or processed after being decrypted using a secret key or password. Encryption typically involves using a mathematical algorithm for combining the original data, referred to as plaintext or Cleartext, with one or more keys—character strings of numbers and/or text known only to the sender and the recipient. The resulting output is *ciphertext*. Encryption is the most effective way to achieve data security.
Cryptography	The science or study of encryption and decryption techniques. A system for encrypting and decrypting data is called a *cryptosystem*. While it's technically possible for encrypted data to be deciphered using *cryptanalysis*, or code-breaking techniques, most modern cryptography techniques are virtually unbreakable.

The Internet, wireless technologies, and other forms of electronic communication are making electronic security technologies increasingly important. Current cryptography technologies being widely used include SSL for web traffic, Secure Shell (SSH) for Telnet sessions, Pretty Good Privacy (PGP) to protect e-mail messages, and IPSec Layer 3 secure communications.

How Encryption Works

Simply encryption uses a mathematical algorithm to scramble a message to create the cipher text. A *key* is used by the algorithm to determine the scramble pattern. The same key can then be used to unscramble or decipher the message. While this isn't a perfect analogy, a deadbolt on a door is nothing more than a chunk of metal until someone inserts a key that can manipulate the tumblers and advance the bolt. The lock and the door are then secure until the correct key is used to reverse the process. Like the door analogy, more than one person might have the key but, without the correct key, the lock won't budge. Also, somewhat like the door analogy, cipher text might require multiple keys both to lock and unlock. Figure 9-14 shows a common, simplified view of encryption.

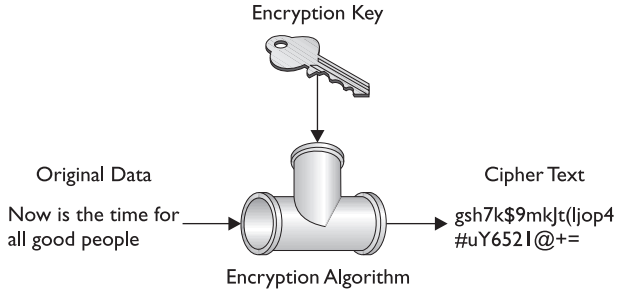

Figure 9-14 Cleartext being processed using a key to produce cipher text

The key terms in the process include data, key, encryption algorithm, and cipher text.

Data

The *data* is the information to be encrypted. This data could be text from an e-mail, text and numbers from a credit card transaction, user name and password combinations, music, graphics, videos, or voice conversations. Data arriving to be encrypted might possibly have already been encrypted one or more times in earlier processes. If cipher text is the result of three encryptions, then it must be deciphered three times in the correct order with the appropriate key(s).

Key

A *key* is a fixed-length digital sequence of characters used to scramble the source data. A longer key typically provides a higher level of security than a shorter key. DES is an industry-standard, symmetric-key encryption method that uses a 56-bit key. Triple-DES (3DES) uses three separate encryptions with unique 56-bit (168-bit) keys for stronger security. Newer AES technology uses 128-, 192-, or 256-bit keys for even greater security.

Government Regulation Recognizing that encryption keys are tightly regulated in many countries is important. Exporting from the United States or Canada technologies using encryption keys longer than 56 bits—even for use by overseas branches of domestic companies—is currently illegal. Similarly, possession of some key technologies could be a serious crime in certain countries. Export controls on strong encryption (greater than 56-bit DES strength) vary according to type, strength, destination territory, end use, and end user. A person would be foolhardy to assume that violations of import or export controls on this type of technology would be overlooked or resolved with an apology. Terrorism, drug trafficking, and racketeering concerns can make even a sloppy transaction to a third party, who then transships the technology illegally, a serious legal problem.

Encryption Algorithm

The mathematical formula used by the encryption process to scramble the data is an *encryption algorithm,* and DES and MD5 are examples. Algorithms are often used in combinations to encryption cryptosystems like IPSec.

Cipher Text

The encrypted result is virtually useless to anyone who lacks the correct key. The length of the key, and the number and types of incorporated encryptions determines the cipher text's capability to withstand unauthorized access. Ultimately, the security of a cryptosystem depends on the secrecy of one or more of the keys, rather than with the secrecy of the encryption algorithm. A strong cryptosystem has a large number of possible keys, making it time- and resource-consuming to try all possible keys (brute force). Brute force attacks, as depicted on television, where a computer generates potential keys, are typically more fantasy than reality.

Cryptography Types

Three basic types of cryptographic functions are used for authentication, integrity, and confidentiality. They include the following:

Symmetric encryption	Often called *secret key encryption*, uses a shared key and the same encryption algorithm to encrypt and decrypt a message.
Asymmetric encryption	Often called *public key encryption*, uses two different, but related, keys for the encryption algorithm. To encrypt a message, both a public key and a private key are used. The sender uses the public key to encrypt the message and the recipient uses the private key to decrypt the encrypted message.
Hash function	A *hash function* is a one-way encryption used to authenticate information. The information is "hashed," and then transmitted with the hash value. The destination peer can perform the same hash, and then compare the result to the transmitted hash value. If they match, the probability is high that the data is complete and unchanged.

Encryption Alternatives

Encryption can occur at various levels within the network to meet the organization's specific objectives. Using the Open Systems Interconnection (OSI) or the Transmission Control Protocol/Internet Protocol (TCP/IP) reference model, the encryption can be placed at the data-link layer, the network layer, or the application layer. Figure 9-15 shows a diagram of encryption at each of the layers.

Data-Link Layer

Data-link layer (L2) encryption can be implemented in devices other than the router on switched links where security is a concern. Because this method encrypts the network layer header, as well as the data, traffic must be decrypted before entering the next router. The concern would be if traffic had to pass through many routers, the decryption and reencryption would need to be repeated each time, leading to latency issues. Another concern involves trust because the data passes through these routers as Cleartext.

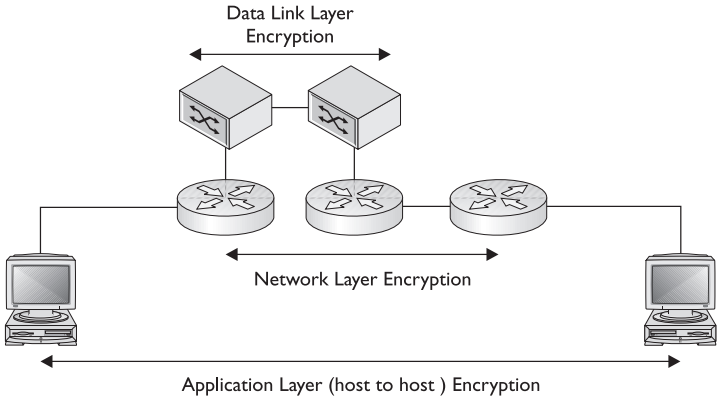

Figure 9-15 Encryption implementation options

Network Layer

Network layer (L3) encryption is done between two layer 3 devices. While Transport mode involves encrypted connections between an end host and a VPN endpoint device, such as routers on the network, the more common implementation is between two VPN endpoints, such as a VPN hardware client device, a router, a firewall, or a VPN concentrator. This means Cleartext data travels between the actual host and the VPN source endpoint where it's encrypted and forwarded on to the VPN destination endpoint, where the transmission is decrypted and forwarded to the destination host as Cleartext. One security concern is that the data is unencrypted during part of the trip, although we hope it would be in the internal secure LAN networks on both ends.

Note, the VPN endpoints might be many routers (hops) away and the interim routers needn't be VPN-aware. Because the packets must be routed through the network, the IP header must remain in Cleartext. While IPSec uses this approach, the layer 3 header information that's vulnerable refers to the VPN tunnel endpoints and not to the source and destination hosts. The original IP header is safely encrypted in the data section of this transport packet.

Layer 3 implementations should use Quality of Service (QoS) features end-to-end to ensure priority delivery through the network, particularly with time-sensitive data.

Application Layer

Application layer (TCP/IP 4 or OSI 7) encryption is implemented on each host using encryption features built in to the user applications. SSL providing secure web transactions or e-mail applications using PGP are both examples of application-layer encryption. While providing the ultimate in end-to-end security, application layer encryption does require all users to share the same encryption tools, even if they work on different hardware and software platforms. And it also relies on each user be somewhat knowledgeable about the encryption tools.

Hashing

As explained earlier, a hash is a one-way encryption method used to authenticate that data received is exactly the same as the data transmitted. The following basics about hashing are important to understand.

- A hash algorithm's function is to provide data integrity verification, not necessarily character-for-character translations of data.
- The hash process typically generates a fixed length number value derived from a variable length string of text.
- The hash output is often substantially smaller than the source text itself.
- The hash algorithm is sufficiently complex to make producing the same hash value with any other text combination extremely unlikely.

Hashed Message Authentication Codes (HMAC)

IPSec supports improved versions of Message Digest 5 (MD5) and Secure Hash Algorithm 1 (SHA-1) hash functions for cryptographically secure authentication. As with everything else in the security arena, hashing algorithms have evolved into HMACs, which build on the proven security of hashing algorithms by implementing additional cryptographic processes.

The HMAC algorithm involves having the original message text encrypted (hashed) using the sender's private key, resulting in a keyed, fixed-length checksum as output. The checksum value is appended to the message as a signature. The destination peer runs the same HMAC on the received message data, using the same private key. The resultant hash is compared with the received hash value, which should match exactly. Figure 9-16 shows this combination of source message and encryption key.

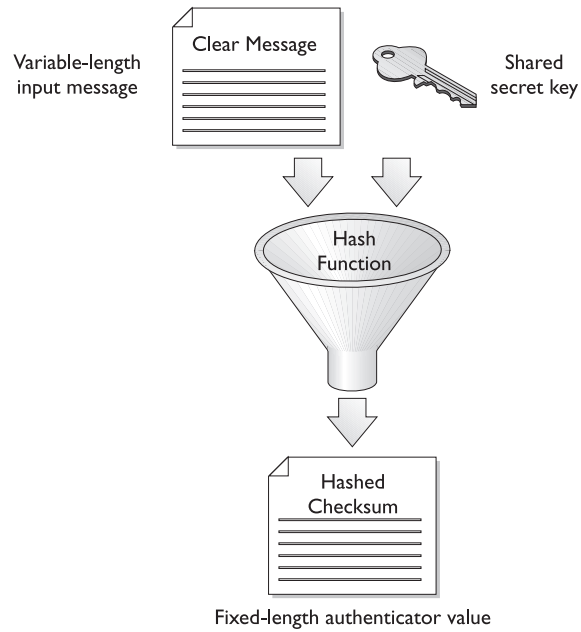

Figure 9-16 HMAC hash process showing the private encryption key

HMAC-MD5-96 Also known as HMAC-MD5 (RFC 1321), this hashing technique is based on MD5 developed by Ronald Rivest of MIT and RSA Data Security Incorporated. HMAC-MD5 uses a 128-bit secret key to produce a 128-bit authenticator value. Once produced, the 128-bit value is truncated to only the first 96 bits and stored in the AH or ESP authenticator field of the new packet. At the destination peer, the 128-bit authenticator value is computed and the first 96 bits are compared to the value stored in the authenticator field.

While the MD5 algorithm used alone was found vulnerable to collision search attacks and other known weaknesses, the vulnerability doesn't impact HMAC-MD5.

HMAC-SHA-1-96 Also known as HMAC-SHA-1 (RFC 2404), this hashing technique is based on SHA-1 specified in FIPS-190-1 combined with HMAC technology. HMAC-SHA-1 uses a 160-bit secret key to produce a 160-bit authenticator value. The 160-bit authenticator value is truncated to the first 96 bits and stored in the AH or ESP authenticator field of the new packet. At the destination peer, the 160-bit authenticator value is computed and the first 96 bits are compared to the value stored in the authenticator field.

Because of the longer key value, HMAC-SHA-1 is cryptographically stronger than HMAC-MD5, but it requires more CPU cycles to compute. Therefore, HMAC-SHA-1 should be the choice if slightly superior security is required. HMAC-MD5 should be the choice if slightly superior performance is required.

Diffie-Hellman Key Agreement (DH)

Diffie-Hellman (DH) is a public key encryption method that provides a mechanism for two IPSec peers to establish a shared secret key that only they know, while communicating over an insecure channel. DH is critical to IPSec operations because the shared secret key is used to encrypt data using the secret key encryption algorithms specified in the IPSec Security Associations, such as DES or MD5. Security Associations are covered later in this section.

DH algorithms and calculations are quite complex, using exponentiation of large numbers, between 768 and 1,536 bytes. Even simple calculations can take up to four seconds on a Cisco 2500. Fortunately, DH boils down to the following relatively simple conceptual steps:

- Each IPSec peer generates a public and private key pair. The public key is a mathematical derivative of the private key, conceptually like a series of unique office door keys might all be derived from a master key that can open every office.
- Each peer keeps its own private key secret and never shares it with anyone.
- Each peer sends its public key over the insecure channel.
- Each peer combines the other peer's public key with its own private key and computes the shared secret number that will be the same on each peer.

- The shared secret number is converted by each peer into a shared secret key. This shared secret key is never exchanged between the peers, which avoids the chance of it being captured in transit on the insecure link.

The end result of the DH process is this: each IPSec peer has the three following keys for each IPSec peer relationship it maintains.

- A *private key*, which is kept secret and never shared. The private key is used to "sign" outgoing messages to the IPSec destination peer.
- A *public key*, which is shared with the IPSec destination peer. The public key is used by the destination peer to verify a "signature" on a transmission.
- A *shared secret key*, which is used to encrypt data using any shared secret encryption algorithm, such as DES, 3DES, MD5, SHA-1, and so forth. This shared secret key is derived from the DH key generation process just described.

Defining a DH Group

The only real configuration in defining the DH process is to assign a DH group level. The groups supported might vary based on the device and the operating system (OS) release. Cisco IOS devices currently support DH Groups 1, 2, and 5. Group 5 provides the greatest security and is the most recently added, in IOS v12.1(3)T. The default group on most devices is Group 1.

DH Group	Description
1	The 768-bit Diffie-Hellman group
2	The 1,024-bit Diffie-Hellman group
5	The 1,536-bit Diffie-Hellman group

Security Association (SA)

The concept of Security Associations (SAs) is fundamental to understanding and configuring IPSec. An *SA* is a relationship between two or more potential VPN endpoints, which describes how those endpoints will use security services (technologies and protocols) to communicate securely. In establishing each secure communication connection, IPSec can provide services for encryption, integrity, and/or authenticity services. Once the services are selected, the two IPSec peers must determine exactly which algorithms to use for each service, such as DES or 3DES for encryption and MD5 or SHA for data integrity.

Once the services are selected and the algorithms chosen to implement those services, the two peers must exchange or implement session keys required by the algorithms. Is this beginning to sound complicated? How can you keep track of all these choices and decisions? The security association is the mechanism IPSec uses to manage these decisions and choices for each IPSec communication session. A basic component of configuring IPSec services on a client, router, firewall, or VPN concentrator is defining SA parameters.

IKE SAs versus IPSec SAs

The next section shows you that two types of SAs are used in configuring IPSec, just as there are two stages in establishing IPSec. *IKE SAs* describe the security parameters between two IKE devices, the first stage in establishing IPSec. *IPSec SAs* pertain to the actual IPSec tunnel, the second stage.

At the IKE level, a single IKE SA is established to handle secure communications both ways between the two peers. The following is an example of the type of information that would be included in an IKE SA.

Description	Example
Authentication method used	MD5
Encryption and hash algorithm	3DES
DH group used	2
Lifetime of the IKE SA in seconds or kilobytes	86,400
Shared secret key values for the encryption algorithms	Preshared

At the IPSec level, SAs are unidirectional—one for each direction. A separate IPSec SA is established for each direction of a communication session. Each IPSec peer is configured with one or more SAs, defining the security policy parameters to use during an IPSec session. To establish an IPSec session, peer 1 sends peer 2 a policy. If peer 2 can accept this policy, it sends the policy back to peer 1. This establishes the two one-way SAs between the peers.

IPSec Security Association (SA)

Each IPSec SA consists of security parameter values, such as a destination address, a unique security parameter index (SPI), the IPSec transforms used, the security keys, and additional attributes, such as IPSec lifetime. The SPI value becomes a unique record identifier (key field) linked to the SA parameters in the Security Parameter Databases in the RAM of peer devices.

Each IPSec SA consists of values such as

Description	Example
Peer (destination) address	10.1.1.23
Security parameter index (SPI)	7C123A9C
IPSec transforms used	AH, HMAC-SHA-1
Security keys	12345CD8765EF432A
Additional attributes (such as IPSec lifetime)	43,200

Five Steps of IPSec Revisited

This section discusses the individual steps required for a successful IPSec data exchange in greater detail. While IPSec incorporates many component technologies and offers multiple encryption options, the basic operation can be broken down into the following five main steps. Figure 9-17 show a graphical representation of the IPSec process.

Step 1	Interesting traffic initiates an IPSec session—Access lists interpret IPSec security policy to determine which traffic will be protected by IPSec.
Step 2	IKE Phase One—IKE authenticates peers and negotiates IKE SAs to determine if a secure channel can be established between the peers.
Step 3	IKE Phase Two—IKE negotiates the stricter IPSec SA parameters between the peers.
Step 4	IPSec data transfer—Qualifying data is transferred between IPSec peers.
Step 5	IPSec tunnel termination—IPSec session terminates through deletion or by timing out.

Step 1—Determine Interesting Traffic

Data communications covers a wide gamut of topics, sensitivity, and security requirements. Just as all rumors you hear aren't worth repeating, often much network traffic isn't worth securing, such as an employee's personal web browsing. Any security measures can come at the expense of network resources and performance, so a method must exist to determine what traffic is protected and what traffic travels normally through the public network.

The *VPN security policy* generally determines which type of traffic is considered "interesting" and, therefore, protected. How the VPN security policy gets implemented depends on the device platform. Cisco routers and the PIX Firewall use access lists to define the traffic to secure. The access lists are then incorporated in a crypto policy, which causes traffic associated with permit statements to be encrypted, while traffic associated with deny statements is sent unencrypted.

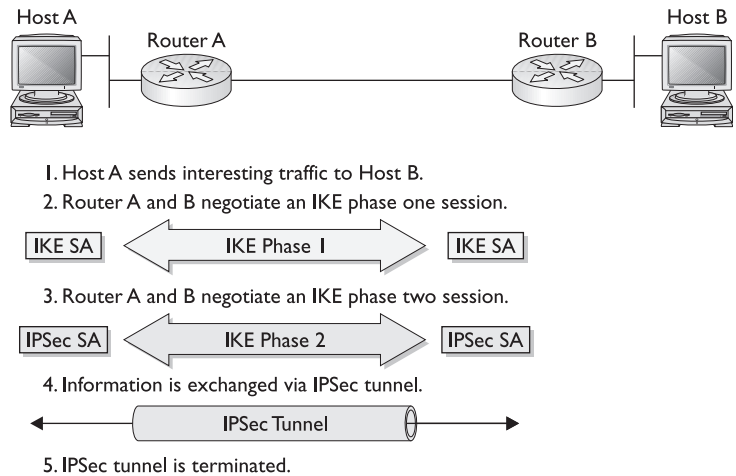

Figure 9-17 IPSec session five steps

Cisco VPN Client and some VPN management software tools allow the use of menu windows to define connections to be secured by IPSec.

Both methods of identifying interesting traffic are covered in Chapters 10 through 16.

Step 2—IKE Phase One

IKE is a key management protocol standard used in conjunction with IPSec. While IPSec can be configured without IKE, the use of IKE enhances the IPSec with additional features and makes it scalable. IKE authenticates each peer in an IPSec session, automatically negotiates two levels of SAs, and handles the exchange of session keys. This is accomplished in two phases: Phase One and Phase Two.

NOTE IKE was formerly known as the Internet Security Association Key Management Protocol/Oakley (ISAKMP/Oakley) or ISAKMP. In Cisco documents, it's typically referred to as IKE, but configuration commands still use the ISAKMP keyword. Just recognize that the two mean the same thing, but might not be interchangeable.

IKE Phase One's main purpose is to authenticate the IPSec peers and to set up a secure channel between the peers. To accomplish this, IKE Phase One performs the following functions:

- Authenticates while protecting the identities of the IPSec peers
- Negotiates a common IKE SA policy between peers to protect the IKE exchange
- Performs a DH exchange to create shared secret keys
- Sets up a secure tunnel to negotiate IKE Phase Two parameters

IKE Phase One occurs in two modes: Main mode and Aggressive mode.

Main Mode

Main mode has three two-way exchanges between the peers to create the secure connection and develop the common SAs, while protecting the identities of the IPSec peers.

- First exchange The security algorithms and hash methods to be used to secure the IKE exchanges are agreed on to create the common IKE SA for each peer.
- Second exchange A DH exchange is performed to generate shared secret keying material to be used by each peer to generate shared secret keys. *Nonces*, pseudorandom numbers, are sent to the other peer, signed, and returned to prove their identity.
- Third exchange The peer's identity is verified using the peer's IP address or fully qualified domain name (FQDN), such as ian@testco.com, in encrypted form.

The resulting IKE SA in each peer is bidirectional and specifies IKE exchange choices for the authentication method, encryption and hash algorithms, DH group, the lifetime of the IKE SA in seconds or kilobytes, and the shared secret key values for the encryption algorithms.

Aggressive Mode

As the name implies, in the *Aggressive* mode, only a single three-way exchange is performed. In the initial exchange, the sender defines the proposed IKE SA values and adds their DH public key, a nonce to be signed by the other party, and an identity packet to verify their identity via a third party.

The peer sends everything back that's required to complete the exchange, plus their DH public key. The only thing left is for the initiator to confirm the exchange.

While Aggressive mode is faster than Main mode, it exposes the peers to discovery because both sides have exchanged information before you have a secure channel. The possibility exists to sniff the wire and discover who formed the new SA.

Peer Authentication

Would-be IPSec peers must authenticate themselves to each other before IKE can proceed. IKE Phase One has three methods to authenticate IPSec peers in Cisco products. The two peers must negotiate a common authentication protocol from the following choices:

- Preshared keys—A key value entered into each peer manually (out of band) and used to authenticate the peer.
- RSA signatures—Uses a digital certificate authenticated by an RSA signature.
- RSA encrypted nonces—Uses RSA encryption to encrypt a nonce value (a random number generated by the peer) and other values.

A common value used by all authentication methods to help identify the peer is the peer identity (ID). Some ID values include the peer's IP address or their FQDN, such as ian.testco.com.

Preshared Key Authentication With this method, the same preshared key is configured on each IPSec peer. These IKE peers can authenticate each other by generating a hash of their ID, plus the key appended to the ID before transmission. If the receiving peer can create the same hash value using its preshared key, it authenticates the source peer. Figure 9-18 shows the preshared key authentication exchange.

While preshared keys are easier to configure than configuring the IPSec policy values on each IPSec peer, they don't scale well and potential security risks can be associated with getting the key to the peer.

RSA Signature Authentication RSA Signature Authentication is a public-key cryptosystem supported by IPSec for IKE Phase One authentication. This technology was developed in 1977 by Ron Rivest, Adi Shamir, and Leonard Adleman. RSA is the first letter of each developer's last name.

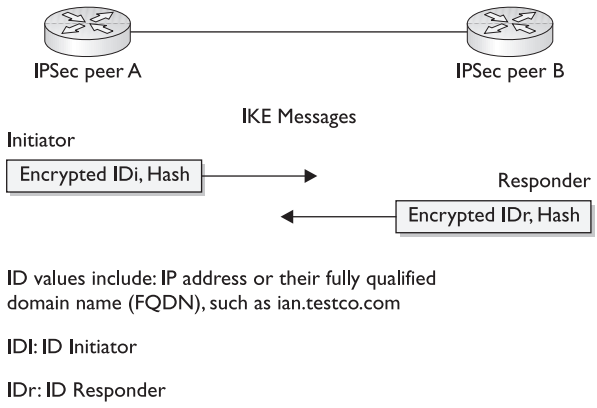

Figure 9-18 Preshared key authentication exchange

RSA signatures use a CA to generate a unique digital-identity certificate for each peer for authentication. The digital-identity certificate is similar in function to the pre-shared key, but provides stronger security.

To authenticate for an IKE session, the initiator and responder send each other their ID values, identity digital certificates, and an RSA signature value consisting of various IKE values, all encrypted using the negotiated IKE encryption method (DES or 3DES). Figure 9-19 shows an RSA signature authentication exchange.

RSA Encryption Authentication The RSA-encrypted nonces authentication method uses the RSA encryption public key cryptography algorithm. This technology requires that each party generate a pseudorandom number (a nonce) and encrypt it (and possibly other publicly and privately available information), using the other party's RSA public key. Authentication occurs when each party decrypts the other party's nonce with its local private key, and then uses the decrypted nonce to compute a keyed hash.

A drawback to this method is this: either side of the exchange can plausibly deny that it took part in the exchange. Cisco IOS software is the only Cisco product that supports this authentication method. Figure 9-20 shows an RSA encryption authentication exchange.

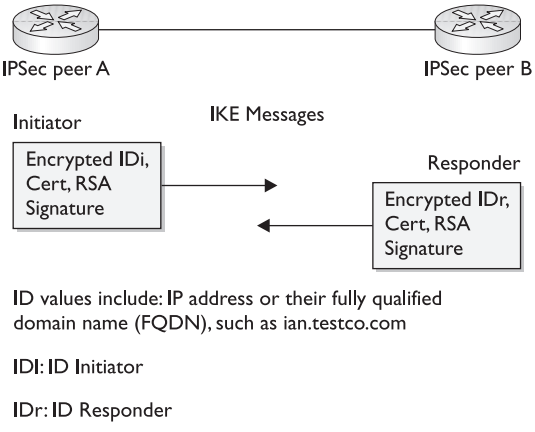

Figure 9-19 RSA signature authentication exchange

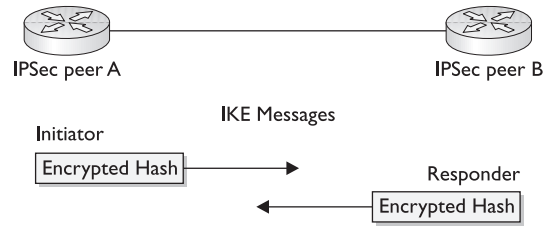

Figure 9-20 RSA encryption authentication exchange

CAs and Digital Certificates CAs and Digital Certificates are covered in greater detail in Chapter 11, but for our purposes here, they represent a digital identification system whereby an independent third party vouches for them. Conceptually, this is similar to most state driver's licenses when used as identification. The merchant accepting the ID accepts that the state has performed adequate verification to be reasonably sure the identity is valid.

The distribution of security keys through an untrusted network, such as the Internet, involves questionable levels of risk.

CAs are trusted third-party organizations, such as Verisign, Entrust, and Netscape, which provide digital certificates. The peers wanting to use digital certificates register with a CA. Once the CA verifies the client's credentials, a digital certificate is issued. The *digital certificate* contains the certificate bearer's identity (name or IP address), the certificate serial number, the certificate expiration date, and a copy of the certificate bearer's public key.

The digital certificate standard format is defined in the X.509 specification. Cisco supports X.509 version 3.

Step 3—IKE Phase Two

IKE Phase Two has only one mode, *Quick mode*, which occurs after IKE has established the secure tunnel in Phase One. In Quick mode, IKE

- Negotiates a shared IPSec policy
- Establishes IPSec SAs
- Derives shared secret keys used for the IPSec security algorithms

To negotiate the IPSec SAs, the sender forwards one or more transformsets, including transform combinations and related settings that represent the sender's security requirements/preferences for the new IPSec session. The receiving peer compares these requirements to its own transform sets (requirements/preferences). If one matches, then the recipient returns that single transform set, indicating a mutually agreed on transform and algorithms for the IPSec session.

Quick mode is also used to renegotiate a new IPSec SA any time the IPSec SA lifetime expires. If PFS (next section) isn't specified, Quick mode refreshes the key generation material used to create the shared secret keys derived from the DH exchange in Phase One.

Perfect Forward Secrecy (PFS)

When Perfect Forward Secrecy (PFS) is specified in the IPSec policy, a new DH exchange is performed for each Quick mode, providing new key material and, thereby, greater resistance to cryptographic attacks. Because each DH exchange requires large exponentiations, this option increases CPU use and can reduce performance.

Nonces

Quick mode exchanges nonces to provide replay protection. These nonces are used to generate new, shared, secret key material and to prevent replay attacks from generating bogus SAs.

Step 4—IPSec Data Transfer

Information is exchanged via the IPSec session based on the method for defining interesting traffic. Packets are encrypted and decrypted at the IPSec peers using any encryption specified in the IPSec SA.

Step 5—Session Termination

The IPSec session can be terminated because the traffic ended and the IPSec SA was deleted or the SA can time –out based on either SA lifetime setting. The SA timeout can be after a specified number of seconds or a specified number of bytes passed through the connection. The keys are discarded when SAs terminate, requiring IKE to perform a new Phase Two and, possibly, a new Phase One negotiation. New SAs can be established before the current ones expire, maintaining uninterrupted data flows.

IPSec Support in Cisco Systems Products

The choice of Cisco VPN technology depends on the type of VPN being developed, remote-access or site-to-site, plus the current and projected size of the resulting network. The Cisco components required to build the VPN might include any or all of the following options:

- Client VPN software, such as Easy VPN Remote, provides secure remote access to central Cisco routers, PIX Firewalls, and VPN Concentrators. The VPN client runs on the Windows OS.
- VPN hardware clients, such as Cisco VPN 3002, provide hardware-based encryption services for each remote user or small branch.
- Cisco VPN 3000 Series Concentrators provide powerful remote access and site-to-site VPN capabilities for the network with many VPN connections. The series supports an easy-to-use management interface.
- Cisco Secure PIX Firewall provides a highly secure VPN gateway alternative to the router or concentrator devices.

- IOS-based VPN-optimized routers for branch and remote user connections. Routers supporting this include the Cisco 800, UBr900 cable access router/modem, 1400 DSL router/modem, the 1700, 2600, 3600, 7100, 7200, and 7500 series routers. The Cisco 7100 Series VPN Router is an integrated VPN router that provides solutions for VPN-centric environments.
- Cisco Secure Intrusion Detection System (CSIDS) and Cisco Secure Scanner can be used to monitor and audit the security of the VPN.
- Cisco Secure Policy Manager and Cisco Works 2000, two Cisco network management software applications, provide VPN-wide system management.

Because VPN technology is so new and no single standard evolved quickly, many companies developed turnkey solutions, which could be implemented and administered easily by the client or consultant services. Cisco developed several VPN solutions built around each of the three following technologies. You learn about each technology in the next eight chapters.

- **VPN 3000 Series Concentrator** Most models use a hardware-based VPN solution with advanced encryption and authentication. *VPN concentrators* are designed and built specifically to provide high availability, high performance, and scalability. The modular design of the 3015 to 3080 models includes Scalable Encryption Processing (SEP) modules that allow the organization to increase capacity and throughput. This scalability supports small businesses, with up to 100 remote-access users, to large organizations, with up to 10,000 simultaneous remote connections.
- **Cisco Secure PIX Firewall** The *PIX Firewall* is a hardware security device that combines features like NAT, proxy server, packet and stateful filtering, and VPN features into a single device. The PIX unit uses a proprietary OS that trades the capability to handle a large variety of protocols and services for extreme robustness and performance by focusing on IP security.
- **VPN-Optimized Router** Cisco routers running an IOS with the VPN feature set provide routing, security, scalability, and QoS for businesses from small office/home office (SOHO) access to central-site VPN aggregation for the large-scale enterprise.

Chapter Review

This chapter looked at how VPNs can be used to extend the corporate networks securely using public networks, such as the Internet. The two basic VPN types are remote access and site-to-site. The three types of VPN connectivity are access VPN, intranet VPN, and extranet VPN. The two VPN modes are transport and tunnel.

While a variety of Layer 2 and Layer 7 VPN implementations exist, IPSec and IETF Layer 3 standards seem to dominate the market today. IPSec technologies include a variety of authentication and encryption methods.

Questions

1. Which is *not* one of the three basic types of VPN connections?
   - **A.** Access VPNs
   - **B.** Intranet VPNs
   - **C.** Internet VPNs
   - **D.** Extranet VPNs
2. Which is *not* one of the concerns in using the Internet for conducting private communications?
   - **A.** Loss of privacy
   - **B.** High cost
   - **C.** Loss of data integrity
   - **D.** Identity spoofing
3. Which one of the following is a Layer 2 tunneling protocol supported by Microsoft and Cisco?
   - **A.** PPTP
   - **B.** L2F
   - **C.** L2TP
   - **D.** GRE
4. With which security protocol is the data *not* encrypted?
   - **A.** AH
   - **B.** EST
   - **C.** ESP
   - **D.** Diffie-Hellman
5. What is the size of the encryption key for DES?
   - **A.** 40 bit
   - **B.** 56 bit
   - **C.** 128 bit
   - **D.** 168 bit
6. Which one of the following is *not* an encryption algorithm?
   - **A.** DES
   - **B.** 3DES
   - **C.** ESP
   - **D.** AES

7. Which is the most secure hashing algorithm?
   - **A.** MD5
   - **B.** SHA-1
   - **C.** HMAC MD5
   - **D.** HMAC SHA-1
8. With which security mode is the original IP header encrypted?
   - **A.** AH Transport
   - **B.** AH Tunnel
   - **C.** ESP Transport
   - **D.** ESP Tunnel
9. Which is *not* a valid transform?
   - **A.** ah-md5-hmac
   - **B.** esp-rfc1829
   - **C.** ah-des
   - **D.** esp-sha-hmac
10. Transform sets can contain how many AH transforms?
    - **A.** 1
    - **B.** 2
    - **C.** 3
    - **D.** None
11. Which cryptography type is also called public key encryption?
    - **A.** Symmetric encryption
    - **B.** Asymmetric encryption
    - **C.** Hash function
    - **D.** Cipher text
12. Which Diffie-Hellman key exchange offers the most security?
    - **A.** 5
    - **B.** 2
    - **C.** 1
    - **D.** 0
13. In an IPSec session, what is the minimum number of SAs that will be created?
    - **A.** 1
    - **B.** 2

C. 3

D. 6

14. At what point are the IPSec peers authenticated?

A. IKE Phase One

B. IKE Phase Two

C. IKE Phase Three

D. Interesting Traffic

15. What is a nonce?

A. A large prime number

B. A random number

C. A pseudorandom number

D. A digital signature

Answers

1. C. Internet VPNs
2. B. High cost
3. C. L2TP
4. A. AH
5. B. 56 bit
6. C. ESP
7. D. HMAC SHA-1
8. D. ESP Tunnel
9. C. ah-des
10. A. 1
11. B. Asymmetric encryption
12. A. 5
13. C. 3—1 IKE and 1 in each direction for IPSec
14. A. IKE Phase One
15. C. pseudorandom number

CHAPTER 10

Cisco IOS IPSec for Preshared Keys

In this chapter, you will learn to:

- Configure IPSec encryption tasks
- Configure IPSec manually

Using Internet Key Exchange (IKE) with preshared keys for authentication of IP Security (IPSec) sessions is relatively easy to configure, but it doesn't scale well. The process for configuring IKE preshared keys in Cisco IOS software for Cisco routers consists of four major tasks. This chapter looks at each configuration task in detail.

Configure IPSec Encryption Tasks

The good news is only four tasks are required to configure IPSec for preshared keys. The bad news is each task has multiple tasks that can initially seem overwhelming. The four tasks Cisco uses, which you can expect on the exam, are as follows:

- Task 1 Prepare for IKE and IPSec
- Task 2 Configure IKE
- Task 3 Configure IPSec
- Task 4 Test and verify IPSec

Don't make this more complicated than necessary. Task 1 is nothing more than making sure you've tested the existing network and gathered the information you need for Tasks 2 and 3. Task 2 is configuring for IKE Phase 1, while Task 3 is configuring for IKE Phase 2. Finally, Task 4 is checking your work.

The following task list shows the four tasks broken down into their individual steps. The steps are numbered to include the task number, as well as to help keep them straight. These steps are repeated in the chapter summary with the key commands listed for each step.

Figure 10-1 shows the networks that provide an example scenario used throughout this chapter. The goal is to create a secure VPN tunnel between Rtr1 at the company main office, and Rtr2 at one of almost 100 branch offices in North America, Europe, and Africa.

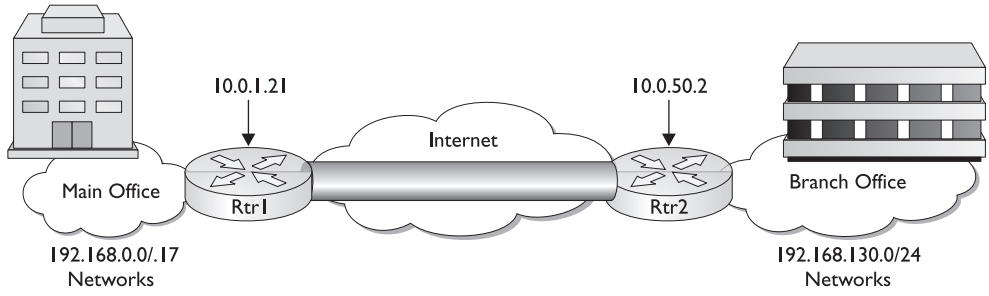

Figure 10-1 Chapter scenario VPN session to be configured

Task 1 Prepare for IKE and IPSec

- Step 1-1 Identify IPSec peers
- Step 1-2 Determine the IKE (IKE Phase 1) policies
- Step 1-3 Determine the IPSec (IKE Phase 2) policies
- Step 1-4 Check the current configuration
- Step 1-5 Ensure the network works without encryption
- Step 1-6 Ensure access control lists are compatible with IPSec

Task 2 Configure IKE

- Step 2-1 Enable or disable IKE
- Step 2-2 Create IKE policies
- Step 2-3 Configure preshared keys
- Step 2-4 Verify the IKE configuration

Task 3 Configure IPSec

- Step 3-1 Configure transform set suites
- Step 3-2 Configure global IPSec security association lifetimes
- Step 3-3 Configure crypto ACLs
- Step 3-4 Configure crypto maps
- Step 3-5 Apply the crypto maps to the interface

Task 4 Test and verify IPSec

- Step 4-1 Display the configured IKE policies
- Step 4-2 Display the configured transform sets
- Step 4-3 Display the current state of the IPSec SAs
- Step 4-4 Display the configured crypto maps
- Step 4-5 Debug IKE events
- Step 4-6 Debug IPSec events

The assumption is this: the main office has reserved networks 192.168.0.0 through 192.168.127.0 for itself and will use one class C for each branch in the remaining 192.168.128.0 to 192.168.255.0 addresses.

The example uses private addresses to avoid using public addresses that might belong to others and to make it easier for those who choose to try to create the configuration in a test lab.

Task 1 Prepare for IKE and IPSec

Successful implementation of an IPSec network requires testing of the existing network and advance planning before any configuration begins. Insufficient testing and planning can lead to troubleshooting problems or configuration errors. Some preparation and planning steps include the following:

- Step 1-1 Identify IPSec peers
- Step 1-2 Determine the IKE (IKE Phase 1) policies
- Step 1-3 Determine the IPSec (IKE Phase 2) policies
- Step 1-4 Check the current configuration
- Step 1-5 Ensure the network works without encryption
- Step 1-6 Ensure Access Control Lists are compatible with IPSec

The actual order of these steps could be varied based on personal preferences and implementation circumstances.

Step 1-1 Identify IPSec Peers

An important part of defining a comprehensive IPSec policy is to identify the IPSec peer pairs that must be configured. In the chapter scenario, expanded in Figure 10-2, each remote site will connect only to the Main Office router and, therefore, requires only simple configuration. The Cisco router at the Main Office must be configured for peer

communications with each of the remote sites and telecommuter(s). Each peer must support IPSec. Because many different types of peer devices exist, it's important to identify all potential peers and determine their VPN capabilities. Possible peer devices could include, but aren't limited to, the following:

- Cisco routers
- Cisco Secure PIX Firewalls
- Cisco Secure VPN 3000 Concentrators
- Cisco Secure VPN 3002 Hardware Client device
- Cisco Secure VPN Software Client
- Other vendor IPSec products that conform to IPSec standards
- CA servers

It's important to recognize that IPSec features supported and default settings can vary between Cisco product families, as well as versions of the operating system (OS) being used. This is most important for the Main Office router in the scenario because it must be able to establish common IKE and IPSec policies with each remote device. This also demonstrates why many companies limit the number of devices supported by defining standards for telecommuters or branch offices.

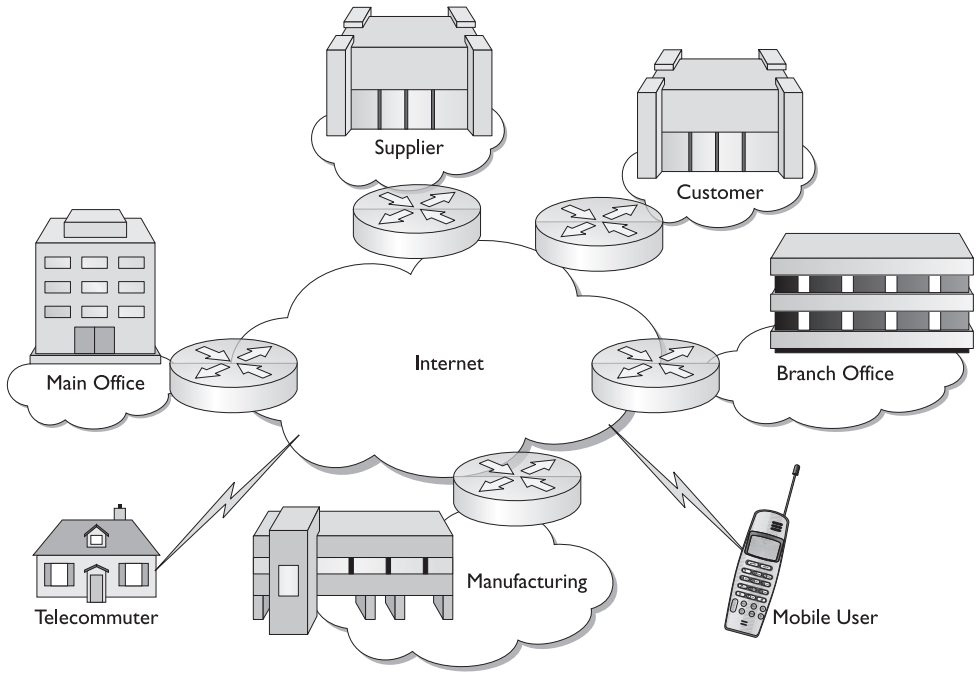

Figure 10-2 Chapter scenario network showing peer connections

The result of this analysis might be a table like the following:

Location	Device	Version (or OS)
Main Office	Cisco 7100 router	12.2(8)T with FW and VPN
Telecommuters (24)	Cisco 900 Cable/DSL router	12.2(8)T with FW and VPN supports 3DES
Mobile users (16)	Cisco VPN Software Client	v3.6-3DES
Branch offices (80) North America	Cisco 2600/3600 routers	12.2(8)T with FW and VPN supports 3DES
Branch offices (20) Europe/Africa	Cisco 2600/3600 routers	12.2(8)T with FW and VPN supports DES only
Manufacturing (1)	Cisco PIX 525 Firewall	OS v6.2-3DES

Step 1-2 Determine the IKE (IKE Phase 1) Policies

IKE is a hybrid protocol that implements the Oakley key exchange and the Skeme key exchange inside the Internet Security Association and Key Management Protocol (ISAKMP) framework. (ISAKMP, Oakley, and Skeme are security protocols implemented by IKE.)

The purpose of IKE Phase 1 is to authenticate the peers and negotiate a secure session between those peers. This creates a stable platform for IKE Phase 2 to negotiate the IPSec tunnel. The word "negotiate" makes this sound like a complex process, which might involve contentious arguments late into the night. In reality, at each phase, the session instigator submits one or more security policies, which will then be compared by the destination peer against its preconfigured security policies. If a match is found, a session is then established. If no match is found, the session is terminated. From a practical standpoint, if a VPN session can't be established between the two peers, then one of the devices must be configured to include a set of policies acceptable to the other.

The objective here is to develop one or more IKE security policies based on the overall company security policy. Each policy will require decisions about five security options: authentication method, encryption algorithm, hashing algorithm, Diffie-Hellman group, and SA lifetime. This can make IKE seem more complicated than it is. To get some perspective, we aren't configuring a connection to an unknown or an untrusted entity. Instead, we're connecting a branch of the company to the main office system. Chances are both are governed by the same security policy so, while some choices must be made, the list of possibilities shouldn't be infinite.

In many cases, the decision could be that 3DES will be the encryption, MD5 will perform hashing functions, preshared keys will be used for authentication, and so forth, and that would be configured on each router. So, why have multiple policies with different option combinations? The answer depends on the router and the number of VPN connections that it needs to support. In our example, if each branch only connects to the main office, then they only need the policy that matches the main office router.

The main office router could need a couple of IKE policy choices because all branch offices might be unable to support the same policy. This might not be a vendor or a product platform issue, but a legal issue. While all the North American branches can be configured with 3DES encryption, export controls could prevent the foreign branches

from using anything but DES. So, a second policy needs to be available for those connections.

Research the Parameter Options The following table shows the five security features that must be decided, including the current options, defaults, and which choice provides the greatest security. A good idea is always to confirm these options and defaults as IOS versions change. It only makes sense that the IOS will assimilate new IPSec standards as they're established.

Parameter	Choices	Keyword	Default	Strongest
Message encryption algorithm	DES 3DES	des 3des	des	3des
Message integrity - hash algorithm	SHA-1 (HMAC variant) MD5 (HMAC variant)	sha md5	sha	sha
Peer authentication method	RSA signatures RSA encrypted nonces preshared keys	rsa-sig rsa-encr pre-share	rsa-sig	rsa-sig
Diffie-Hellman group (key exchange)*	768 bit 1,024 bit	1 2	1	2
IKE SA lifetime	60 to 86,400 seconds		86,400	60

*Some VPN devices support D-H group 5 (1,568 bit), but it might be a while before IOS devices support this level of key exchange.

Develop the Parameter Preferences To complete the IKE planning process, what would make sense is to create a table of the preferred combination of security features, plus one or more fallback options for those devices or locations that can't support the preferred package. The resulting table might look like the following:

Parameter	Preferred (stronger)	2nd Choice	3rd Choice
Encryption algorithm	3des	des	des
Hash algorithm	sha	sha	md5
Authentication method	preshare	preshare	preshare
DH key exchange group	2	2	1
IKE SA lifetime	43,200	43,200	86,400

In comparing the two tables, you should see that RSA-SIG would be a preferred authentication method over preshared keys. This chapter deals with using preshared keys, while Chapter 12 covers Certificate Authority (RSA-SIG). RSA encrypted nonces are addressed in the last section of this chapter.

NOTE You can't use SHA together with DES encryption on Cisco's VPN software client version 3.6. Part of the problem with determining a set of preferences is being aware of the different platform limitations.

Using our scenario for this chapter, the Preferred column represents our configuration preferences for all North American branches and would be the only options configured. The 2nd Choice column would be for those overseas branches that can support the Preferred configuration, except for the export restriction on triple DES. Assuming they don't form VPN sessions with anyone else, this would be the only choice configured on those devices. The 3rd Choice column might represent extremely small branches or telecommuters using small personal routers with limited resources and options. This also represents the lowest level of parameter options the main office will accept for a VPN connection. The only device that will have all three choices configured is the main office router.

Because any parameter that isn't defined will use the existing default value, the actual table could look like the following example. Either table can be used in Task 2 to configure the IKE parameters.

Parameter	Preferred (stronger)	2nd Choice	3rd Choice
Encryption algorithm	3des		
Hash algorithm			md5
Authentication method	preshare	preshare	preshare
DH key exchange group	2	2	
IKE SA lifetime	43,200	43,200	

Step 1-3 Determine the IPSec (IKE Phase 2) Policies

Once the choices are made for IKE Phase 1, it's time to turn to those parameters required to complete IKE Phase 1. This is where the IPSec tunnel is negotiated and, ultimately, the IPSec SAs established. As in Phase 1, the goal is to define one or more sets of IPSec security parameters defining the IPSec security policy, based on the overall company security policy.

The information gathered is required in Task 3 when IPSec is configured. To facilitate that process, gather the following information into a table similar to the one developed for Phase 1 for each peer for which sessions will be established.

Transforms (or transform sets)	IPSec transforms or build transform sets, plus the VPN mode to provide the optimal security and performance balance
Peer address	Peer IP address for this session
Peer host name	What is the peer host name for this session?
IP address to protect	Which source hosts are to be protected for this session?
Traffic to protect	What applications should be protected for this session?
SA establishment method	Will SAs be manually established or established via IKE?

Most of the choices to be configured, such as peer address and host name, are straight-forward. The transforms or, if necessary, the transform sets, involve determining the security features required, and then striking a balance between security level and performance implications. Chapter 9 introduced Transforms and Task 3, later in this chapter, looks at configuring them.

The following table might represent the IPSec values for the chapter scenario:

Parameter	Site 1	Site 2
Transform (set)	esp-des, tunnel	esp-des, tunnel
Peer host name	Rtr2	Rtr1
Peer IP address	10.0.50.2	10.0.1.21
Hosts to be encrypted	192.168.0.0/25	192.168.30.0/24
Traffic to be encrypted	TCP	TCP
SA establishment	ipsec-isakmp	ipsec-isakmp

Step 1-4 Check the Current Configuration

It's important to check the current Cisco router configuration to see if any existing IPSec policies are configured that could be useful for, or interfere with, the new IPSec policies. If appropriate, previously configured IKE and IPSec policies can be used to save configuration time. This section looks at three commands useful in discovering existing IKE and IPSec policies.

The show running-config Command The basic **show running-config** command is always a good starting point. In the following partial output, the lines in boldface are the IKE and IPSec configuration statements.

```
Rtr2#show run
!
hostname Rtr2
!
enable secret 5 $1$ojfF$eAp7xa0hU4E678XaibAQa
!
crypto isakmp policy 1
hash md5
authentication pre-share
crypto isakmp key test_key1 address 10.0.1.21
!
crypto ipsec transform-set trans1 esp-3des
crypto ipsec transform-set trans2 ah-md5-hmac esp-des
crypto ipsec transform-set trans3 esp-sha-hmac esp-null
!
crypto map test_map1 1 ipsec-isakmp
set peer 10.0.1.21
set security-association lifetime seconds 43200
set transform-set trans1 trans2 trans3
match address 101
!
interface Ethernet0
 ip address 192.168.130.1 255.255.255.0
 no mop enabled
!
interface Serial0
 ip address 10.0.50.2 255.255.255.252
 no ip mroute-cache
 no fair-queue
 crypto map test_map1
```

```
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.0.50.1
access-list 101 permit tcp 192.168.130.0 0.0.0.255 192.168.0.0 0.0.127.255
dialer-list 1 protocol ip permit
!
```

The show crypto isakmp policy Command The following is an example of the output from the **show crypto isakmp policy** command, which can be used to view all existing IKE policies. Notice the default parameters are displays at the bottom of the output.

```
Rtr1#show crypto isakmp policy
Protection suite of priority 100
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #2 (1024 bit)
        lifetime:               43200 seconds, no volume limit
Default protection suite
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Rivest-Shamir-Adleman Signature
        Diffie-Hellman group:   #1 (768 bit)
        lifetime:               86400 seconds, no volume limit
Rtr1#
```

The show crypto map Command The **show crypto map** command is useful for displaying any previously configured crypto maps. If possible, these can be used to save configuration time, but they can interfere with the new IPSec policy.

```
Rtr1#show crypto map
Crypto Map "testmap" 50 ipsec-isakmp
        Description: VPN Link to branch office in Tacoma, WA
        Peer = 10.0.50.2
        Extended IP access list 125
            access-list 125 permit tcp 192.168.0.0 0.0.127.255 192.168.130.0
  0.0.0.255
        Current peer: 10.0.50.2
        Security association lifetime: 2300000 kilobytes/1800 seconds
        PFS (Y/N): Y
        DH group:  group2
        Transform sets={ CPU-HOG, }
        Interfaces using crypto map testmap:
                    Serial0
Rtr1#
```

The show crypto ipsec transform-set Command The **show crypto ipsec transform-set** command can be used to view previously configured transform sets. Whenever possible, these transforms can, and should, be used to save configuration time.

```
Rtr1# show crypto ipsec transform-set
Transform set MD5-DES: { esp-des esp-md5-hmac  }
   will negotiate = { Tunnel,  },

Transform set DES-ONLY: { esp-des  }
   will negotiate = { Tunnel,  },
```

```
 Transform set CPU-HOG: { ah-md5-hmac  }
   will negotiate = { Tunnel,  },
   { esp-des esp-md5-hmac  }
   will negotiate = { Tunnel,  },

Transform set AH-ONLY: { ah-sha-hmac  }
   will negotiate = { Tunnel,  },
Rtr1#
```

Step 1-5 Ensure the Network Works Without Encryption

All peer-to-peer connectivity must be verified before configuring IPSec encryption. Basic troubleshooting techniques become more difficult, if not impossible, once encryption is in place.

While the router **ping** command can be used to verify basic connectivity between peers, equally important is that all important applications and related data are confirmed to be accessible. This can only be accomplished by testing those applications before beginning IPSec configuration.

Skipping this step can lead to wasted time trying to debug IPSec configuration only to discover eventually that the underlying connectivity was never tested. A simple rule of thumb is this type of mistake is always exposed when it can embarrass you the most.

Step 1-6 Ensure Access Control Lists Are Compatible with IPSec

Make certain any existing access lists on VPN device and perimeter router don't block IPSec traffic. Perimeter routers frequently implement restrictive security policies using ACLs. These policies often deny all inbound traffic that isn't responding to a specific outbound request. This restriction often blocks inbound IPSec traffic necessary to establish a VPN session.

Specific permit statements must be added to the inbound ACL to allow IPSec traffic. The following ports and protocol access should be left open:

- **IKE/ISAKMP** UDP port 500/keyword: isakmp
- **Encapsulating Security Payload (ESP)** IP protocol number 50/keyword: esp
- **Authentication Header (AH)** IP protocol number 51/keyword: ahp

The first step is to use the **show access-lists** command to determine if any ACLs will block IPSec traffic.

Assuming an inbound ACL is causing problems, it must be edited, placing entries at the beginning to permit IPSec traffic. The following shows an example of the lines that should be added to Rtr2's existing ACL 110.

```
Rtr2# show running-config
!
interface Serial 0
 ip address 10.0.50.2 255.255.255.252
 ip access-group 110 in
!
access-list 110 permit ahp host 10.0.1.21 host 10.0.50.2
access-list 110 permit esp host 10.0.1.21 host 10.0.50.2
```

```
access-list 110 permit udp host 10.0.1.21 host 10.0.50.2 eq isakmp
    ---balance of ACL statements---
```

Task 2 Configure IKE

The second major task in configuring the IPSec VPN is to configure the IKE parameters gathered in Task 1, Step 2. Configuring IKE involves the following four steps:

- Step 2-1 Enable or disable IKE
- Step 2-2 Create IKE policies
- Step 2-3 Configure preshared keys
- Step 2-4 Verify the IKE configuration

This section looks at these four steps and the associated commands.

Step 2-1 Enable or Disable IKE

The first step in configuring IKE is to make sure IKE is globally enabled. While IKE is enabled globally by default, the **crypto isakmp enable** command makes sure the feature is on. Use the no form of the command to disable IKE. Because IKE is a global **mode** command, it needn't be enabled or disabled for individual interfaces. Some circumstances might warrant using an ACL entry to block ISAKMP access (UDP port 500) on interfaces not being used for IPSec to reduce the chances of a denial of service (DoS) attack. The command syntax and an example follow.

```
Rtr2(config)# crypto isakmp enable
Rtr2(config)# no crypto isakmp enable
```

Step 2-2 Create IKE Policies

Now, it's time to define a suite of ISAKMP (IKE) policy details gathered during the planning task, Task 1, Step 2 (repeated in the following with defaults in italics). The goal is to create a suite of IKE policies, so you can establish IKE peering between two IPSec endpoints.

Parameter	Preferred (stronger)	2nd Choice	3rd Choice
Encryption algorithm	3des	*des*	*des*
Hash algorithm	*sha*	*sha*	md5
Authentication method	preshare	preshare	preshare
DH key exchange group	2	2	*1*
IKE SA lifetime	43,200	43,200	*86,400*

Using the scenario example, Rtr2 only needing to connect to Rtr1, might be configured with only the preferred (most secure) set of parameters because that set will also be on Rtr1. On the other hand, Rtr1, needing to peer with various devices both domestically and internationally, might need multiple policies to accommodate the different peers. The trick is to configure the policies, so the preferred policy is evaluated by the

peer first, and then the second and third choices. This ensures the agreed policy is always the highest mutually agreeable choice.

The crypto isakmp policy Command Use the global **crypto isakmp policy** command to define an IKE policy containing a set of parameters used during IKE negotiation. Use the no form of this command to delete an IKE policy. The command syntax and an example follow.

```
Rtr1(config)#[no] crypto isakmp policy priority
Rtr1(config)#crypto isakmp policy 100
Rtr1(config-isakmp)#
```

priority	Unique identifier for the IKE policy used to order, or prioritize, the policies. Acceptable range 1 to 10,000 with 1 being the highest possible priority.

This command invokes the ISAKMP policy configuration (config-isakmp) mode, which allows the six IKE parameters to be defined. The no form of each parameter removes that setting. Exit is used to end the session and any undefined parameters use the default values. The actual choices ware based on the table made in Task 1, Step 2. The following shows the syntax for each parameter, followed by the entries for the preferred policy.

```
Rtr1#conf t
Rtr1(config)#crypto isakmp policy 100
Rtr1(config-isakmp)#authentication {pre-share | rsa-encr | rsa-sig}
Rtr1(config-isakmp)#encryption {des | 3des}  (depending on feature set)
Rtr1(config-isakmp)#hash {md5 | sha}
Rtr1(config-isakmp)#group {1 | 2}
Rtr1(config-isakmp)#lifetime <60 to 86400 seconds>

Rtr1(config)#crypto isakmp policy 100
Rtr1(config-isakmp)#authentication pre-share
Rtr1(config-isakmp)#encryption 3des
Rtr1(config-isakmp)#hash sha
Rtr1(config-isakmp)#group 2
Rtr1(config-isakmp)#lifetime 43200
Rtr1(config-isakmp)#exit
```

Because the preferred policy was assigned 100 as a priority, the second and third choices must be progressively larger to control the order in which they're offered to a peer. A possible choice might be 200 and 300, respectively, enabling newer policies to be assigned priorities that would allow them to be prioritized according to changes in the company security policy. Using 1, 2, and 3 might initially be easy to understand, but it could force reconfiguring the devices if a new feature like AES encryption is implemented.

The following shows the steps to configure the second and third preference, taking advantage of default values to save configuration:

```
Rtr1#conf t
Rtr1(config)#crypto isakmp policy 200
Rtr1(config-isakmp)#authentication pre-share
Rtr1(config-isakmp)#group 2
```

```
Rtr1(config-isakmp)#lifetime 43200
Rtr1(config-isakmp)#crypto isakmp policy 300
Rtr1(config-isakmp)#authentication pre-share
Rtr1(config-isakmp)#hash md5
Rtr1(config-isakmp)#exit
```

Branch Office Configuration Implications Assuming the chapter scenario, if each branch will only be connecting VPNs to the main office, configuring only one of the IKE policies for each branch should be necessary. Assuming the domestic branches can all support 3DES, policy 100 would seem appropriate. For any foreign branches that can't use 3DES because of export restrictions, then policy 200 seems the likely candidate. If any branch initiates a session with the main office, both peers would have the same policy. The same is true for a session initiated by the main office to any branch; even the foreign branches could accept the second policy offered.

This whole discussion is based on using a single IKE policy between peers for any encrypted traffic. If the security policy called for 3DES for only certain types of traffic and DES for others, then configuring a second policy on the remote peers would be necessary.

The show crypto isakmp policy Command Use the **show crypto isakmp policy** command to verify the entries, display the default policy, and confirm that the nonconfigured parameters took on the default values.

```
Rtr1#show crypto isakmp policy
Protection suite of priority 100
        encryption algorithm:   3DES - Triple DES (168 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #2 (1024 bit)
        lifetime:               43200 seconds, no volume limit
Protection suite of priority 200
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #2 (1024 bit)
        lifetime:               43200 seconds, no volume limit
Protection suite of priority 300
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Message Digest 5
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #1 (768 bit)
        lifetime:               86400 seconds, no volume limit
Default protection suite
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Rivest-Shamir-Adleman Signature
        Diffie-Hellman group:   #1 (768 bit)
        lifetime:               86400 seconds, no volume limit
Rtr1#
```

NOTE Although the output shows "no volume limit" for the lifetimes, currently it is only possible to configure a time IKE lifetime (such as 86,400 seconds); volume limit lifetimes aren't used.

Step 2-3 Configure Preshared Keys

IPSec peers authenticate each other during IKE negotiations using the preshared key and their IKE identity. The *IKE identity* can either be the device IP address or the host name. The router IOS defaults to the IP address and doesn't add a line to the configuration, unless the host name option is used.

To use the host name identity method, use the global **crypto isakmp identity** command. The no form of this command resets the default value (address). The command syntax is as follows:

```
Rtr1(config)#[no] crypto isakmp identity {address | hostname}
Rtr1(config)#crypto isakmp identity hostname
```

Using the host name identity method makes the most sense if a DNS server is available on the network to resolve the name. If not, the following example shows using an **IP Host** command entry on the router to handle the name resolution:

```
Rtr1(config)#ip host Rtr2.domain.com 10.0.50.2
```

Use the global **crypto isakmp key** command to configure a preshared authentication key whenever you specify preshared keys in an ISAKMP policy. Use the no form of the command to delete a preshared authentication key. The command syntax for both the address and the host name option is as follows:

```
Rtr1(config)#[no] crypto isakmp key keystring address peer-address [mask]
Rtr1(config)#[no] crypto isakmp key keystring hostname peer-hostname
```

keystring	The preshared key. Any combination of alphanumeric characters up to 128 bytes. Must be identical on both peers.
mask	(Optional) Defines a subnet address for the remote peer. Indicates the remote peer ISAKMP identity will be established using the preshared key only. With the mask, keys are no longer restricted to only two users. Added v12.1(1)T.

A given preshared key is shared between two peer devices. While it's possible for Rtr1 in the scenario to use the same key with all remote peers, the more secure approach is to use different keys with each pair of peers.

The next example shows IKE and preshared key (cisco123) for Rtr1 and Rtr2 using the address identity method. The IKE policies are compatible. Because the hash algorithm SHA-1 is the default value, it needn't be configured.

```
Rtr1(config)#crypto isakmp key cisco123 address 10.0.50.2
Rtr1(config)#crypto isakmp policy 100
Rtr1(config-isakmp)#authentication pre-share
Rtr1(config-isakmp)#encryption 3des
Rtr1(config-isakmp)#group 2
Rtr1(config-isakmp)#lifetime 43200
Rtr1(config-isakmp)#exit
```

```
Rtr2(config)#crypto isakmp key cisco123 address 10.0.1.21
Rtr2(config)#crypto isakmp policy 100
Rtr2(config-isakmp)#authentication pre-share
Rtr2(config-isakmp)#encryption 3des
Rtr2(config-isakmp)#group 2
Rtr2(config-isakmp)#lifetime 43200
Rtr2(config-isakmp)#exit
```

Because both devices are under the administrative policies of a single entity, it makes sense that the policy priority value might be the same on both devices. That isn't a requirement, though, and it would be unlikely when the two peers are from different organizations. The important thing is the parameters are the same.

NOTE One exception to matching parameters is the lifetime can be shorter on the destination peer and that value will be used. The shorter of the two lifetimes is always used. A shorter lifetime, in theory, increases the security.

Step 2-4 Verify the IKE Configuration

Use the **show crypto isakmp policy** command to display configured and default policies. An example of IKE policy for Rtr1 is shown in the Step 2-2 section. The **show running-config** command can also be useful, but the parameters using default values are omitted for brevity.

Task 3 Configure IPSec

The next major task in configuring Cisco IOS IPSec is to configure the IPSec parameters gathered in Task 1, Step 3. The general steps used to configure IPSec encryption on Cisco routers are summarized as follows. This chapter looks at each configuration step in detail.

- Step 3-1 Configure transform set suites
- Step 3-2 Configure global IPSec security association lifetimes
- Step 3-3 Configure crypto ACLs
- Step 3-4 Configure crypto maps
- Step 3-5 Apply the crypto maps to the interface

Step 3-1 Configure Transform Set Suites

In Chapter 9, an IPSec transform was defined as a single IPSec security protocol—AH or ESP—with its associated security algorithms and mode. Technically, a *transform* is the list of operations done on a dataflow to provide data authentication, data confidentiality, and since IOS v12.2(8) data compression. Some of the supported transform choices are shown in Figure 10-3.

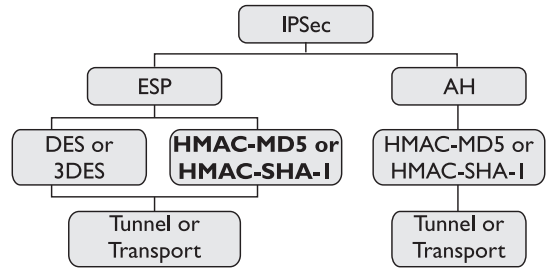

Figure 10-3 IPSec transform options

The actual transforms supported might vary by device type. VPN devices could support a slightly different group than the PIX Firewall. The following are transforms supported by the IOS-based devices:

AH Transforms—Pick up to one:

Transform	Description
ah-md5-hmac	AH with the MD5 (HMAC variant) authentication algorithm
ah-sha-hmac	AH with the SHA-1 (HMAC variant) authentication algorithm

ESP Encryption Transforms—Pick up to one:

Transform	Description
esp-des	ESP with the 56-bit DES encryption algorithm
esp-3des	ESP with the 168-bit DES (Triple DES) encryption algorithm
Esp-null	ESP without cipher. Can be used with esp-md5-hmac or esp-sha-hmac for ESP authentication without encryption. Shouldn't be used in a production network because of the lack of security

ESP Authentication Transform—pick up to one, but only if you also selected the esp-des or esp-3des transform (not esp-rfc1829):

Transform	Description
esp-md5-hmac	ESP with the MD5 (HMAC variant) authentication algorithm
esp-sha-hmac	ESP with the SHA-1 (HMAC variant) authentication algorithm

IP Compression Transform—pick up to one:

Transform	Description
comp-lzs	IP compression with the LZS algorithm. (IOS v12.2((8)) or later)

The compression algorithm used in the *comp-lzs* transform is Lempel-Ziv-Stac (LZS) compression, the most widely accepted lossless compression algorithm at this time. Compression is implemented on Layer 3, like hardware encryption, and can considerably reduce bandwidth requirements to support IP Security (IPSec).

Transform Sets A *transform set* is a combination of up to three individual IPSec transforms designed to implement a specific security policy to protect a particular dataflow. The transform sets represent the security choices available during IPSec SA negotiation between two IPSec peers in IKE Phase 2 quick mode. The following are two examples of transform sets:

- **ah-sha-hmac and esp-des and esp-sha-hmac** AH protocol with SHA-1 authentication, ESP DES encryption, plus ESP SHA-1 authentication.
- **esp-3des and esp-md5-hmac** ESP protocol with DES encryption, plus ESP MD5 authentication.

The IKE Phase 2 quick mode operates much like the Phase 1 negotiation. The initiator sends one or more transform sets to the destination peer. The destination peer compares them in order against its own preferred transform sets. If a match is found, the destination peer returns the transform to the initiator peer to complete the "negotiation." The peers must share a common transform set or the exchange can't occur. If the peers can't agree on a transform set, this means one or the other will need to configure a matching set of policies.

Transform sets are limited to three transforms with no more than one AH transform, plus no more than two ESP transforms. Transform sets combine the following IPSec features:

- Payload authentication, example: AH transform
- Payload encryption, example: ESP transform
- IPSec mode (transport or tunnel)

You can define multiple transform sets, and then specify one or more of these sets using a crypto map entry. Use the global **crypto ipsec transform-set** command to define a transform set. Use the no form of the command to delete a transform set. The command syntax is as follows:

Rtr1(config)#crypto ipsec transform-set *transform-set-name trans1* [*trans2* [*trans3*]]
Rtr1(config)#no crypto ipsec transform-set *transform-set-name*

The command invokes the Crypto-Transform Configuration mode that can be used to define up to four transform sets. Use the **exit** command to end the configuration. Once defined, the transform sets can be specified in a crypto map entry. The following example shows the creation of four transform sets:

```
Rtr1#conf t
Rtr1(config)#crypto ipsec transform-set MD5-DES esp-md5-hmac esp-des
Rtr1(cfg-crypto-trans)#crypto ipsec transform-set DES-ONLY esp-des
Rtr1(cfg-crypto-trans)#crypto ipsec transform-set AH-ONLY ah-sha-hmac
Rtr1(cfg-crypto-trans)#crypto ipsec transform-set CPU-HOG ah-md5-hmac
  esp-md5-hmac esp-des
Rtr1(cfg-crypto-trans)#exit
Rtr1(config)#
```

When IKE isn't being used to establish SAs, a single transform set must be configured and the transform set isn't negotiated.

Step 3-2 Configure Global IPSec Security Association Lifetimes

IPSec security associations that use the shared secret keys time out together, based on the SA lifetime configured. This section looks at configuring a global lifetime value to be used when a particular crypto map entry doesn't have lifetime values configured. When a router receives a negotiation request from a peer, it uses the smaller of the lifetime value proposed by the peer or the locally configured global lifetime value for the new security associations.

Two lifetimes exist—a "timed" lifetime and a "traffic-volume" lifetime that can be configured. The SA and keys expire when the first of these lifetimes expires.

Any changes made to a global lifetime are only applied when the crypto map entry doesn't have a lifetime value specified. The change isn't applied immediately to existing security associations, but will be used in subsequent SA negotiations.

Use the global **crypto ipsec security-association lifetime** command to set the global lifetime. For a timed lifetime, use the *seconds* form of the command. The *kilobytes* form causes the security association to time out after the specified traffic limit (in kilobytes) has been protected. The lifetime values are ignored for manually established security associations installed using an **ipsec-manual crypto map** command entry. To reset a lifetime to the default value, use the no form of the command. The command syntax is as follows:

```
Rtr1(config)#crypto ipsec security-association lifetime {seconds seconds |
kilobytes kilobytes}
Rtr1(config)#no crypto ipsec security-association lifetime {seconds | kilobytes}
```

The default value for a timed lifetime is 3,600 seconds (one hour) with the range of acceptable values between 120 and 86,400 (one day) seconds. The default value for a traffic volume lifetime is 4,608,000 kilobytes (10 megabytes per second for one hour) with the range of acceptable values between 2,560 and 536,870,912 kilobytes. Using shorter lifetimes can make it harder for an intruder to mount a successful key recovery attack because the attacker has less data encrypted under the same key to work with. The downside is shorter lifetimes require more CPU processing time for establishing new security associations.

The following example shortens the timed lifetime to 1,800 seconds (30 minutes) and the traffic-volume lifetime to 2,300,000 kilobytes. The **show crypto ipsec security-association-lifetime** command confirms the result.

```
Rtr1(config)#crypto ipsec security-association lifetime seconds 1800
Rtr1(config)#crypto ipsec security-association lifetime kilobytes 2300000


Rtr1#show crypto ipsec security-association-lifetime
Security association lifetime: 2300000 kilobytes/1800 seconds
Rtr1#
```

How These Lifetimes Work The SA and associated keys expire when either lifetime expires. To help ensure a new SA is ready when the old one expires, the new SA negotiation starts either 30 seconds or 256 kilobytes before the existing SA lifetime expires. If no traffic has passed through the tunnel during the entire life of the SA, a new one isn't negotiated until IPSec sees another packet that should be protected.

Step 3-3 Configure Crypto ACLs

Access lists, specifically crypto ACLs, play a substantial roll in how IPSec processes data flows. Crypto ACLs perform the following functions:

- Define the data flow to be protected by IPSec (permit statements).
- Define the data flow to pass unprotected by IPSec (deny statements).
- Check outbound traffic to see if it should be protected by IPSec.
- Check inbound traffic to discard traffic that should have been protected by IPSec. The assumption is if it's unprotected, it isn't from the source.
- Determine whether to accept requests for IPSec SAs for requested data flows when processing IKE negotiations.

Only extended IP ACLs can be used to create crypto ACLs. These extended IP ACLs select the IP traffic to encrypt by protocol, IP address, network, subnet, and port. While the basic ACL syntax is unchanged, the result is slightly different for crypto ACLs. Crypto ACLs are only deciding which traffic to protect and which to allow access through unprotected. A *permit statemen*t means the matching packet must be protected, while a *deny statement* means the matching packet needn't be protected. Any actual packet filtering must be done with another ACL. The crypto ACL determines only *if* the packet is protected. The level and type of protection is determined by the IKE and IPSec SAs.

Allowing one type of data traffic to receive one combination of IPSec protection is possible—such as authentication only—and other traffic to receive a different combination, such as both authentication and encryption. To do this, you must create two different crypto ACLs to define each of the two different types of traffic. The resulting ACLs are then used in different crypto map entries defining the different IPSec policies.

Using the Any Keyword Cisco strongly recommends avoiding the **any** keyword for specifying source or destination addresses whenever possible. A permit any any statement causes all outbound traffic to be protected and discards any unprotected inbound traffic, including packets for routing protocols, NTP, echo, echo response, and so forth.

If using the **any** keyword in a permit statement is unavoidable, be sure to precede the statement with as many deny statements as necessary to identify any traffic to be allowed through unprotected.

Mirrored ACLs The same outbound IPSec ACL that determines which outbound traffic will be protected by IPSec is also used to "catch" noncompliant inbound traffic. For example, if all TCP traffic is to be protected, then any unprotected inbound TCP traffic must have been introduced at some point other than the IPSec source. This is both an

elegant and a simple system, but it relies on having ACLs at each end, which are virtually mirror images of each other. This means the source addresses/ports and destination addresses/ports are literally reversed on the peer.

In the following example, using the chapter scenario, each peer will protect outbound TCP traffic, discard any unprotected inbound TCP traffic, and, because of the implicit deny any statement, each will allow all other traffic to pass unprotected and unfiltered.

```
Rtr1(config)#access-list 125 permit tcp 192.168.0.0 0.0.127.255
   192.168.130.0 0.0.0.255

Rtr2(config)#access-list 150 permit tcp 192.168.130.0 0.0.0.255
   192.168.0.0 0.0.127.255
```

Compare that to the next example, where the Rtr1 Admin decided to protect outbound UDP traffic without adding the same ACL statement to Rtr2. Now, as requested, all outbound UDP traffic will be protected. But maybe what isn't expected is the inbound UDP traffic continues to be unprotected and, therefore, mistaken as interlopers and discarded. While this is a somewhat glaring example and would probably be noticed quite quickly, it's an easy and troublesome mistake to make when new protocols or services are added to a network.

```
Rtr1(config)#access-list 125 permit tcp 192.168.0.0 0.0.127.255
   192.168.130.0 0.0.0.255
Rtr1(config)#access-list 125 permit udp 192.168.0.0 0.0.127.255
   192.168.130.0 0.0.0.255

Rtr2(config)#access-list 150 permit tcp 192.168.130.0 0.0.0.255
   192.168.0.0 0.0.127.255
```

In the next step, the crypto ACL is associated to a crypto map assigned to a specific interface.

Step 3-4 Configure Crypto Maps

Crypto map entries must be created for IPSec to set up SAs for traffic flows that must be protected. Crypto map entries created for IPSec set up security association parameters, tying together the various parts configured for IPSec, including

- Which traffic should be protected by IPSec (crypto ACL)
- Which IPSec peers the protected traffic can be forwarded to—the peers with which a security association can be established
- Which transform sets are to be used with the protected traffic
- How keys and security associations should be used or managed (or what the keys are, if IKE isn't used)
- Other parameters that might be necessary to define an IPSec SA

As with applying packet filtering, it's only possible to apply one crypto map set to a particular interface. You can apply the same crypto map to multiple interfaces to apply

the same policy to all included interfaces. The crypto map set can include a combination of IPSec using IKE and IPSec with manually configured SA entries.

To create more than one crypto map entry for a particular interface, use a unique sequence number (seq-num) of each map entry. The number assigned to the seq-num shouldn't be arbitrary; it's used to rank multiple crypto map entries within a crypto map set. Within a crypto map set, a crypto map entry with a lower seq-num has a higher priority and is evaluated before a map entry with a higher seq-num. Be sure to space the sequence numbers to allow for future technology to be inserted without having to completely reconfigure.

Use the global **crypto map** command to create a crypto map entry and enter the Crypto Map Configuration mode. Use the no form of the command to delete a crypto map entry or set. Use the **crypto map** ***map-name seq-num*** command without a keyword to modify an existing crypto map entry. Once a crypto map entry is created, you can't change the global configuration parameters because those parameters determine which **configuration** commands are valid at the crypto map level. For example, once a map entry is created as ipsec-isakmp, it can't be changed to ipsec-manual. It must be deleted and reentered.

The command syntax follows:

Rtr1(config)#crypto map *map-name seq-num* ipsec-manual
Rtr1(config)#crypto map *map-name seq-num* ipsec-isakmp [dynamic *dynamic-map-name*][discover]
Rtr1(config)#no crypto map *map-name* [*seq-num*]

map-name	Identifies the crypto map set. Assigned when the crypto map was created.
seq-num	A sequencing number assigned to a crypto map entry that determines its processing order. A lower number is used before a higher number.
ipsec-manual	Indicates IKE will *not* be used to establish the IPSec SAs for protecting the traffic specified by this crypto map entry.
ipsec-isakmp	Indicates IKE *will* be used to establish the IPSec SAs for protecting the traffic specified by this crypto map entry.
dynamic	(Optional) Specifies this crypto map entry is to reference a preexisting dynamic crypto map. *Dynamic crypto maps* are policy templates used in processing negotiation requests from a peer IPSec device. None of the **crypto map configuration** commands are available with this option.
dynamic-map-name	(Optional) Name of the dynamic crypto map set to use as the policy template.
discover	(Optional) Enables peer discovery. Added in v12.0(5)T to support Tunnel Endpoint Discovery (TED).

By default, no crypto maps exist and peer discovery isn't enabled.

After typing the **crypto map** command, the Crypto Map Configuration mode is invoked, as shown in the following output:

```
Rtr1(config)#crypto map testmap 50 ipsec-isakmp
% NOTE: This new crypto map will remain disabled until a peer
        and a valid access list have been configured.
Rtr1(config-crypto-map)#
```

From here, the **set** command can be used to define the crypto map parameters. You can also add a description, a match statement to link the crypto ACL that defines the traffic, and, finally, exit to end the session. The following example demonstrates pulling choices from the earlier exercises:

```
Rtr1(config)#crypto map testmap 50 ipsec-isakmp
Rtr1(config-crypto-map)#description VPN Link to branch in Tacoma, WA
Rtr1(config-crypto-map)#set peer 10.0.50.2
Rtr1(config-crypto-map)#set security-association lifetime seconds 1800
Rtr1(config-crypto-map)#set pfs group2
Rtr1(config-crypto-map)#set transform-set CPU-HOG
Rtr1(config-crypto-map)#match address 125
Rtr1(config-crypto-map)#exit
Rtr1(config)#
```

This next example demonstrates creating a second crypto map entry for *testmap* that creates a VPN to San Antonio, TX. Because of the larger seq-num (100), this choice is always considered after the original entries.

```
Rtr1(config)#crypto map testmap 100 ipsec-isakmp
Rtr1(config-crypto-map)#description VPN Link to branch in San Antonio, TX
Rtr1(config-crypto-map)#set peer 10.1.195.130
Rtr1(config-crypto-map)#set security-association lifetime seconds 2700
Rtr1(config-crypto-map)#set pfs group2
Rtr1(config-crypto-map)#set transform-set DES-ONLY
Rtr1(config-crypto-map)#match address 150
Rtr1(config-crypto-map)#exit
```

After crypto map entries are defined, the crypto map set is assigned to interfaces using the **crypto map** (interface IPSec) command, covered in the next section.

Step 3-5 Apply the Crypto Maps to the Interface

The last step in configuring IPSec is to apply the crypto map set to an interface with the **crypto map** command while in Interface Configuration mode. Use the no form of the command to remove the crypto map set from the interface. The command syntax is as follows:

Rtr1(config-if)#crypto map *map-name*

The IPSec SAs initialize as soon as the crypto map statement is entered. Because only one crypto map set can be assigned to an interface, crypto maps using the same map-name could have multiple entries designated by a different seq-num. The crypto map entry with the lowest seq-num is considered the highest priority and is evaluated first.

In the following example, crypto map testmap will be applied to the serial interface:

```
Rtr1(config)#interface serial 0
Rtr1(config-if)#crypto map testmap
Rtr1(config-if)#exit
```

Task 4 Test and Verify IPSec

The Cisco IOS software contains several **show, clear,** and **debug** commands that can be useful in testing and verifying IPSec and ISAKMP configuration and operations. The following steps provide a checklist to test and verify the correct configuration and operations of the VPN connection:

- Step 4-1 Display the configured IKE policies
- Step 4-2 Display the configured transform sets
- Step 4-3 Display the current state of the IPSec SAs
- Step 4-4 Display the configured crypto maps
- Step 4-5 Debug IKE events
- Step 4-6 Debug IPSec events

Step 4-1 Display the Configured IKE Policies

Use the **show crypto isakmp policy** command to view the parameters for each IKE policy, as the following example shows. Multiple policies are listed in priority order with the highest priority (first compared) listed first. The IKE default settings are shown at the end.

```
Rtr1#show crypto isakmp policy
Protection suite of priority 100
        encryption algorithm:   3DES - Triple DES (168 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #2 (1024 bit)
        lifetime:               43200 seconds, no volume limit
Protection suite of priority 200
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #2 (1024 bit)
        lifetime:               43200 seconds, no volume limit
Protection suite of priority 300
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Message Digest 5
        authentication method:  Pre-Shared Key
        Diffie-Hellman group:   #1 (768 bit)
        lifetime:               86400 seconds, no volume limit
Default protection suite
        encryption algorithm:   DES - Data Encryption Standard (56 bit keys).
        hash algorithm:         Secure Hash Standard
        authentication method:  Rivest-Shamir-Adleman Signature
        Diffie-Hellman group:   #1 (768 bit)
        lifetime:               86400 seconds, no volume limit
Rtr1#
```

Step 4-2 Display the Configured Transform Sets

Use the **show crypto ipsec transform-set** command to view the transform sets configured on the device. The command syntax shows an optional parameter that allows limiting the output to only a single transform set.

Rtr1#show crypto ipsec transform-set [tag *transform-set-name*]

Any time the keyword tag is used in **IOS** commands, it can be interpreted as "named" or "labeled" and is followed by a name. If no keyword and name are used, all transform sets configured are displayed as in the following example:

```
Rtr1#show crypto ipsec transform-set
Transform set DES-ONLY: { esp-des  }
   will negotiate = { Tunnel,  },

Transform set MD5-DES: { esp-des esp-md5-hmac  }
   will negotiate = { Tunnel,  },

Transform set CPU-HOG: { ah-md5-hmac  }
   will negotiate = { Tunnel,  },
   { esp-des esp-md5-hmac  }
   will negotiate = { Tunnel,  },

Transform set AH-ONLY: { ah-sha-hmac  }
   will negotiate = { Tunnel,  },
Rtr1#
```

Step 4-3 Display the Current State of the IPSec SAs

Use the **show crypto ipsec sa** command to display the settings of the current security associations. The command supports three parameters to limit the display and an option to generate a more detailed output of the selected choices. The syntax is as follows:

Rtr1#show crypto ipsec sa [map *map-name* | address | identity | interface-id] [detail]

map *map-name*	Limits the display to any existing security associations created for the crypto map set named *map-name*. (Optional)
Address	Lists all existing security associations, sorted by the destination address, and then by protocol (AH or ESP). (Optional)
Identity	Displays only the flow information, not the SA information. (Optional)
Interface-id	Limits the display to a single interface.
Detail	Displays detailed error counters. (Optional)

When no optional keyword is used, all security associations are displayed, sorted by interface, and then by traffic flow. Within a flow, the SAs are listed by protocol (ESP/AH) and direction (inbound/outbound).

```
Rtr1#show crypto ipsec sa

interface: Serial0
```

```
    Crypto map tag: testmap, local addr. 10.0.1.21

   local  ident (addr/mask/prot/port): (192.168.0.0/255.255.128.0/6/0)
   remote ident (addr/mask/prot/port): (192.168.130.0/255.255.255.0/6/0)
   current_peer: 10.0.50.2
     PERMIT, flags={origin_is_acl,}
    #pkts encaps: 60, #pkts encrypt: 60, #pkts digest 60
    #pkts decaps: 60, #pkts decrypt: 60, #pkts verify 60
    #pkts compressed: 0, #pkts decompressed: 0
    #pkts not compressed: 0, #pkts compr. failed: 0, #pkts decompress failed: 0
    #send errors 60, #recv errors 0

     local crypto endpt.: 10.0.1.21, remote crypto endpt.: 10.0.50.2
     path mtu 1500, media mtu 1500
     current outbound spi: 20890A6F

     inbound esp sas:
      spi: 0x36724AA4(913459876)
        transform: ah-md5-hmac esp-md5-hmac esp-des
        in use settings ={Tunnel, }
        slot: 0, conn id: 58, crypto map: testmap
        sa timing: remaining key lifetime (k/sec): (4607461/90)
        IV size: 8 bytes
        replay detection support: Y

     inbound ah sas:

     inbound pcp sas:

     outbound esp sas:
      spi: 0x338E791F(864975135)
        transform: ah-md5-hmac esp-md5-hmac esp-des
        in use settings ={Tunnel, }
        slot: 0, conn id: 59, crypto map: testmap
        sa timing: remaining key lifetime (k/sec): (4607461/90)
        IV size: 8 bytes
        replay detection support: Y

     outbound ah sas:

     outbound pcp sas:

Rtr1#
```

Step 4-4 Display the Configured Crypto Maps

Use the **show crypto map** command to display the crypto map configuration. The two optional parameters allow showing only the crypto map set applied to the specified interface or the crypto map set associated with the specified map-name. The syntax is as follows:

Rtr1#show crypto map [interface *interface* | tag *map-name*]

```
Rtr1#show crypto map
Crypto Map "testmap" 50 ipsec-isakmp
        Description: VPN Link to branch office in Tacoma, WA
        Peer = 10.0.50.2
        Extended IP access list 125
            access-list 125 permit tcp 192.168.0.0 0.0.127.255 192.168.130.0
```

```
   0.0.0.255
        Current peer: 10.0.50.2
        Security association lifetime: 2300000 kilobytes/1800 seconds
        PFS (Y/N): Y
        DH group:  group2
        Transform sets={ CPU-HOG, }
        Interfaces using crypto map testmap:
                Serial0
Rtr1#
```

Step 4-5 Debug IKE Events

```
Rtr1#debug crypto isakmp
Crypto ISAKMP debugging is on
Rtr1#
```

Step 4-6 Debug IPSec Events

```
Rtr1#debug crypto ipsec
Crypto IPSEC debugging is on
Rtr1#
```

Crypto System Error Messages for ISAKMP

Cisco IOS software can generate many useful system error messages for IKE/IPSec. The IOS sends these error messages to the console and, optionally, to a Syslog server. Understanding that not all system error messages indicate problems with the system is important. Some are purely informational and others could help diagnose problems with communications lines, internal hardware, or the system software.

> %CRYPTO-6-IKMP_AUTH_FAIL: Authentication method failed with host 10.0.50.2

How to Read System Error Messages System error messages begin with a percent sign (%) and are structured as follows:

> %FACILITY-SEVERITY-MNEMONIC: Message-text

FACILITY	Code consisting of two or more uppercase letters that indicate the facility to which the message refers. A facility can be a hardware device, a protocol, or a module of the system software.
SEVERITY	A single-digit code (0 to 7) reflecting severity. The lower the number, the more serious the situation. Logging traps are based on these values.
MNEMONIC	Code that uniquely identifies the error message.
Message-text	Text string describing the condition. Can contain detailed event information, including port numbers, network addresses, or memory addresses. Variable fields change from message to message and are represented here by short strings enclosed in square brackets ([ ]). A decimal number, for example, is represented as [dec].

For a detailed description of the error messages, search on the Cisco site for Cisco IOS System Error Messages.

Pay close attention to these messages during initial configuration and trouble shooting because they can provide the clues to what's going wrong. Once the configuration is up and running, if these messages are being logged to a server, they can help identify and solve problems that arise with data handling. The following is a sample IPSec system error message:

> %CRYPTO-6-IKMP_AUTH_FAIL: Authentication method [dec] failed with host [*IP_address*]

Explanation: The IKE process was unable to authenticate its SA with its remote peer.

> %CRYPTO-6-IKMP_MODE_FAILURE: Processing of [chars] mode failed with peer at [*IP_address*]

Explanation: Negotiation with the remote peer failed.

> %CRYPTO-6-IKMP_NOT_ENCRYPTED: IKE packet from [*IP_address*] was not encrypted and it should've been.

Explanation: A portion of the IKE is unencrypted and a portion is encrypted. This message should have been encrypted, but wasn't. This would occur if the ACLs aren't mirror images of each other on the peers. See Task 3, Step 3 for details.

Configuring IPSec Manually

You can configure the IPSec keys manually. This section provides a brief overview of how this is done and why manual key use isn't generally recommended.

Use the **set session-key** command in Crypto Map Configuration mode to specify the IPSec session keys manually within a crypto map entry. The command is only available for ipsec-manual crypto map entries. They aren't used with ipsec-isakmp entries because the SAs and corresponding keys are automatically established via the IKE negotiation. Use the no form of this command to remove IPSec session keys from a crypto map entry. The command uses the following syntax:

```
Rtr1(config)#crypto map map-name seq-num ipsec-manual
Rtr1(config-crypto-map)#set session-key {inbound | outbound} ah spi
hex-key-string
Rtr1(config-crypto-map)#set session-key {inbound | outbound} esp spi cipher
hex-key-string [authenticator hex-key-string]
Rtr1(config-crypto-map)#no set session-key {inbound | outbound} ah
Rtr1(config-crypto-map)#no set session-key {inbound | outbound} esp
```

inbound	Sets the inbound IPSec session key. Both inbound and outbound must be set.
outbound	Sets the outbound IPSec session key. Both inbound and outbound must be set.
ah	Sets the IPSec session key for the AH protocol. Use if the crypto map entry transform set includes an AH transform.
esp	Sets the IPSec session key for the ESP. Use if the crypto map entry transform set includes an ESP transform.
spi	Specifies the security parameter index (SPI), a number used to uniquely identify a security association. *SPI* is an arbitrary value assigned by the admin between 256 and 4,294,967,295 (FFFF FFFF).
hex-key-string	The session key in HEX format. The key is an arbitrary HEX string of 8, 16, or 20 bytes. If the crypto map's transform set includes A DES algorithm—use at least 8 bytes per key An MD5 algorithm—use at least 16 bytes per key An SHA algorithm—use at least 20 bytes per key Keys longer than the previous sizes are truncated.
cipher	Indicates the key string is to be used with the ESP encryption transform.
authenticator	(Optional) Indicates the key string is to be used with the ESP authentication transform. This argument is required only when the crypto map entry's transform set includes an ESP authentication transform.

Use the following rules to determine when to create keys:

- If the crypto map transform set includes an AH protocol, IPSec keys must be defined for AH for both inbound and outbound traffic.
- If the crypto map transform set includes an ESP encryption protocol, IPSec keys must be defined for ESP encryption for both inbound and outbound traffic.
- If the crypto map transform set includes an ESP authentication protocol, IPSec keys must be defined for ESP authentication for inbound and outbound traffic.

The SPI is used to identify the security association used with the crypto map. When defining multiple IPSec session keys within a single crypto map, it's permissible to assign the same SPI number to all keys. Session keys must be the same for both peers. Before assigning a SPI, confirm with the peer's admin that the same SPI isn't used more than once for the same destination address/protocol combination.

If a session key is changed, the SA using the key will be deleted and reinitialized.

```
Rtr1(config)#crypto map testmap2 50 ipsec-manual
Rtr1(config-crypto-map)#set peer 10.0.10.133
Rtr1(config-crypto-map)#match address 100
Rtr1(config-crypto-map)#set transform-set encrypt-des
Rtr1(config-crypto-map)#set session-key inbound esp 300 32a87b445c98701c
Rtr1(config-crypto-map)#set session-key outbound esp 300 32a87b445c98701c
```

Configuring IPSec Manually Is Not Recommended

Cisco recommends using IKE to set up the SAs because it's difficult to ensure the SA values match between peers and D-H is a more secure method for generating secret keys between peers. Other reasons not to configure IPSec manually include the following:

- Manual keying doesn't scale well.
- The result can be insecure because of difficulty in creating secure keying material manually.
- Manually established SAs never expire.
- ACLs used for crypto map entries using ipsec-manual are restricted to a single permit statement. Any other entries are ignored.

Chapter Review

This chapter looked at steps involved in configuring IPSec with preshared keys. The steps and related commands are summarized in the following task list.

Task 1 Prepare for IKE and IPSec

- Step 1-1 Identify IPSec peers
- Step 1-2 Determine the IKE (IKE Phase 1) policies
- Step 1-3 Determine the IPSec (IKE Phase 2) policies
- Step 1-4 Check the current configuration
  show running-config
  show isakmp
  show crypto map
- Step 1-5 Ensure the network works without encryption
  ping
- Step 1-6 Ensure access control lists are compatible with IPSec
  show access-lists

Task 2 Configure IKE

- Step 2-1 Enable or disable IKE
  crypto isakmp enable
- Step 2-2 Create IKE policies
  crypto isakmp policy
  authentication
  encryption
  hash
  lifetime
- Step 2-3 Configure preshared keys
  crypto isakmp key
- Step 2-4 Verify the IKE configuration
  show crypto isakmp policy

Task 3 Configure IPSec

- Step 3-1 Configure transform set suites
  crypto ipsec transform-set
- Step 3-2 Configure global IPSec security association lifetimes
  crypto ipsec security-association lifetime
- Step 3-3 Configure crypto ACLs
  access-list
- Step 3-4 Configure crypto maps
  crypto map
- Step 3-5 Apply the crypto maps to the interface
  interface
  crypto map

Task 4 Test and verify IPSec

- Step 4-1 Display the configured IKE policies
  show crypto isakmp policy
- Step 4-2 Display the configured transform sets
  show crypto ipsec transform set
- Step 4-3 Display the current state of the IPSec SAs
  show crypto ipsec sa
- Step 4-4 Display the configured crypto maps
  show crypto map
- Step 4-5 Debug IKE events
  debug crypto isakmp
- Step 4-6 Debug IPSec events
  debug crypto ipsec

Questions

1 Which one of the following is *not* one of the tasks required to configure IPSec for Preshared Keys?

A. Configure IPSec

B. Prepare for IKE and IPSec

C. Test and verify IPSec

D. Configure the crypto map

2. Which of the following VPN products would be common for mobile users?

A. Cisco 1700 router

B. Cisco 900 Cable/DSL router

C. Cisco VPN Software Client

D. Cisco VPN Hardware Client

3. Which one of the following is a hybrid protocol that implements the Oakley key exchange?
   - **A.** IPSec
   - **B.** Crypto map
   - **C.** IKE
   - **D.** Hash algorithm
4. Which of the following is a peer authentication method?
   - **A.** 3DES
   - **B.** SHA-1
   - **C.** MD5
   - **D.** Preshared keys
5. Which of the following preparation steps is done using the **ping** command?
   - **A.** Identify IPSec peers
   - **B.** Check the current configuration
   - **C.** Ensure the network works without encryption
   - **D.** Ensure access control lists are compatible with IPSec
6. Which one of the following is *not* an IKE Phase 1 parameter?
   - **A.** Encryption algorithm
   - **B.** Traffic to protect
   - **C.** Authentication method
   - **D.** DH key exchange group
7. To make sure the router ACLs are IPSec-compatible, which is *not* required to be permitted?
   - **A.** Port 500
   - **B.** Port 510
   - **C.** Protocol 51
   - **D.** Protocol 50
8. If the **crypto isakmp policy** command were used to create policies with the following priorities, which would be processed first?
   - **A.** 1000
   - **B.** 500
   - **C.** 12
   - **D.** 25

9. If the crypto isakmp policy lifetime is set to 43,200, to what does the 43,200 refer?
   A. 43,200 bytes of protected throughput
   B. 43,200 hours
   C. half a day
   D. 43,200 lines of protected throughput
10. Which command shows the IKE policies and default values?
   A. show running-config
   B. show isakmp policy
   C. show crypto ike policy
   D. show crypto isakmp policy
11. A transform set can contain up to how many transforms?
   A. 4
   B. 6
   C. 3
   D. 1
12. Which is *not* a function of a crypto ACL?
   A. Define the dataflow to be protected by IPSec
   B. Discard inbound traffic that should have been protected by IPSec
   C. Filter outbound traffic for access to the Internet
   D. Define the data flow to pass unprotected by IPSec
13. Which of the following is *not* true?
   A. The crypto ACL determines the traffic to be protected
   B. The global **crypto map** command ties together the IPSec parameters
   C. The interface **crypto map** command applies the crypto map to an interface
   D. The global **crypto map policy** command sets the implementation priority
14. Which command shows IPSec performance indicators?
   A. show crypto map
   B. show crypto ipsec sa
   C. show crypto ipsec transform-set
   D. show crypto isakmp policy

15. Which statement is *not* true about the ipsec-manual form of the **crypto map** command?

    A. It doesn't scale well

    B. The result can be insecure because of difficulty in manually creating secure keying material

    C. It enhances the flexibility of the crypto ACLs

    D. Manually established SAs never expire

Answers

1. **D.** Configure the crypto map
2. **C.** Cisco VPN Software Client
3. **C.** IKE
4. **D.** Preshared keys
5. **C.** Ensure the network works without encryption
6. **B.** Traffic to protect
7. **B.** Port 510
8. **C.** 12
9. **C.** half a day
10. **D.** show crypto isakmp policy
11. **C.** 3
12. **C.** Filter outbound traffic for access to the Internet
13. **D.** The global **crypto map policy** command sets the implementation priority
14. **B.** show crypto ipsec sa
15. **C.** It enhances the flexibility of the crypto ACLs

CHAPTER 11

Cisco IOS IPSec Certificate Authority Support

In this chapter, you will learn to:

- Learn about CA support overview
- Configure CA support tasks
- Understand RSA encrypted nonces overview

This chapter looks at configuring Cisco IOS IPSec using a certificate authority (CA). After presenting an overview of the configuration process, the chapter covers planning before beginning configuration, and the major tasks and steps required for successful implementation of an IPSec network. These steps involve how to configure CA support, as well as IKE policies and IPSec. Descriptions of all required steps are provided.

CA Support Overview

With public/private key encryption technology, each key is a unique encryption device. If, in fact, no two keys are ever identical, it's possible to use a key to identify its owner. Keys are always created and used in pairs: one is the private key; the other is the public key. You can then encrypt data with the public key, which only the corresponding private key can decrypt.

The public key can be distributed to anyone who needs to share secure information with the owner of the private key. The private key, on the other hand, is never duplicated or distributed to anyone: it remains secure on the owner's computer or server. While not as secure as limiting keys to pairs, the same public key can be duplicated and distributed to multiple users who need to share information with the private key owner. The public key holders would be unable to exchange information using the shared key. The downside is the public key no longer identifies a unique host.

A fundamental problem with public key cryptosystems is this: all users must be careful to make sure they're always encrypting to the correct person's key. In an environment where keys can be physically exchanged in person or configured by a single administrator, this shouldn't be a problem. Unfortunately, neither of these key-exchange methods scales well, particularly over large geographical areas. One common solution for key exchanges is to use public servers, such as e-mail servers. Unfortunately, the threat of man-in-the-middle attacks must be considered a potential threat. Figure 11-1 shows

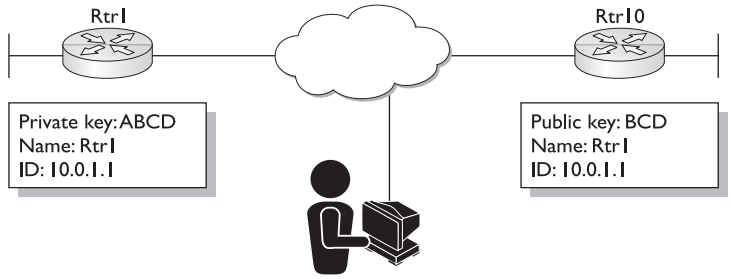

Figure 11-1 Rtr1 sharing secure data with Rtr10 using a public key

Rtr1 communicating securely with Rtr10, using a public key BCD sent to Rtr10. The attacker can see the traffic, but can't decrypt it.

In the man-in-the middle scenario, the culprit posts a phony public key using the real name and user ID of the legitimate holder of the private key with whom the user expects to share data. Now any data encrypted using the phony public key can only be decrypted by the holder of the matching private key—the attacker. In Figure 11-2, the hacker has posted a new public key using Rtr1's name and ID. Rtr10, not realizing the ruse, implements the new key. The attacker can now decrypt the traffic, but Rtr1 can't.

In a public key environment, it's absolutely critical that the public key used to encrypt data is, in fact, a legitimate public key of the intended recipient and not a forgery. Common sense indicates that requiring an administrator to configure both devices or physically deliver the public keys won't scale well. Even if this is physically possible, the process would present a significant cost burden, even in regional networks. There must be another way. Digital certificates (certs) can simplify the task of sharing public keys, while verifying the public key belongs to the "real" owner.

Digital Certificates

A *digital certificate* is a form of credential much like a driver's license or a passport in the paper-based world. Like its paper counterparts, the digital certificate has information on

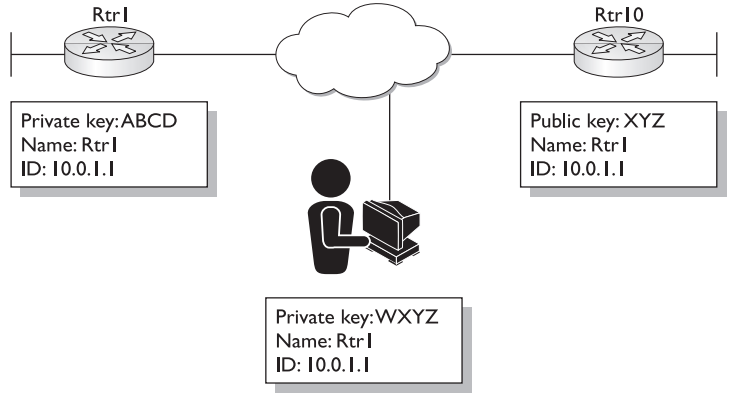

Figure 11-2 The attacker and Rtr10 now share a key, allowing the attacker to decrypt the messages.

it that identifies the holder, plus some uninvolved third-party authorization, which indicates they confirmed the holder's identity.

A digital certificate has additional information included with the holder's public key that helps others to verify the key is genuine. This additional information, like a person's picture on a driver's license or a passport, can thwart attempts to substitute an unauthorized public key.

A digital certificate contains the following three items:

- Public key
- Certificate information—Identifying information about the holder, such as ID, name, and so forth
- One or more digital signatures

The *digital signatures* indicate the certificate information was verified and attested to by a trusted independent third party. Understanding that the digital signature doesn't guarantee the authenticity of the certificate as a whole is important. What it verifies is only that the signed identity information belongs to or is bound to the attached public key. Don't let this get too complicated. In the case of a passport, the government is verifying that the picture and identifying information belong to the passport number. But, the passport can be expired or revoked by a court order. Similarly, digital certificates can expire or be revoked.

Conceptually, a digital certificate is a public key with a tag containing one or more forms of ID, plus a seal of approval by a trusted third party. Figure 11-3 shows a conceptual representation of a digital certificate.

Certificate Distribution

Chapter 10 looked at preshared public keys, but this manual distribution method is only practical to a certain point. After that, having an exchange system that can provide the necessary security, storage, and exchange mechanisms becomes necessary, so coworkers, business partners, and even strangers could establish secure communications, if necessary.

These public key exchange systems can take the form of storage-only repositories called *Certificate Servers*, or they can be much more structured systems providing additional key management features. These latter systems are called *Public Key Infrastructures*.

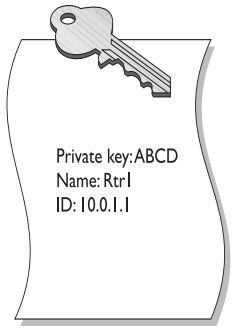

Figure 11-3 Conceptual representation of a digital certificate

Certificate Servers

A certificate server, also called a *cert server* or a *key server*, is a database service running on an existing or dedicated server that allows users to submit and retrieve digital certificates. The cert server typically offers additional administrative features that allow the company to implement and maintain its security policies. For example, the cert server can be configured to allow only certain types of keys to be stored.

Public Key Infrastructures (PKI)

A Public Key Infrastructures (PKI) includes the storage capabilities of a certificate server, but it also provides additional certificate management functions, such as the capability to issue, store, retrieve, revoke, and trust certificates. A powerful feature of a PKI system is the concept of a Certification Authority (CA). *CA*s are responsible for managing certificate requests and issuing certificates to participating IPSec network peers. These services provide centralized key management for the participating peers. CAs simplify the administration of IPSec network devices (peers) in networks containing multiple IPSec-compliant devices, such as with the Cisco Secure PIX Firewall units and Cisco routers.

The CA creates the digital certificate, and then digitally signs it using its own private key. Any CA client can then use the shared CA public key to authenticate a certificate's digital signature and, thereby, the integrity of the certificate contents, including the certificate holder public key and the identity of the certificate holder.

Conceptually, the CA functions in the digital environment much like the government passport or driver's license office in the nondigital world.

IPSec with CAs

When using IPSec with a CA, configuring keys among all the potential IPSec peers isn't necessary. Instead, each participating device is individually enrolled with the CA, requesting a certificate for the device. Once this is done, each participating device can dynamically authenticate itself to all the other participating peers.

Adding a new IPSec device to an existing network is similarly simplified. In each case, it's only necessary to configure the new device to request a certificate from the CA.

How CA Certs Are Used by IPSec Peers

As you learned in Chapter 9, any time two IPSec peers want to establish an IPSec-protected communication session between them, they must first authenticate each other. This authentication is done with IKE and, without it, IPSec protection can't occur.

Without a CA, the peers authenticate themselves to the remote peer, using either preshared keys (Chapter 10) or RSA-encrypted nonces (see the section "RSA Encrypted Nonces Overview.") Both methods require previously configured keys on both peers. PIX Firewall devices currently don't support RSA-encrypted nonces.

With a CA, each peer authenticates itself by sending a valid digital certificate to the remote peer. The peer's unique certificate was issued and validated by the CA, encapsulates the peer's public key and IDs, and works because all participating peers recognize the CA as an authenticating authority. This process is called *IKE with an RSA signature*.

Each peer can continue sending its own certificate for additional IPSec sessions, as well as to additional IPSec peers, until the certificate expires. When a certificate expires, the peer administrator must request a new one from the CA.

Certificate Revocation List (CRL)

CAs can revoke otherwise valid certificates for peers that will no longer participate in IPSec. These revoked certificates are listed in a certificate revocation list (CRL), which each peer can check before accepting another peer's certificate. Revoked certificates, like expired passports or driver's licenses, aren't recognized as valid by other peers.

Registration Authority (RA)

Some CAs extend their reliability and availability by supporting registration authority (RA) as part of their implementation. An *RA* is a server that acts as a proxy for the CA, so CA functions can continue when the CA is offline or otherwise unavailable.

Cisco IOS CA Standards

Cisco IOS supports the following open CA standards:

- **Internet Key Exchange (IKE)** A hybrid protocol that implements Oakley and Skeme key exchanges inside the ISAKMP framework. IKE is covered in Chapters 9 and 10.
- **Public-Key Cryptography Standard #7 (PKCS #7)** A standard developed by RSA Data Security, Inc. used to encrypt and sign certificate enrollment messages.
- **Public-Key Cryptography Standard #10 (PKCS #10)** A standard syntax developed by RSA Data Security, Inc. for certificate requests.
- **RSA Keys** *RSA* is the public-key cryptographic system developed by RSA Data Security, Inc. RSA keys come in pairs: one public key and one private key.
- **X.509v3 Certificates** An *X.509 certificate* is a collection of a standard set of fields containing information about a user or device and their corresponding public key. The X.509 standard defines what information goes into the certificate and describes how to encode it (the data format). The encoded information is used with the IKE protocol when authentication requires public keys.

Simple Certificate Enrollment Protocol (SCEP)

Developed by a Cisco, Verisign, Entrust, Microsoft, Netscape, and Sun Microsystems initiative, Simple Certificate Enrollment Protocol (SCEP) provides a standard way of managing the certificate lifecycle. This initiative is important for furthering open development for certificate handling protocols that can help ensure interoperability with devices from many vendors.

SCEP provides the following two authentication methods:

- Manual authentication
- Authentication based on a preshared secret

Manual Mode

In the Manual mode, the entity that submits the request is required to wait until the CA operator can verify its identity, using any reliable out-of-band method. An MD5 hash "fingerprint" generated by and included in the PKCS10 must be compared out-of-band between the SCEP clients and CAs (or RAs, if appropriate) to enable verification.

Preshared Secret Mode

With a preshared secret method, the CA server distributes a shared secret to the end entity, which can then be used to associate an enrollment request uniquely with the end entity. To maintain the integrity of the method, the distribution of the shared secret must be private, allowing only the end entity to know the secret.

Challenge Passwords

When an enrollment request is initiated, the end entity is asked to provide a challenge password. With the *preshared secret* method, the end entity must type in the distributed secret as the password. With the *manual authentication* method, the challenge password is also required because the server might challenge an end entity for a password before any certificate can be revoked. Eventually, the challenge password is included as a PKCS#10 attribute and is sent to the CA server as encrypted data. The PKCS#7 envelope protects the privacy of the challenge password using DES encryption.

CA Servers Interoperable with Cisco Routers

CA interoperability permits Cisco IOS devices, PIX Firewalls, Cisco VPN Hardware devices, and CA servers to communicate so the VPN device can obtain and use digital certificates supplied by the CA. While IPSec can be implemented in the network without using a CA, the CA with SCEP provides enhanced manageability and scalability for IPSec.

The list of CAs supported could vary from VPN platform to platform, for example, PIX Firewalls might not support the same options as IOS devices or VPN hardware devices. A good idea is always to check the Cisco online documentation for the particular device and version of the operating system (OS) to confirm support and appropriate version numbers. The following are CA providers that support SCEP to interoperate for enrolling Cisco IOS routers:

- VeriSign
- Entrust Technologies, Inc.
- Baltimore Technologies
- Microsoft Corporation

VeriSign OnSite 4.5

The VeriSign's OnSite 4.5 solution provides a fully integrated enterprise PKI to issue, control, and manage IPSec certificates for PIX Firewalls and Cisco routers. VeriSign administers the CA themselves, providing the certificates as a service. VeriSign OnSite service specifics include the following:

- **Requirements**—No local server requirements. Routers must be running IOS v 12.0.6 or higher and must have a retry count set to greater than 60.
- **Standards supported**—SCEP, x509 certificate format, and PKCS# 7, 10, 11, and 12.

For more information on VeriSign products and features, consult their web site at www.verisign.com.

Entrust Technologies

One major difference between CA servers is where the server is located and who administers it. *Entrust* is a software solution that's installed and administered by the user on a local network server. The Entrust/PKI service is able to issue digital IDs to any device or application supporting the X.509 certificate standard. This implementation offers a secure, flexible, and low-cost solution from one PKI. Entrust/PKI service specifics include the following:

- **Requirements**—Entrust software runs on the Windows servers (for Cisco interoperability), Solaris 2.6, HP-UX 10.20, and AIX 4.3 OSs. Requires RSA usage-keys configured on the routers. Cisco IOS release 11.(3)5T and later.
- **Standards supported**—CA services, RA capabilities, SCEP, and PKCS#10.

For more information on Entrust products and features, consult their web site at www.entrust.com.

Baltimore Technologies

Baltimore Technologies supports SCEP in UniCERT, its CA server, as well as the PKI Plus toolkit, making it easy for customers to implement certificates in their networks. *UniCERT* is a software solution installed and administered by the user on a local network server. Baltimore Technologies service specifics include the following:

- **Requirements**—The current release of the UniCERT CA software module is available for Windows servers. Must use Cisco IOS release 12.0(5)T and later.
- **Standards Supported**—The following standards are supported with this CA server: X509 v3, CRL version 2, PKCS# 1, 7, 10, 11, and 12, and many more protocols.

For more information on Baltimore products and features, consult their web site at www.baltimore.com.

Microsoft Windows 2000 Certificate Services Microsoft has integrated SCEP support into the Windows 2000 CA server via the Security Resource Kit for Windows 2000. This SCEP support allows Microsoft clients to obtain certificates and certificate revocation information from Microsoft Certificate Services for all of Cisco's VPN security solutions.

- **Requirements**—Intel-based PC running Windows 2000 Server. Cisco IOS release 12.0(5)T or later.

- **Standards Supported**—The following standards are supported with this CA server: X.509 version 3, CRL version 2, PKCS #7, #10, and #12, and many more protocols.

For more information on Microsoft products and features, consult their web site at www.microsoft.com.

Enroll a Device with a CA

The basic process for enrolling an IPSec peer with a CA includes the following steps. Most of these steps have been automated using the SCEP protocol supported by many CA server providers. Each CA vender determines how long certificates are valid.

- **Step 1** Configure the router for CA support.
- **Step 2** Generate a public and private key pair on the router.
- **Step 3** The router authenticates the CA server.
- **Step 4** The router sends a certificate request to the CA.
- **Step 5** The CA creates and signs an identity (ID) and root certificate and, optionally, an RA certificate.
- **Step 6** The CA sends the certificates to the router and posts the certificates in its public repository (directory).

The configuration commands are covered under Task 2 "Configure CA support" activities.

Configure CA Support Tasks

Configuring for RSA signatures consists of five major tasks, each requiring multiple steps. This chapter covers the CA configuration tasks and steps in detail, but those tasks and steps identical to those covered in Chapter 10 for preshared keys aren't repeated. Refer to Chapter 10 for a detailed explanation of these steps.

The five major tasks Cisco uses and which you can expect on the exam to configure CA support are as follows. The major difference is the insertion of a new Task 2 specifically for configuring CA support.

- Task 1 Prepare for IKE and IPSec
- Task 2 Configure CA support
- Task 3 Configure IKE
- Task 4 Configure IPSec
- Task 5 Test and verify IPSec

A summary task list showing the five tasks broken down to their individual steps and key commands is included in the Summary. As in Chapter 10, the steps are numbered to include the task number, as well as to help keep them straight.

Figure 11-4 shows the example networks from Chapter 10. These networks provide an example scenario used throughout this chapter. The goal is to create a secure VPN tunnel between Rtr1 company main office and Rtr2 at one of almost 100 branch offices in North America, Europe, and Africa. The assumption is this: the main office has reserved networks192.168.0.0 through 192.168.127.0 for itself and will use one class C for each branch in the remaining 192.168.128.0 to 192.168.255.0 addresses.

Because not all configuration steps are repeated in this chapter, using the same scenario as Chapter 10 means a person performing the configurations along with the text can refer to Chapter 10 for assistance.

Task 1—Prepare for IKE and IPSec

Successful implementation of an IPSec network requires testing of the existing network and advance planning before any configuration begins. Insufficient testing and planning can lead to troubleshooting problems or configuration errors. Some preparation and planning steps include the following:

Step 1–1 and Step 1–2 are covered in detail in this chapter. The other steps are presented for review purposes, but aren't covered in this chapter. Chapter 10 has detailed coverage of Steps 1–3 through 1–6.

- Step 1–1 Plan for CA support
- Step 1–2 Determine the IKE (IKE phase one) policies
- Step 1–3 Determine the IPSec (IKE phase two) policies
- Step 1–4 Check the current configuration

  show running-configuration

  show isakmp [policy]

  show crypto map
- Step 1–5 Ensure the network works without encryption

  ping

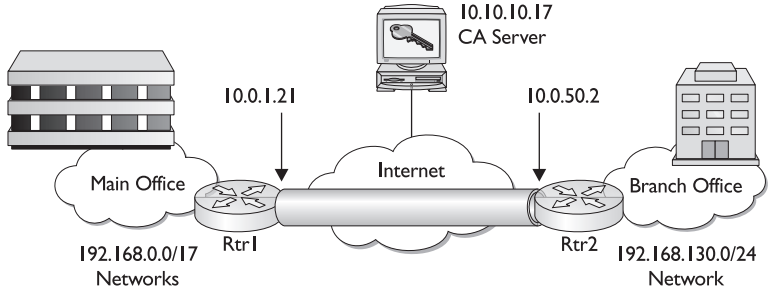

Figure 11-4 Chapter scenario VPN session to be configured

- Step 1–6 Ensure access control lists are compatible with IPSec

  show access-lists

The order of these steps could be varied based on personal preferences and implementation circumstances.

Step 1–1 Plan for CA support

Configuring CA is a complicated process, but having a detailed plan reduces the chances of improper configuration. Some basic information to be decided or collected includes the following items:

- **Determine which CA server is being used.** The organizations providing CA services vary in capabilities, requirements, and the resulting configuration. Once the service provider is selected, the requirements could include the RSA key type required, CRL capabilities, and support for RA mode.
- **Identify the CA server IP address, host name, and URL.** Information to use Lightweight Directory Protocol (LDAP).
- **Identify the CA server administrator contact information.** This arranges for the certificates to be validated if the process isn't automatic.

The actual information required for all CA features varies with the organization providing the CA support and can change over time with a specific organization. A good idea is always to check the CA supplier and Cisco documentation. Figure 11-5 shows the information that would be useful in configuring VeriSign CA support.

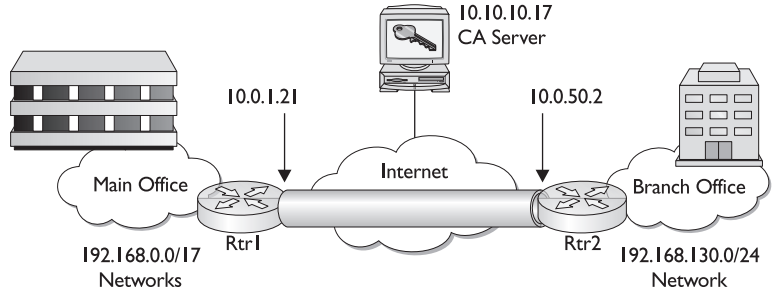

Parameter	CA Server
Type of CA Server	Verisign
Domain Name	ca-test.com
CA Server IP Address	10.10.10.17
Authentication URL	http://ca-ipsec.verisign.com/
Administrator Contact	(800) 555-1111

Figure 11-5 Configuration parameters for VeriSign CA support

Step 1–2 Determine the IKE (IKE phase one) policies

The objective here is to develop one or more IKE security policies, based on the overall company security policy. Each policy is going to require decisions about five security options: authentication method, encryption algorithm, hashing algorithm, Diffie–Hellman group, and SA lifetime.

Defining the IKE phase one policies is basically the same as with preshared keys except the authentication method will use RSA Signatures. Cisco IOS software supports preshared keys, RSA encrypted nonces, or RSA signatures to authenticate IPSec peers.

To complete the IKE planning process, what would make sense is to create a table of the preferred combination of security features, plus one or more fallback options for those devices or locations that can't support the preferred package. The resulting table might look like the following:

Parameter	Preferred (stronger)	2nd Choice	3rd Choice
Encryption algorithm	3des	des	des
Hash algorithm	sha	sha	md5
Authentication method	rsa-sig	rsa-sig	rsa-sig
DH key exchange group	2	2	1
IKE SA lifetime	43,200	43,200	86,400

NOTE You can't use SHA together with DES encryption on Cisco's VPN software client version 3.6. Part of the problem with determining a set of preferences is being aware of the different platform limitations.

Using our scenario for this chapter, the *Preferred* column represents our configuration preferences for all North American branches and would be the only options configured. The *2nd Choice* column would be for those overseas branches that can support the Preferred configuration, except for the export restriction on triple DES. Assuming they don't form VPN sessions with anyone else, this would be the only choice configured on those devices. The *3rd Choice* column might represent small branches using small routers with limited resources and options or telecommuters using Cisco's VPN software. This would also represent the lowest level of parameter options that the main office will accept for a VPN connection. The only device that will have all three choices configured will be the main office router.

Task 2—Configure CA Support

Configuring Cisco IOS CA support can be quite complicated and requires ten steps to complete. The planning steps performed in Task 1 and the first three commands in this section ensure the actual router CA support configuration goes smoothly. The process necessary to configure CA support on Cisco routers includes the following steps:

- Step 2–1 Manage the NVRAM memory usage (optional)
- Step 2–2 Set the router time and date

- Step 2–3 Configure the router host name and domain name
- Step 2–4 Generate an RSA key pair
- Step 2–5 Declare a CA
- Step 2–6 Authenticate the CA
- Step 2–7 Request your own certificate
- Step 2–8 Save the configuration
- Step 2–9 Monitor and maintain CA interoperability (optional)
  Request a CRL. Delete your router's RSA keys.

  Delete both public and private certificates from the configuration.

  Delete peer's public keys.
- Step 2–10 Verify the CA support configuration

Step 2–1 Manage the NVRAM Memory Usage (Optional)

Certificates and CRLs are stored in NVRAM on the router, and each certificate and CRL requires a moderate amount of memory. The types of certificates stored locally on the router include the following:

- The router's own certificate
- The CA's root certificate
- Two RA certificates if the CA supports RAs

If the CA doesn't support an RA, there will be one CRL stored on the router. Otherwise, there will be multiple CRLs.

Depending on the router model, the NVRAM available, and whether the CA supports RAs, storing certificates and CRLs locally shouldn't be a problem. But memory might become an issue with smaller routers, limited NVRAM, and a CA that supports RAs. If the CA supports an RA and a large number of CRLs end up being stored on the router, it can consume a large amount of NVRAM space.

To preserve limited NVRAM, you can specify that certificates and CRLs won't be stored locally, but can be retrieved from the CA as needed. While this method can save NVRAM, it could cause a slight performance reduction and puts additional traffic out on the network.

The crypto ca certificate query Command Use the global configuration mode **crypto ca certificate query** command to specify that certificates and CRLs shouldn't be stored locally but, instead, retrieved from the CA as needed. This command puts the router into Query mode. To cause certificates and CRLs to be stored locally (the default), use the no form of this command. The command syntax is as follows:

```
Rtr1(config)#crypto ca certificate query

Rtr1(config)#no crypto ca certificate query
```

NOTE Remember, this command should only be used if NVRAM resources are inadequate to service the CA support locally.

Step 2–2 Set the Router Time and Date

For CA support to function properly, the router must have an accurate time and date to enroll with a CA server. The clock must be set accurately before generating RSA key pairs and enrolling with the CA server because the keys and certificates are time-sensitive.

Like many time-dependent processes, CA support features use the router software clock. The *software clock* runs from when the system powers up or a **reload** command is used and keeps track of the current date and time. The software clock can be set from a number of sources and, in turn, can be used to distribute the current time through various mechanisms to other systems. If no other system or process initializes the clock, it uses a date in the past as the seed value. (My lab 1720 uses 16:00:00 PST Sun Feb 28 1993).

If a router has a built-in hardware clock, the software clock is initially set at power up or reload, based on the time in the hardware clock. The software clock can then be updated from the following time sources:

- Network Time Protocol (NTP)
- Simple Network Time Protocol (SNTP)—some models
- Manual configuration

Because the software clock can be dynamically updated, it has the potential to be more accurate than the hardware clock. In fact, updating the hardware clock from the software clock after synchronization is common.

The software clock can provide time to the following services:

- CA support features
- Access lists
- Logging and debugging messages
- User show commands
- NTP
- The hardware clock

The software clock stores time internally based on Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT). The next sections look at configuring information about the local time zone and summer time (daylight saving time), so the time is always displayed correctly relative to the local time zone.

Network Time Protocol (NTP) NTP is an industry protocol developed to facilitate time synchronization of network devices. NTP Version 3 uses UDP transport and is documented in RFC 1305.

A specified device in a NTP network usually gets its time from an *authoritative time source,* such as a radio clock or an atomic clock attached to a time server. NTP then periodically distributes this time across the network. NTP is an extremely efficient protocol,

requiring no more than one packet per minute to synchronize two machines to the accuracy of within a millisecond of one another.

NTP uses the concept of a stratum to describe how many NTP "hops" a device is away from an authoritative time source. A stratum 1 time server typically has an authoritative time source directly attached, a stratum 2 time server receives its time via NTP from a stratum 1 time server, and so on.

Devices in a Cisco network can be configured to be a NTP server, master, or peer. The Cisco implementation of NTP doesn't support stratum 1 service, recommending instead that time service be synchronized with a public NTP server available in the IP Internet. Note, it's important to recognize that not all Cisco devices support NTP fully. PIX firewalls just added NTP support with v6.2 of the PIX OS, while IOS devices added support with IOS 10.0.

NTP devices avoid synchronizing to a device with inaccurate time in two ways. First, NTP will never synchronize to a device that isn't synchronized itself. Second, NTP will compare the time reported by several devices and won't synchronize to a device with a time significantly different than the others, even if its stratum is lower. This strategy effectively builds a self-organizing tree of NTP servers.

NTP devices are usually statically configured to create associations; each device is given the IP address of any devices with which it should form associations. In a LAN environment, configuring NTP to use IP broadcasts is often advantageous. This implementation reduces configuration complexity because each device is configured to send or receive broadcast messages.

Use the global configuration **ntp peer** command to configure the software clock to synchronize a peer or to be synchronized by a peer. To disable the feature, use the no form of this command. The syntax is as follows:

Rtr1(config)#ntp peer ip-address [version *ver-num*] [key *keyid*]
[source *interface*] [prefer]

Rtr1(config)#no ntp peer ip-address

ip-address	IP address of the peer providing, or being provided, the clock synchronization
ver-num	NTP version number (1 to 3)* (Optional)
keyid	Authentication key to use when sending packets to this peer (Optional)
source	Names the interface (Optional)
interface	Name of the interface from which to pick the IP source address (Optional)
prefer	Makes this peer the preferred peer that provides synchronization (Optional)

*If the default version (3) doesn't result in NTP synchronization, try NTP version 2 (NTPv2).

The following example allows a router to synchronize its software clock with the peer at IP address 192.168.1.45 using the default settings over the FastEthernet 0 interface.

```
Rtr1(config)#ntp peer 192.168.1.45 source fastethernet 0
```

Use the interface configuration **ntp broadcast client** command to receive NTP broadcast packets on a specified interface. To disable this capability, use the no form of this command.

Rtr1(config-if)#ntp broadcast client

Rtr1(config-if)no ntp broadcast client

Simple Network Time Protocol (SNTP) SNTP is a simplified, client-only version of NTP supported on smaller Cisco routers, such as the 1000, 1600, and 1700 platforms. SNTP devices can only receive the time from NTP servers. They can't be used to provide time services to other systems.

While SNTP is simple to implement, it's more vulnerable to a misbehaving server than an NTP client and should be used only when strong authentication isn't required for the following reasons:

- Accuracy to within 100 milliseconds versus 2 milliseconds for NTP
- Doesn't provide complex filtering and statistical mechanisms of NTP
- Doesn't authenticate traffic, although extended ACLs can be configured

SNTP clients can be configured to request and accept packets from NTP servers or to accept NTP broadcast packets from any source.

Use the global configuration **sntp broadcast client** command to use the SNTP to accept NTP traffic from any broadcast server. To prevent the router from accepting broadcast traffic, use the no form of this command.

Rtr1(config)#sntp broadcast client

Rtr1(config)#no sntp broadcast client

The following example demonstrates using the **sntp broadcast client** command and the **show sntp** command to display the results.

```
Router(config)#sntp broadcast client
Router(config)#end
Router#show sntp
SNTP server    Stratum   Version   Last Receive
10.0.115.2        4         3        00:00:21   Synced   Bcast
Broadcast client mode is enabled.
```

The clock timezone Command To specify the router time zone, use the configuration mode **clock timezone** command. The command sets the time zone and an offset from UTC, displayed by the router. The command syntax is as follows:

Rtr1(config)#clock timezone *zone hour-offset* [*minute-offset*]

zone	Name of the time zone to be displayed (typically a standard acronym).
hour-offset	Hours offset from UTC.
minute-offset	Minutes offset from UTC. (Optional). This argument is for those cases where a local time zone is a percentage of an hour different from UTC/GMT.

The following example demonstrates setting the time zone to Central Standard Time (CST) in the United States. The -6 indicates CST is 6 hours behind UTC.

```
Rtr1(config)#clock timezone CST -6
```

The *zone* entry isn't a keyword. What you type is what you'll get, including case usage. Spaces aren't allowed. The following example demonstrates creating a custom entry. The **show clock** command displays the current settings. Notice the date and time setting is old, indicating the router hasn't had the software clock set. This is covered in the next section, "The clock set Command."

```
Rtr1(config)#clock timezone Pacific.Standard.Time -8
Rtr1(config)#^Z
Rtr1#show clock
*17:51:09.799 Pacific.Standard.Time Sun Feb 28 1993
Rtr1#
```

For more information on UTC, time zones, and common time zone abbreviations, go to http://www.timeanddate.com/worldclock.

The clock set Command If the system time is synchronized by a valid outside timing mechanism, such as a Network Time Protocol (NTP), or if it can be set using the router hardware clock, it shouldn't be necessary to set the software clock. Use this command if no other time sources are available.

To set the router's time and date manually, use the privileged EXEC **clock set** command. This command isn't available in Configuration mode. The time specified in this command is relative to the configured time zone. The command syntax is either of the following:

Rtr1#clock set *hh:mm:ss day month year*

Rtr1#clock set *hh:mm:ss month day year*

hh:mm:ss	Current time in hours, minutes, and seconds (24-hour format)
day	Current day of the month
month	Current month (by name)
year	Current year (4-digit format)

There are no defaults and all elements are required, even the seconds. The following example sets the time to 6:15 P.M. on December 15, 2004. While the full month can be typed out, it's only necessary to type the first three letters. The output is the same, regardless of whether the entry is abbreviated or the case of the entry.

```
Rtr1#clock set 18:15:00 15 december 2004
Rtr1#show clock
18:15:28.279 PST Wed Dec 15 2004
Rtr1#
```

Setting the Software Clock from the Hardware Clock If the router has a hardware-based clock, it's possible to set the software clock to the new hardware clock

setting by using the following **privilege EXEC mode** command. If the router doesn't have a hardware clock, the command is rejected.

```
Rtr1#clock read-calendar
```

Setting the Hardware Clock To update the hardware clock with the current software clock setting, use the following **privilege EXEC mode** command. If the router doesn't have a hardware clock, the command is rejected.

```
Rtr1#clock update-calendar
```

You can periodically update the hardware clock (calendar) from a NTP time source, using the global configuration **ntp update-calendar** command. To disable the periodic updates, use the no form of this command.

```
Rtr1(config)#ntp update-calendar
Rtr1(config)#no ntp update-calendar
```

Configuring Daylight Saving Time To configure daylight saving time (summer time) in areas where it starts and ends on a particular day of the week each year, use the following global configuration **mode** command. To avoid automatically switching to summer time, use the no form of this command.

> Rtr1(config)#clock summer-time *zone* recurring [*week day month hh:mm week day month hh:mm* [*offset*]]

zone	Time zone name, such as PDT for Pacific Daylight Time, to be displayed when summer time is in effect
recurring	Summer time starts and ends on the corresponding specified days every year
week	Week of the month, 1 to 5 or last (Optional)
day	Day of the week (Sunday, Monday, and so on) (Optional)
month	Month (January, February, and so on) (Optional)
offset	Number of minutes to add during summer time (default is 60) (Optional)

The first *week day month hh:mm* is when daylight saving time starts, while the last one represents when it ends.

The following example specifies that summer time starts on the first Sunday in April at 2 A.M. and ends on the last Sunday in October at 2 A.M.:

```
Rtr1(config)# clock summer-time PDT recurring 1 Sunday April 2:00 last Sunday
October 2:00
```

For those areas that have a varying daylight saving time, it's possible to configure the exact date and time of the next summer time event by using either of the following global configuration **mode** commands. To avoid automatically switching to summer time, use the no form of this command.

> Rtr1(config)#clock summer-time *zone* date *month date year hh:mm month date year hh:mm* [*offset*]

Rtr1(config)#no clock summer-time *zone* date *month date year hh:mm month date year hh:mm* [*offset*]

Rtr1(config)#clock summer-time *zone* date *date month year hh:mm date month year hh:mm* [*offset*]

Rtr1(config)#no clock summer-time *zone* date *date month year hh:mm date month year hh:mm* [*offset*]

date	Summer time starts on the first specific date listed in the command and ends on the second specific date in the command
date	Date of the month (1 to 31)
month	Text month (January, February, and so on) (Optional)
year	Four-digit year (1993 to 2035)
offset	Number of minutes to add during summer time (default is 60) (Optional)

In the following example, daylight saving time (summer time) for Pacific Daylight Time is configured to start on April 26, 1997 at 2 A.M. and end on October 12, 1997at 2 A.M. The **show clock** command with the Detail option displays the current settings:

```
Rtr1(config)#clock summer-time PDT date 26 Apr 2004 2:00 12 Oct 2004 2:00
Rtr1(config)#^Z
Rtr1#show clock detail
18:23:40.399 PST Thu Apr 15 2004
Time source is user configuration
Summer time starts 02:00:00 PST Mon Apr 26 2004
Summer time ends 02:00:00 PDT Tue Oct 12 2004
Rtr1#
```

Monitoring Time and Calendar Services To monitor clock, calendar, and NTP EXEC services, use the following commands in EXEC mode, as needed:

Rtr1#show calendar	Displays the current *hardware* clock time.
Rtr1#show clock [detail]	Displays the current *software* clock time.
Rtr1#show ntp associations [detail]	Displays the status of NTP associations.
Rtr1#show ntp status	Displays the status of NTP.

Step 2–3 Configure the Router Host Name and Domain Name

If the router's host name and domain name haven't been configured, it's necessary to configure them for CA support to work correctly. The host name is used in prompts and default configuration file names. The domain name is used to define a default domain name, which the Cisco IOS software uses to complete unqualified host names.

Use the global configuration **hostname** command to specify or modify the host name for the network server. The command syntax is as follows:

Rtr1(config)#hostname *host-name*

Use the global configuration **ip domain-name** command to define a default domain name that the Cisco IOS software uses to complete unqualified host names (names

without a dotted-decimal domain name). To disable the default domain name, use the no form of this command. The command syntax is as follows:

Rtr1(config)#ip domain-name *dom-name*

The following example shows configuring a host name and default domain name:

```
Router(config)#hostname Rtr1
Rtr1(config)#ip domain-name ca-test.com
```

Name Resolution Alternative In a system without a DNS server, using the global configuration **ip host** command to define a static host name-to-address mapping is possible. Each host name can be associated with up to eight IP addresses separated by spaces. To remove the name-to-address mapping, use the no form of the command. The command syntax is as follows:

Rtr1(config)#ip host *name address1* [*address2...address8*]

The following example shows creating an IP Host entry to resolve the CA server name because it can't be resolved using DNS features:

Rtr1(config)#ip host ca-server 192.168.1.100

Step 2–4 Generate a RSA Key Pair

RSA signatures support two mutually exclusive types of RSA key pairs.

- General-purpose keys
- Special-usage keys

When generating RSA key pairs, you can indicate whether to generate general-purpose keys or special-usage keys. If the usage-keys keyword isn't specified in the **crypto key generate rsa** command, the general-purpose keys are generated.

General Purpose Keys With general purpose keys, only one pair of RSA keys is generated. The generated pair is used for any IKE policies specifying either RSA signatures or RSA-encrypted nonces. Therefore, a general purpose key pair might get used more frequently than a special usage-key pair, increasing the chances of compromise.

Special-Usage Keys If the *usage-keys* option is specified, two separate pairs of RSA keys are generated for signing and encryption. One pair is used with any IKE policy that includes RSA signatures as the authentication method. The other pair is used with any IKE policy that specifies RSA encrypted nonces as the authentication method. RSA encrypted nonces are covered in the "RSA Encrypted Nonces Overview" section.

If the security policy specifies to have both types of RSA authentication methods in the IKE policies, it would be best to generate special-usage keys. By using the special-usage keys, each key isn't unnecessarily exposed. Otherwise, the single key is used for both authentication methods, increasing that key's exposure.

Modulus Length The key generation process can be both a lengthy and resource-intensive process, depending on the router platform and the length of the key modulus specified. When generating RSA keys, the system will prompt for the number of bits to use in the modulus key. The default is 512 bits, with a range of choices from 360 to 2,048 bits. A longer modulus provides stronger security, but takes longer to generate and use. A modulus of less than 512 bits is generally considered too vulnerable. Under some circumstances, a short modulus might not function properly with IKE policies. Cisco recommends using a minimum modulus of 1,024 bits.

The following table compares the processing time for generating general-purpose key pairs with various-length modulus values.

Platform	360 bits	512 bits	1,024 bits	2,048 bits
Cisco 2500	11 seconds	20 seconds	4 minutes, 38 seconds	more than 1 hour
Cisco 4700	less than 1 second	1 second	4 seconds	50 seconds

The crypto key generate rsa command The global configuration **crypto key generate rsa** command is used to generate RSA key pairs for a Cisco device, such as a router. If the router already has RSA keys, a warning is displayed, along with a prompt to replace the existing keys with new keys. The command syntax is as follows:

```
Rtr1(config)#crypto key generate rsa [usage-keys]
```

usage-keys	Specifies two RSA special-usage key pairs should be generated (that is, one encryption pair and one signature pair), instead of one general-purpose key pair. (Optional)

NOTE If the host name and the IP domain name weren't created in Step 2–3, the **crypto key generate rsa** command will fail.

This command isn't saved in the router configuration. The keys generated are saved in the private configuration in NVRAM, however, which is never displayed to the user or backed up to another device.

The following example shows generating a special-usage key pair followed by the **show crypto key mypubkey rsa** command demonstrating the result:

```
Rtr1(config)#crypto key generate rsa usage-keys
The name for the keys will be: Rtr1.ca-test.com
Choose the size of the key modulus in the range of 360 to 2048 for your
  Signature Keys. Choosing a key modulus greater than 512 may take
  a few minutes.

How many bits in the modulus [512]: 2048
Generating RSA keys ...
[OK]

Rtr1(config)#^Z
```

```
Rtr1#show crypto key mypubkey rsa
% Key pair was generated at: 00:21:00 PST Dec 16 2004
Key name: Rtr1.ca-test.com
 Usage: Signature Key
 Key Data:
  30820122 300D0609 2A864886 F70D0101 01050003 82010F00 3082010A 02820101
  009407CE 65AF10C8 A8D0023F 4CF157FE C359B0BF B21D86C8 9415D8EB B62878AB
  C559AD80 B6B5FD2F F2D4DEAD B6181350 9F937A17 BD23F2C0 CA0759DF 2D51F12C
  B60B2E9F 4F0ECD86 4B3C2B7E 5E30A4E7 D1C47D22 6ADEF189 BB99197E BAA741FF
  6014E0D5 426F37A8 6722A920 93DCB02F 6D5D6670 D3E4DAD1 7ECC6DF2 02D52A45
  53B45062 BF1685BC DF042887 79349774 82BC82E6 CA276477 65645CF6 E78E350E
  0C259D9F 366FD38F B78D16E2 F48839D5 DBA0D138 99CD26C1 3333C03D 1BBBD5B4
  7509E413 38D25619 EB5C4138 4B539EDB B0080E40 5FC179BA B1A5A5EF 0CD1CC2C
  BB6D018A D3BA4CD9 8F502553 7FEFB1D0 070177C5 DBAC2EBE 19D9E1A2 052F2B44
  A5020301 0001
% Key pair was generated at: 00:25:16 PST Dec 16 2004
Key name: Rtr1.ca-test.com
 Usage: Encryption Key
 Key Data:
  30820122 300D0609 2A864886 F70D0101 01050003 82010F00 3082010A 02820101
  00C50963 248DD7E2 E21191E4 1AD6A0D7 2D8EF5E5 8952B071 1FF8C18E C6E5F53B
  37DBB9E2 52A3BD1B 9C02444B EABC0E85 96C9866B 164AE96F CD27F448 C1C93028
  5B8A024D 7939E533 66D0CC37 6F512BF9 1432E6EC 791DC05D 4059A552 17EF7326
  664DB438 847461E4 8ECBBEB0 622E0CA3 1EA97ECA DF9B5C8F 65A0D887 23CADB33
  6F3D3BA4 C1962E66 C7753335 A41D63AE 1DB69086 6A63AF1C F008B7FD 94F41240
  087150CA 1422730E 3D06F070 09F4C630 09F6B3A0 8838E2F7 9CEDC701 DCBB7D6E
  0178D788 5235DBE0 033681C8 23CF2134 2DDC4151 CF86152B 391F456A 74135B61
  EB15D797 CECE8493 FEC16D9C 47BE7410 D0EF2353 61A3C108 08325630 8519D89D
  07020301 0001
% Key pair was generated at: 00:25:22 PST Dec 16 2004
Key name: Rtr1.ca-test.com.server
 Usage: Encryption Key
 Key Data:
  307C300D 06092A86 4886F70D 01010105 00036B00 30680261 00B7ED8C E4C06C08
  C8DE36F5 DDEEEB78 A0F1FA0E 4C8CAC32 2B4AC0A4 B813F56F FEA8B7F7 5C372C66
  55A1189E 0A373898 846B0679 497F9474 29EB1088 E31ED19E 771C10C6 C35A98D0
  226AEDB3 FE42066A A81FD230 4FFCBCD0 0159D879 6F632FFA 1F020301 0001
Rtr1#
```

Step 2–5 Declare a CA

Use the global configuration **crypto ca identity** command to specify what CA the router will use. The CA might require a particular name, such as its domain name. Use the no form of this command to delete all identity information and certificates associated with the CA. The command syntax is as follows:

Rtr1(config)#crypto ca identity *name*

Rtr1(config)#no crypto ca identity *name*

If the CA has been previously declared and updating its characteristics is all that's required, just specify the name previously created. Performing the **crypto ca identity** command initiates the ca-identity Configuration mode, where CA characteristics can be defined, as shown in the following output:

```
Rtr1(config)#crypto ca identity ca-ipsec
Rtr1(ca-identity)#?
```

```
CA identity configuration commands:
  crl         CRL option
  default     Set a command to its defaults
  enrollment  Enrollment parameters
  exit        Exit from certificate authority identity entry mode
  no          Negate a command or set its defaults
  query       Query parameters

Rtr1(ca-identity)#
```

Some of the more common CA characteristics that can be defined in ca-identity Configuration mode include the following. Use the no form of the commands to remove the feature.

enrollment url	URL of the CA—required
enrollment mode ra	RA mode—required if the CA system supports a registration authority (RA).
query url	URL of the Lightweight Directory Access Protocol (LDAP) server—required if the CA supports an RA and the LDAP protocol.
enrollment retry period	Minutes the router will wait between sending certificate request retries. Range 0 to 60 minutes. Default is one minute. (Optional)
enrollment retry count	Number of certificate request retries the router will send before giving up. Range 1 to 100. Default is 0, which allows an infinite number of retries. (Optional)
crl optional	Specify whether the router can accept other peers' certificates if the certificate revocation list isn't accessible. Without this command, CRL checking is mandatory. (Optional)

The following is an example of the identity options that might be used with Verisign as the CA service provider. The choices and options vary with provider.

```
Rtr1(config)#crypto ca identity ca-ipsec
Rtr1(ca-identity)#enrollment url http://ca-ipsec.verisign.com/
Rtr1(ca-identity)#query url ldap://ca-ipsec.verisign.com/
Rtr1(ca-identity)#crl optional
Rtr1(ca-identity)#enrollment mode ra
Rtr1(ca-identity)#enrollment retry count 10
Rtr1(ca-identity)#enrollment retry period 5
Rtr1(ca-identity)#exit
Rtr1(config)#
```

The crypto ca identity is only significant locally. It needn't match the identity defined on any VPN peer.

Step 2–6 Authenticate the CA

The router needs to authenticate the validity of the CA by obtaining its self-signed certificate containing the CA's public key. Because the CA's certificate is self-signed, its public key should be manually authenticated by contacting the CA administrator.

If the CA support is using RA mode, the **enrollment mode ra** command in Step 2–5, then RA signing and encryption certificates will accompany the CA certificate.

To authenticate the certification authority (and RA, if appropriate), use the global configuration **crypto ca authenticate** command. Use the same name used to declare the CA with the **crypto ca identity** command:

Rtr1(config)#crypto ca authenticate *name*

If the CA doesn't respond within the timeout period defined in the **ca identity** commands, the command will fail. If this happens, you must reenter the command. The following show a representation of a successful CA authentication:

```
Rtr1(config)#crypto ca authenticate ca-ipsec
Certificate has the following attributes:
Fingerprint: 9876 5432 1098 7654 3210
Do you accept this certificate? [yes/no]
```

While the command isn't saved to the router configuration, the public keys embedded in the received CA (and RA) certificates are saved to the configuration as part of the RSA public key record called the "RSA public key chain."

Step 2–7 Request Your Own Certificate

The next step requires the local router to obtain its certificate from the CA using the global configuration **crypto ca enroll** command. Use the same name as when you declared the CA using the **crypto ca identity** command. Use the no form of this command to delete a current enrollment request:

Rtr1(config)#crypto ca enroll *name*

Rtr1(config)#no crypto ca enroll *name*

The router needs a signed certificate from the CA for each of the router's RSA key pairs. With a RSA general purpose key request in Step 2–4, this command will obtain one certificate corresponding to the general purpose RSA key pair. If the RSA key request was for special-usage keys, this command will return two certificates corresponding to each of the special-usage key pairs. This task is also known as *enrolling* with the CA.

If certificates already exist for each key pair, the command will fail and a prompt will appear, advising the existing certificates be removed first. Removing the existing certificates is accomplished with the **no certificate** command.

Responding to Prompts During the enrolling process with the CA, several prompts appear. The actual prompts can vary with the CA provider and over time. The following example is included as representative of what you can expect.

First, you must create and verify a challenge password. This *challenge password* is important because it's required if you ever need to revoke the router's certificate(s). The password can be up to 80 characters in length. This password isn't stored anywhere by the router, so you must remember it or secure it in a safe place. A lost password might not absolutely prevent revocation, but the CA administrator will undoubtedly require additional authentication of the router administrator identity.

Second, a prompt will ask for a subject name in the certificate. In the following example, the fully qualified domain name of the local device was used.

Third, a prompt will ask whether to include the router's serial number in the certificate. This is an internal board serial number rather than the one on the case. While the serial number isn't used by IPSec or the IKE process, it can be used by the CA to authenticate certificates or to associate a certificate with that particular router.

Fourth, a prompt will ask whether to include the router's IP address. This might not be a good idea in some cases because the IP address binds the certificate more tightly to a specific entity. If the router is moved, requiring a new IP address, then requesting a new certificate would also be necessary. If including the IP address is selected, a prompt will ask for the interface of the IP address. The interface should correspond to the interface to which the crypto map is applied.

The following example shows an RSA general-purpose key pair request for a certificate from the CA. Ultimately, the router displays the certificate fingerprint, at which time the local administrator verifies the fingerprint by calling the CA administrator. If the fingerprint is correct, the router administrator accepts the certificate.

A delay might occur between when the certificate request is sent and when the certificate is received by the router. The length of any delay depends on the CA authorization process.

```
Rtr1(config)#crypto ca enroll ca-ipsec
%
% Start certificate enrollment ..
% Create a challenge password. You will need to verbally provide this
password to the CA Administrator in order to revoke your certificate.
For security reasons your password will not be saved in the configuration.
Please make a note of it.
Password: ca-password
Re-enter password: ca-password
% The subject name in the certificate will be: rtr1.ca-test.com
% Include the router serial number in the subject name? [yes/no]: yes
% The serial number in the certificate will be: 01234567
% Include an IP address in the subject name [yes/no]? no
Interface: fastethernet 0
Request certificate from CA [yes/no]? yes
% Certificate request sent to Certificate Authority
% The certificate request fingerprint will be displayed.
% The 'show crypto ca certificate' command will also show the fingerprint.
Rtr1(config)#
```

At a later time, the router receives the certificate from the CA and displays a confirmation message, such as the following:

```
Rtr1(config)#Fingerprint: 1234 9876 5432 1098 7654
%CRYPTO-6-CERTRET: Certificate received from Certificate Authority
Rtr1(config)#
```

If the certificate request can't be granted for any reason, a message like the following is displayed at the router instead:

```
%CRYPTO-6-CERTREJ: Certificate enrollment request was rejected by
  Certificate Authority
```

Step 2–8 Save the Configuration

After completing the router configuration for CA support, the configuration should be saved using the **copy running-config startup-config** command.

Step 2–9 Monitor and Maintain CA Interoperability (Optional)

The following housekeeping measures are optional and will vary, based on the CA server requirements and operational circumstances:

- Request a CRL
- Delete your router's RSA keys
- Delete both public and private certificates from the configuration
- Delete peer's public keys

Request a CRL When the router receives a certificate from a peer, it will download a CRL from either the CA or from a CA-designated CRL distribution point. The router then looks for the certificate on the CRL to make sure it hasn't been revoked. The router won't accept the certificate and it won't authenticate the peer if the certificate appears on the CRL. A CRL can continue to be reused with other certificates until the CRL expires. If the router receives a certificate after the CRL has expired, it will download the latest CRL.

If the CA system supports RAs, multiple CRLs can exist. The certificate in question will indicate which CRL applies and should be downloaded by the router for authentication. If the router doesn't have or is unable to download the appropriate CRL, the certificate will normally be rejected. The exception to this would be if the router configuration contains the **crl optional** feature under the **crypto ca identity** command. With the **crl optional** command, the router will still try to get the required CRL, but failing that, it can accept the peer's certificate anyway. The router will continue to try to get the CRL.

If the CA server requires a local CRL, use the global configuration **crypto ca crl request** command to request an immediate download of the latest CRL. Use the same name used to declare the CA with the **crypto ca identity** command:

Rtr1(config)#crypto ca crl request *name*

An example for the chapter scenario would look like the following:

```
Rtr1(config)#crypto ca crl request ca-ipsec
```

Delete Router's RSA Keys While this could seem like a harsh step, it might be necessary to delete a router's RSA keys if you had reason to believe the RSA keys were compromised. This process can't be undone and all new key exchanges would have to wait until new RSA key requests could be processed.

Use the global configuration **crypto key zeroize rsa** command to delete all the router's RSA keys. The syntax and an example follow:

```
Rtr1(config)#crypto key zeroize rsa
```

Delete Both Public and Private Certificates from the Configuration The **crypto key zeroize rsa** command in the last step deletes all RSA keys generated by the router. In addition to this command, it would be necessary to perform two additional tasks:

- Request the CA administrator revoke the router's certificates at the CA. The challenge password created when the router's certificates were originally enrolled must be supplied.
- The router's certificates must be manually removed from the configuration using the **certificate** command in Certificate Chain Configuration mode. The syntax would be as follows:
  Rtr1(config-cert-chain)#no certificate certificate-serial-number

In the following example, the **show** command is used to determine the serial number(s) of the certificate to be deleted:

```
Rtr1#show crypto ca certificates
Certificate
  Subject Name
    Name: rtr1.ca-test.com
    IP Address: 10.0.1.21
  Status: Available
  Certificate Serial Number: 9876543210ABCDEF9876543210ABCDEF
  Key Usage: General Purpose

CA Certificate
  Status: Available
  Certificate Serial Number: 0192837465A0B1C2D3E4F50192837465A
  Key Usage: Not Set
Rtr1#config t
Rtr1(config)#crypto ca certificate chain ca-ipsec
Rtr1(config-cert-chain)#no certificate 0192837465A0B1C2D3E4F50192837465A
% Are you sure you want to remove the certificate [yes/no]? yes
% Be sure to ask the CA administrator to revoke this certificate.
Rtr1(config-cert-chain)#exit
Rtr1(config)#
```

To delete the CA's certificate, you must remove the entire CA identity, which also removes all certificates associated with the CA—the router's certificate, the CA certificate, and any RA certificates. Use the global configuration **no crypto ca identity** command to remove the CA identity. The syntax is

Rtr1(config)#no crypto ca identity *name*

The following example demonstrates using the **no crypto ca identity** command to remove the CA identity:

```
Rtr1(config)#no crypto ca identity ca-ipsec
% Removing an identity will destroy all certificates received from
```

```
 the related Certificate Authority.

Are you sure you want to do this? [yes/no]: y
% Be sure to ask the CA administrator to revoke your certificates.
Rtr1(config)#
```

Delete Peer's Public Keys If a peer's RSA public keys have been compromised, it could be necessary to delete the peer's public keys. Use the following commands to accomplish this task:

Rtr1(config)#crypto key pubkey-chain rsa

Rtr1(config-pubkey-chain)#no named-key key-name [encryption | signature]

or

Rtr1(config-pubkey-chain)#no addressed-key key-address [encryption | signature]

The global configuration **crypto key pubkey-chain rsa** command enters the Public Key Chain Configuration mode.

The choice of *named-key* or *addressed-key* depends on which keys were created originally. If only general purpose RSA keys were generated, don't use either of the encryption or signature keywords. If the IPSec remote peer generated special-usage keys resulting in two pairs of keys, then you must perform the appropriate command twice and use the encryption and signature keywords in turn. Similarly, using the encryption or signature keywords makes it possible to delete only one type of key.

Step 2–10 Verify the CA Support Configuration

The following commands can be used to verify any configured CA certificates:

- show crypto ca certificates
- show crypto key mypubkey rsa
- show crypto key pubkey-chain rsa

The show crypto ca certificates Command The **show crypto ca certificates** command is used to display the details of any CA certificates. The syntax reveals other interesting options:

Rtr1#show crypto ca [certificates | crls | roots]

```
Rtr1#show crypto ca certificates
Certificate
  Subject Name
    Name: rtr1.ca-test.com
    IP Address: 10.0.1.21
  Status: Available
  Certificate Serial Number: 9876543210ABCDEF9876543210ABCDEF
  Key Usage: General Purpose

CA Certificate
  Status: Available
```

```
  Certificate Serial Number: 0192837465A0B1C2D3E4F50192837465A
  Key Usage: Not Set
Rtr1#
```

The show crypto key mypubkey rsa Command Use the **show crypto key mypubkey rsa** EXEC command to view the RSA public keys associated with this router. The following output from the **show crypto key mypubkey rsa** command shows the results of RSA keys generated, using the **crypto key generate rsa usage-keys** command with a 2,048-bit modulus.

```
Rtr1#show crypto key mypubkey rsa
% Key pair was generated at: 00:21:00 PST Dec 16 2004
Key name: Rtr1.ca-test.com
 Usage: Signature Key
 Key Data:
  30820122 300D0609 2A864886 F70D0101 01050003 82010F00 3082010A 02820101
  009407CE 65AF10C8 A8D0023F 4CF157FE C359B0BF B21D86C8 9415D8EB B62878AB
  C559AD80 B6B5FD2F F2D4DEAD B6181350 9F937A17 BD23F2C0 CA0759DF 2D51F12C
  B60B2E9F 4F0ECD86 4B3C2B7E 5E30A4E7 D1C47D22 6ADEF189 BB99197E BAA741FF
  6014E0D5 426F37A8 6722A920 93DCB02F 6D5D6670 D3E4DAD1 7ECC6DF2 02D52A45
  53B45062 BF1685BC DF042887 79349774 82BC82E6 CA276477 65645CF6 E78E350E
  0C259D9F 366FD38F B78D16E2 F48839D5 DBA0D138 99CD26C1 3333C03D 1BBBD5B4
  7509E413 38D25619 EB5C4138 4B539EDB B0080E40 5FC179BA B1A5A5EF 0CD1CC2C
  BB6D018A D3BA4CD9 8F502553 7FEFB1D0 070177C5 DBAC2EBE 19D9E1A2 052F2B44
  A5020301 0001
% Key pair was generated at: 00:25:16 PST Dec 16 2004
Key name: Rtr1.ca-test.com
 Usage: Encryption Key
 Key Data:
  30820122 300D0609 2A864886 F70D0101 01050003 82010F00 3082010A 02820101
  00C50963 248DD7E2 E21191E4 1AD6A0D7 2D8EF5E5 8952B071 1FF8C18E C6E5F53B
  37DBB9E2 52A3BD1B 9C02444B EABC0E85 96C9866B 164AE96F CD27F448 C1C93028
  5B8A024D 7939E533 66D0CC37 6F512BF9 1432E6EC 791DC05D 4059A552 17EF7326
  664DB438 847461E4 8ECBBEB0 622E0CA3 1EA97ECA DF9B5C8F 65A0D887 23CADB33
  6F3D3BA4 C1962E66 C7753335 A41D63AE 1DB69086 6A63AF1C F008B7FD 94F41240
  087150CA 1422730E 3D06F070 09F4C630 09F6B3A0 8838E2F7 9CEDC701 DCBB7D6E
  0178D788 5235DBE0 033681C8 23CF2134 2DDC4151 CF86152B 391F456A 74135B61
  EB15D797 CECE8493 FEC16D9C 47BE7410 D0EF2353 61A3C108 08325630 8519D89D
  07020301 0001
% Key pair was generated at: 00:25:22 PST Dec 16 2004
Key name: Rtr1.ca-test.com.server
 Usage: Encryption Key
 Key Data:
  307C300D 06092A86 4886F70D 01010105 00036B00 30680261 00B7ED8C E4C06C08
  C8DE36F5 DDEEEB78 A0F1FA0E 4C8CAC32 2B4AC0A4 B813F56F FEA8B7F7 5C372C66
  55A1189E 0A373898 846B0679 497F9474 29EB1088 E31ED19E 771C10C6 C35A98D0
  226AEDB3 FE42066A A81FD230 4FFCBCD0 0159D879 6F632FFA 1F020301 0001
Rtr1#
```

The show crypto key pubkey-chain rsa Command Use the **show crypto key pubkey-chain rsa EXEC** command to view any peer RSA public keys stored on the router:

Rtr1#show crypto key pubkey-chain rsa [name *key-name* | address *key-address*]

name *key-name*	Name of the particular public key to view (Optional)
address *key-address*	Address of the particular public key to view (Optional)

The following is sample output from the **show crypto key pubkey-chain rsa** command:

```
Rtr1#show crypto key pubkey-chain rsa
Codes: M - Manually configured, C - Extracted from certificate
Code    Usage         IP-Address     Name
M       Signature     10.0.1.21      rtr1.ca-test.com
M       Encryption    10.0.1.21      rtr1.ca-test.com
C       Signature     10.0.50.2      rtr2.ca-test.com
C       Encryption    10.0.50.2      rtr2.ca-test.com
```

Task 3—Configure IKE

Task 3 is configuring IKE parameters determined in Task 2. The steps required are identical to configuring preshared keys in Chapter 10, except for setting the authentication method parameter in Step 2. Configuring IKE consists of the following steps and commands:

- Step 3–1 Enable or disable IKE
  crypto isakmp enable
- Step 3–2 Create IKE policies
  crypto isakmp policy
  authentication
  encryption
  hash
  group
  lifetime
- Step 3–3 Set the IKE identity to address or host name
  crypto isakmp identity
- Step 3–4 Verify the IKE configuration
  show crypto isakmp policy
  show crypto isakmp sa

Step 3–2 Create IKE Policies

Now, it's time to define a suite of ISAKMP (IKE) policy details gathered during the planning task, Task 1 Step 2 (repeated in the following with the defaults in italics). The goal is to create a suite of IKE policies, so it's possible to establish IKE peering between two IPSec endpoints.

Parameter	Preferred (stronger)	2nd Choice	3rd Choice
Encryption algorithm	3des	*des*	*des*
Hash algorithm	*sha*	*sha*	md5
Authentication method	*rsa-sig*	*rsa-sig*	*rsa-sig*
DH key exchange group	2	2	*1*
IKE SA lifetime	43,200	43,200	*86,400*

Using the scenario example, Rtr2 only needing to connect to Rtr1, might be configured with only the preferred (most secure) set of parameters because that set will also be on Rtr1. On the other hand, Rtr1, needing to peer with various devices both domestically and internationally, might need multiple policies to accommodate the different peers. The policies are configured using priority numbers, so the preferred policy is evaluated by the peer first, and then by the second and third choices. This ensures the agreed policy will always be the highest mutually agreeable choice.

The crypto isakmp policy Command Use the global **crypto isakmp policy** command to define an IKE policy containing a set of parameters used during IKE negotiation. Use the no form of this command to delete an IKE policy. The command syntax and example are as follows:

Rtr1(config)#[no] crypto isakmp policy *priority*

Rtr1(config)#crypto isakmp policy 100

Rtr1(config-isakmp)#

priority	Unique identifier for the IKE policy used to order, or prioritize, the policies. Acceptable range 1 to 10,000, with 1 being the highest possible priority.

This command invokes the ISAKMP Policy Configuration (config-isakmp) mode, which allows the six IKE parameters to be defined. The no form of each parameter removes that setting. Exit is used to end the session and any undefined parameters will use the default values. The actual choices are based on the table made in Task 1 Step 2. The following shows the syntax for each parameter, followed by the entries for the preferred policy.

```
Rtr1#conf t
Rtr1(config)#crypto isakmp policy 100
Rtr1(config-isakmp)#authentication {pre-share | rsa-encr | rsa-sig}
Rtr1(config-isakmp)#encryption {des | 3des}  (depending on feature set)
Rtr1(config-isakmp)#hash {md5 | sha}
Rtr1(config-isakmp)#group {1 | 2}
Rtr1(config-isakmp)#lifetime <60 to 86400 seconds>

Rtr1(config)#crypto isakmp policy 100
Rtr1(config-isakmp)#authentication rsa-sig
Rtr1(config-isakmp)#encryption 3des
Rtr1(config-isakmp)#hash sha
Rtr1(config-isakmp)#group 2
Rtr1(config-isakmp)#lifetime 43200
Rtr1(config-isakmp)#exit
```

Because the preferred policy was assigned 100 as a priority, the second and third choices must be progressively larger to control the order in which they're offered to a peer. A possible choice might be 200 and 300, respectively, allowing newer policies to be assigned priorities that would allow them to be prioritized according to changes in the company security policy.

The following shows the steps to configuring the second and third preference, taking advantage of default values to save configuration:

```
Rtr1#conf t
Rtr1(config)#crypto isakmp policy 200
Rtr1(config-isakmp)#group 2
Rtr1(config-isakmp)#lifetime 43200
Rtr1(config-isakmp)#crypto isakmp policy 300
Rtr1(config-isakmp)#hash md5
Rtr1(config-isakmp)#exit
```

Step 3–3 Set the IKE Identity to Address or Host Name

Use the global **crypto isakmp identity** command to set the ISAKMP identity of the remote peer to be used during IKE negotiations. The IKE identity can either be the device IP address or the host name. The router IOS defaults to the IP address and doesn't add a line to the configuration, unless the host name option is used.

To use the host name identity method, use the global **crypto isakmp identity** command. The no form of this command resets the default value (address). The command syntax is as follows:

Rtr1(config)#[no] crypto isakmp identity {address | hostname}

Rtr1(config)#crypto isakmp identity hostname

Using the host name identity method makes the most sense if a DNS server is available on the network to resolve the name. If not, the following example shows using an **IP Host** command entry on the router to handle the name resolution:

```
Rtr1(config)#ip host Rtr2.domain.com 10.0.50.2
```

Task 4—Configure IPSec

Task 4 is configuring Cisco IOS IPSec to the IPSec parameters developed in Task 1. The steps are identical to those used to configure preshared keys in Chapter 10. The steps and commands used to configure IPSec encryption on Cisco routers are summarized as follows:

- Step 4–1 Configure transform set suites
  crypto ipsec transform-set
- Step 4–2 Configure global IPSec security association lifetimes
  crypto ipsec security-association lifetime
- Step 4–3 Configure crypto ACLs
  access-list
- Step 4–4 Configure crypto maps
  crypto map
- Step 4–5 Apply the crypto maps to the interface
  interface
  crypto map

Task 5—Test and Verify IPSec

Cisco IOS software contains a number of **show, clear,** and **debug** commands useful for testing and verifying IPSec and ISAKMP. While many of these commands are used the same as when configuring preshared keys, Step 5–7 introduces two debug options that are unique to RSA signatures. To test and verify that Cisco IOS IPSec VPN is correctly configured, use the following commands:

- Step 5–1 Display the configured IKE policies
  show crypto isakmp policy
- Step 5–2 Display the configured transform sets
  show crypto ipsec transform set
- Step 5–3 Display the current state of the IPSec SAs
  show crypto ipsec sa
- Step 5–4 Display the configured crypto maps
  show crypto map
- Step 5–5 Debug IKE events
  debug crypto isakmp
- Step 5–6 Debug IPSec events
  debug crypto ipsec
- Step 5–7 Debug CA events
  debug crypto key-exchange
  debug crypto pki commands

RSA Encrypted Nonces Overview

The *RSA-encrypted nonces authentication method* uses the RSA-encryption public key cryptography algorithm. This technology requires each party to generate a pseudorandom number (a nonce) and encrypt it (and possibly other publicly and privately available information), using the other party's RSA public key. Authentication occurs when each party decrypts the other party's nonce with their local private key, and then uses the decrypted nonce to compute a keyed hash.

The major drawback to implementing this technology is it's somewhat difficult to configure and, therefore, more difficult to scale to a large number of VPN peers. RSA-encrypted nonces require peers to possess each other's public keys, but they don't use a CA. Two methods can be used for peers to get each others' public keys:

- Manually configure and exchange RSA keys
- Use the RSA signatures used previously during a successful ISAKMP negotiation with the remote peer

Another potential drawback to this authentication method is this: either side of the exchange can plausibly deny they took part in the exchange. Cisco IOS software is

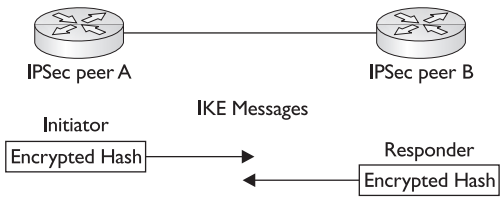

Figure 11-6 RSA-encryption authentication exchange

the only Cisco product that supports this authentication method. Figure 11-6 shows a RSA-encryption authentication exchange.

This section provides a short overview of configuring IPSec using RSA-encrypted nonces. Only those tasks and steps that are unique to RSA-encrypted nonces are presented. Configuring RSA encryption is similar to preshared keys and CA support. It's introduced using the same outline with commands introduced when those technologies were covered in Chapter 10 and earlier in this chapter. The following are the major tasks for configuring RSA-encrypted nonces:

- Task 1 Prepare for IKE and IPSec
- Task 2 Configure RSA keys manually
- Task 3 Configure IKE for IPSec to select RSA encryption
- Task 4 Configure IPSec (typically, the same as preshare keys)
- Task 5 Test and verify IPSec

Task 2—Configure RSA Keys

The steps and commands used in Task 2 are included in the following items. While this display is intended to demonstrate the similarities to technologies already covered, a thorough coverage of the tasks and steps to configure RSA-encrypted nonces can be found in the Cisco IOS Security Configuration Guide online.

Configuring RSA keys involves the following six steps:

- Step 2–1 Plan for RSA keys.
- Step 2–2 Configure the router's host name and domain name.

  hostname

  ip domain-name
- Step 2–3 Generate the RSA keys.

  crypto key generate rsa usage-keys
- Step 2–4 Enter peer RSA public keys.

  crypto key pubkey-chain rsa

  addressed-key ***ip-addr***

  named-key ***name***

- Step 2–5 Verify the key configuration.
  show crypto key mypubkey rsa
  show crypto key pubkey-chain rsa
- Step 2–6 Manage RSA keys—Remove old keys.
  crypto key zeroize rsa

Chapter Review

This chapter looked at steps involved in configuring IPSec for CA support. The steps and related commands are summarized in the following task list:

Task 1 Prepare for IKE and IPSec

- Step 1–1 Plan for CA support
- Step 1–2 Determine the IKE (IKE phase one) policies
- Step 1–3 Determine the IPSec (IKE phase two) policies
- Step 1–4 Check the current configuration
  show running-configuration
  show isakmp [policy]
  show crypto map
- Step 1–5 Ensure the network works without encryption
  ping
- Step 1–6 Ensure access control lists are compatible with IPSec
  show access-lists

Task 2 Configure CA support

- Step 2–1 Manage the NVRAM memory usage (Optional)
  crypto ca certificate query
- Step 2–2 Set the router's time and date
  ntp broadcast client
  sntp broadcast client
  clock set
- Step 2–3 Configure the router's host name and domain name
  hostname
  ip domain-name
  ip host
- Step 2–4 Generate an RSA key pair
  crypto key generate rsa

- Step 2–5 Declare a CA
  crypto ca identity
  enrollment url
  query url
  crl optional
  enrollment mode ra
  enrollment retry count
  enrollment retry period
- Step 2–6 Authenticate the CA
  crypto ca authenticate
- Step 2–7 Request your own certificate
  crypto ca enroll
- Step 2–8 Save the configuration
  copy running-config startup-config
- Step 2–9 Monitor and maintain CA interoperability (Optional)
  Request a CRL
  crypto ca crl request
  Delete your router's RSA keys:
  crypto key zeroize rsa
  Delete both public and private certificates from the configuration:
  no certificate certificate-serial-number
  no crypto ca identity
  Delete peer's public keys:
  no named-key key-name
  no addressed-key key-address
- Step 2–10 Verify the CA support configuration
  show crypto ca certificates
  show crypto key mypubkey rsa
  show crypto key pubkey-chain rsa

Task 3 Configure IKE

- Step 3–1 Enable or disable IKE
  crypto isakmp enable
- Step 3–2 Create IKE policies
  crypto isakmp policy

authentication
encryption
hash
group
lifetime

- Step 3–3 Configure preshared keys
  crypto isakmp key
- Step 3–4 Verify the IKE configuration
  show crypto isakmp policy

Task 4 Configure IPSec

- Step 4–1 Configure transform set suites
  crypto ipsec transform-set
- Step 4–2 Configure global IPSec security association lifetimes
  crypto ipsec security-association lifetime
- Step 4–3 Configure crypto ACLs
  access-list
- Step 4–4 Configure crypto maps
  crypto map
- Step 4–5 Apply the crypto maps to the interface
  interface
  crypto map

Task 5 Test and verify IPSec

- Step 5–1 Display the configured IKE policies
  show crypto isakmp policy
- Step 5–2 Display the configured transform sets
  show crypto ipsec transform set
- Step 5–3 Display the current state of the IPSec SAs
  show crypto ipsec sa
- Step 5–4 Display the configured crypto maps
  show crypto map
- Step 5–5 Debug IKE events
  debug crypto isakmp
- Step 5–6 Debug IPSec events
  debug crypto ipsec

Questions

1. A digital certificate is conceptually most like which type of document?
   - **A.** Event admission ticket
   - **B.** Vehicle license plate
   - **C.** Passport
   - **D.** Social Security card
2. Which of the following is *not* a common name for a database service running on an existing or dedicated server that allows users to submit and retrieve digital certificates?
   - **A.** Certificate server
   - **B.** CRL
   - **C.** Cert server
   - **D.** Key server
3. What does the acronym PKI stand for?
   - **A.** PIX Key Interchange
   - **B.** Private Key Interchange
   - **C.** Public Key Infrastructures
   - **D.** PIX Key Interface
4. Digital certificates are generated by which of the following?
   - **A.** Sending peer
   - **B.** Certificate authority
   - **C.** Receiving peer
   - **D.** The government
5. When checking a certificate against a CRL, what happens if a match occurs?
   - **A.** The certificate is accepted
   - **B.** A new CRL is requested
   - **C.** The certificate is rejected
   - **D.** The request is sent to the CA
6. Which of the following is a server that acts as a proxy for the CA, so CA functions can continue when the CA is offline or otherwise unavailable?
   - **A.** CRL
   - **B.** CAR
   - **C.** CA
   - **D.** RA

7. Which of the following is an initiative for furthering open development for certificate-handling protocols that can help ensure interoperability with devices from many vendors?
   - **A.** PKI
   - **B.** CA
   - **C.** LDAP
   - **D.** SCEP
8. Which of the following is *not* a CA provider supported by the Cisco IOS?
   - **A.** Entrust Technologies, Inc.
   - **B.** Symantic
   - **C.** VeriSign
   - **D.** Microsoft
9. Which is the IKE keyword for CA support authentication method?
   - **A.** rsa-sig
   - **B.** pki
   - **C.** rsa-encr
   - **D.** preshare
10. Which command specifies that certificates and CRLs should *not* be stored locally, but should be retrieved from the CA as needed?
    - **A. no ntp peer ip-address**
    - **B. crypto key generate rsa**
    - **C. crypto ca identity**
    - **D. crypto ca certificate query**
11. In the following command, what does the word "six" represent?

```
Rtr1(config)#clock timezone CST -6
```

    - **A.** The number six is a sequence number
    - **B.** Six hours behind NY standard time
    - **C.** Six hours behind UTC/GMT
    - **D.** Six hours ahead of UTC/GMT
12. Given the following command, how many RSA key pairs will be generated?

```
Rtr1(config)#crypto key generate rsa usage-keys
```

    - **A.** 1
    - **B.** 2
    - **C.** 3
    - **D.** 4

13. Which command is used to define the CA?
    A. **crypto ca enroll**
    B. cryp**to ca identity**
    C. **crypto ca authenticate**
    D. **crypto key zeroize rsa**
14. Which command removes all certificates associated with the CA—the router's certificate, the CA certificate, and any RA certificates?
    A. **no named-key key-name**
    B. **no crypto ca identity**
    C. **crypto key zeroize rsa**
    D. **no certificate**
15. Which of the following is *not* required for CA support on Cisco IOS devices?
    A. Hostname defined
    B. Special-usage keys ordered
    C. Domain name defined
    D. Software clock set

Answers

1. **C.** Passport
2. **B.** CRL
3. **C.** Public Key Infrastructures
4. **B.** Certificate authority
5. **C.** The certificate is rejected—it has been revoked
6. **D.** RA
7. **D.** SCEP—Simple Certificate Enrollment Protocol
8. **B.** Symantic
9. **A.** rsa-sig
10. **D.** **crypto ca certificate query**
11. **C.** Six hours behind UTC/GMT
12. **B.** 2
13. **B.** **crypto ca identity**
14. **B.** **no crypto ca identity**
15. **B.** Special-usage keys ordered

CHAPTER 12

Cisco IOS Remote Access Using Cisco Easy VPN

In this chapter, you will learn to:

- Understand the Cisco Easy VPN
- Use the Cisco Easy VPN server
- Use the Cisco Easy VPN Remote
- Make use of the Cisco VPN 3.6 Client
- Do easy VPN server configuration tasks
- Preconfigure the Cisco VPN 3.6 client
- Understand the new Management Center for VPN routers
- Use the easy VPN Remote Phase Two
- Know about the Cisco VPN Firewall feature for VPN Client

This chapter looks at the Cisco VPN Client software as part of the Cisco Easy VPN strategy to deliver easy to use, reliable VPN connectivity. Many features require supporting configuration in an IOS router, a PIX Firewall, or a Cisco VPN Concentrator. That specific configuration was addressed in Chapters 10 and 11.

Introduction to Cisco Easy VPN

Cisco Easy VPN is a component of the Cisco Unified Client Framework, in which VPN management is centralized across all Cisco VPN devices. This strategy simplifies VPN deployment for remote offices and telecommuters, reduces deployment complexity, and conserves human resources. The Cisco Easy VPN strategy incorporates all VPN platforms, including VPN software client, VPN 3002 hardware client, Cisco routers, Cisco PIX Firewalls, and Cisco VPN 3000 Concentrators.

The Cisco Easy VPN strategy is comprised of two types of implementations:

- Cisco Easy VPN Server
- Cisco Easy VPN Remote

Cisco Easy VPN Server

The *Cisco Easy VPN Server service* allows a growing number of Cisco IOS routers, PIX Firewalls, and Cisco VPN 3000 Concentrators to act as VPN head-end devices in site-to-site or remote access VPNs. These head-end devices provide the VPN connections, as well as a configuration source for the Cisco Easy VPN Remote sites. The Cisco Easy VPN Server is available on numerous Cisco VPN routers, including Cisco UBR900, 1700, 2600, 3600, 7100, 7200, and 7500 series routers running Cisco IOS Release 12.2(8)T, Cisco PIX Firewalls, and all Cisco VPN 3000 Concentrators. Use the Cisco online documentation to verify any new platforms supported because new supported platforms have been added with each new release.

Basic feature configuration and security policies defined on the head-end device are pushed to the remote VPN site before the connection is established. This insures those connections have current configurations and policies in place. In the case of the VPN 3002 Hardware Device, firmware updates can be maintained using the same feature.

Cisco Easy VPN Server-enabled devices can provide VPN tunnel termination for mobile remote workers running Cisco VPN client software, allowing them to access the headquarters' intranet.

Some of the benefits of the Cisco Easy VPN Server implementation for rolling out remote VPN connections include the following:

- Centrally stored configurations allow dynamic configuration of end-user VPN connections and require less manual configuration by end users and field technicians, reducing errors and additional service calls.
- VPN security policy management can be centralized.
- Supports large-scale, rapid deployments with minimal remote user provisioning.
- When used with a remote VPN device, such as a router, firewall, or VPN hardware client, this eliminates the need for end-user VPN devices or client software.

Client Connection Process

In general terms, an Easy VPN Remote device or a VPN Software Client version 3.x/4.x initiates a connection with a Cisco router configured as a Easy VPN Server. During the connection establishment the exchange includes device authentication via IKE, user authentication via IKE Extended Authentication (Xauth), VPN policies that are pushed down to the client, and then the IPSec SA is established.

The following is a more detailed look at the client/server session establishment.

1. The client initiates IKE Phase One exchange. If a preshared key is to be used for authentication, the exchange is initiated via IKE Aggressive mode. In this case, the group name entered while configuring the client with the web application is used to identify the group profile. If digital certificates are used, the exchange will be via IKE Main mode and the organizational unit field of a distinguished name will be used to identify the appropriate group profile.

2. The client attempts to negotiate an IKE SA with the Easy VPN Server. To reduce client configuration policies weren't defined, so all supported combinations of encryption and hash algorithms for authentication, plus supported Diffie-Hellman (DH) group sizes, are proposed.
   The Easy VPN Server Device accepts the first proposal received that matches its configured policies. Assuming a policy match is achieved, device authentication is completed and user authentication can begin.
3. If the Easy VPN Server is configured for Xauth, the server issues a user name/password challenge to the client. The resulting entries are verified against using AAA supported protocols, such as TACACS+, RADIUS, or one-time password token cards using AAA proxy. This step is particularly important if the peer is a remote client or a remote device is configured as a remote client.
4. The system parameters are pushed from the server to the client. These parameters can be configured to include an IP address (required) and the following optional information: DNS address(es), domain name, WINS address(es), local NAT pool name, access list, split tunnel attributes, and so forth. The access list defines the traffic to be protected through the VPN tunnel.
5. The Easy VPN Server can use reverse route injection (RRI) to create static routes and inject them into any dynamic routing protocols for distribution to surrounding devices. With dynamic crypto maps, a static route is created for each subnet or host protected by the remote peer when the peer establishes its IPSec security association. With *static crypto maps*, a static route is created for each destination using an extended access-list rule.
6. Once all parameters are transferred to the client, IKE Phase Two Quick mode is used to negotiate IPSec SAs to complete the connection.

Cisco Easy VPN Remote

The Cisco Easy VPN Remote provides a remote VPN client feature, which is currently supported on the following platforms: Cisco 800, 1700, and UBR900 Series Routers, Cisco PIX 501 Firewalls, and Cisco VPN 3002 Hardware Client device. As with all new technologies, the Cisco online documentation includes any new platforms supported.

The Cisco Easy VPN Remote feature eliminates much of the basic VPN configuration by implementing Cisco's Unity Client protocol, which allows most VPN parameters to be defined at a VPN remote access server. These basic configuration and security policies are "pushed" down from the Cisco Easy VPN Server during the initial connection. Configuration changes, software upgrades, and, in some cases, firmware updates can also be pushed down to the remote site. This can reduce local VPN configuration requirements and ongoing VPN support at the remote location.

The Cisco Easy VPN Remote feature allows for "push" configuration and automatic management of the following details:

- Managing security keys for encryption and decryption.
- Negotiating tunnel parameters, including IP addresses, algorithms, lifetime, and so forth.

- Establishing tunnels according to the defined security parameters.
- Enabling and configuring NAT/PAT translation, plus any related access lists.
- User authentication based on user names, group names, and passwords, as well as X.509 digital certificates.
- Authenticating, encrypting, and decrypting data through the tunnel.

The Cisco Easy VPN Remote feature supports the following two modes of operation:

Client	The private hosts protected behind the Easy VPN Remote device are a separate network that remains invisible and nonroutable to the central site. The local hosts are assigned their IP addresses from the Easy VPN Remote device DHCP server feature. The Cisco Easy VPN Remote feature automatically configures the NAT/PAT translation and any required access lists to implement the VPN tunnel. Because all traffic to the central network has the Public interface IP address, PAT both supplies and manages unique port number mappings to use in combination with the IP address.
Network Extension	The Cisco Easy VPN Remote device establishes a secure site-to-site connection with the central site device. PAT isn't used, allowing the client hosts to have direct access to the hosts on the corporate network. The local stations behind the VPN Remote device are fully routable and the local network is visible to the central site. As the name implies, the local network becomes a part of the organization's intranet.

Split Tunneling

Split tunneling is a useful feature that provides the capability to have a secure tunnel to the central site, while simultaneously maintaining a Cleartext connection to the Internet through the Internet service provider (ISP). The Cisco Easy VPN Remote device uses PAT to protect the local workstations during split tunneling to the Internet. Figure 12-1 shows secure a VPN tunnel and split tunneling for web browsing.

If the organization security policy prohibits split tunneling, it can be blocked by creating a policy on the central site device, which is then pushed down to the remote device.

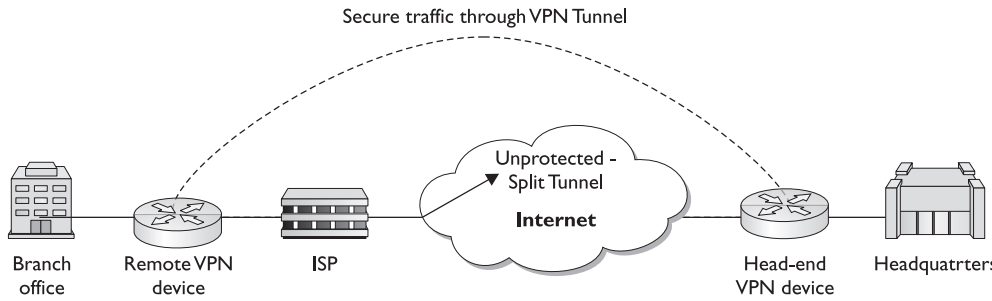

Figure 12-1 A secure VPN tunnel and split tunneling for web browsing

Cisco VPN 3.6 Client

The VPN Client is a thin design software program, which is easy to deploy and operate providing secure, end-to-end encrypted tunnels. The client software is available from the Cisco web site for use with any Cisco central site VPN product and is included free of charge with any Cisco VPN 3000 Series Concentrator.

The client software can be preconfigured for large, rapid deployments that require little end-user intervention. VPN access policies and configurations are downloaded and updated using push technology from the central VPN head-end device when a connection is established. This centralized push technology makes deployment and management quite simple, as well as highly scalable.

The Cisco VPN Client software is available in versions to support most Windows versions, Sun Solaris (UltraSparc 32 and 64 bit) UNIX, Linux (Intel), as well as Mac OS X 10.1 and 10.2 (Jaguar).

The Cisco VPN Client software, regardless of the host operating system (OS), is compatible with the following Cisco products:

- Cisco IOS Software Releases 12.2 T and later
- Cisco PIX Firewall Software Version 6.0 and later
- Cisco VPN 3000 Series Concentrator with Software Version 3.0 and later

How the VPN Client Works

The VPN Client works with a central Cisco VPN device to create a secure tunnel connection between the host computer and the corporate network. VPN Client uses Internet Key Exchange (IKE) and Internet Protocol Security (IPSec) tunneling protocols to make and manage the secure connection. Some of the steps include the following:

- Negotiating tunnel parameters, such as IP addresses, algorithms, lifetime, and so on.
- Establishing tunnels according to the defined parameters.
- Authenticating users based on user names, group names, passwords, and X.509 digital certificates.
- Establishing user access rights, including hours of access, connection time, allowed destinations, allowed protocols, and so forth.
- Managing security keys for encryption and decryption.
- Authenticating, encrypting, and decrypting data through the tunnel.

Connection Technologies

The VPN Client supports each of the following technologies for connecting to the Internet and, thereby, accessing the destination VPN endpoint.

- Plain Old Telephone Service (POTS)

- Integrated Services Digital Network (ISDN)
- Cable/modem
- Digital Subscriber Line (DSL)
- LAN connection

Easy VPN Server Configuration Tasks

The specific commands to configure the Easy VPN Server features vary, depending on the hardware platform. Other command options, such as the Auto Upgrade feature, apply to VPN Hardware Client devices only. Examples of the type of configuration tasks that need to be implemented on the VPN Server to support Easy VPN Clients include the following. The first three are required, while all others are optional.

- Enabling Policy Lookup via AAA (required)
- Defining Group Policy Information for Mode Configuration Push (required)
- Applying Mode Configuration and Xauth (required)
- Enabling Reverse Route Injection for the Client (optional)
- Enabling IKE Dead Peer Detection (optional)
- Configuring RADIUS Server Support (optional)
- Verifying Easy VPN Server (optional)

In addition, the Easy VPN Server feature enables Cisco IOS routers to push new and/or enhanced VPN policy parameters to any remote access Easy VPN client (hardware or software). This feature adds support for the following functions:

- Mode Configuration Version 6 Support—based on an IETF draft submission.
- Xauth Version 6 Support—based on an IETF draft submission.
- IKE Dead Peer Detection (DPD) —a new keepalive scheme.
- Split Tunneling Control—enables clients to have intranet and Internet access at the same time, without requiring Internet access to use the VPN tunnel.
- Initial Contact—to facilitate reestablishing lost connections.
- Group-Based Policy Control—Policy attributes, such as IP addresses, DNS, and split tunnel access can be enabled on a per-group or per-user basis.

Preconfiguring the Cisco VPN 3.6 Client

To use the VPN Client, at least one connection entry must be configured to define the following information:

- The Cisco VPN Remote Server to connect to.

- Preshared keys—The IPSec group to which the system administrator assigned the client. The group determines the level and type of access that will be available once connected. The group membership could define the user authentication method used, the hours the client can access the network, if and how many simultaneous logins are allowed, and even the IKE and IPSec algorithms the VPN Client will use.
- Certificates—The name of the certificate being used for authentication (if used).
- Optional parameters.

You can create multiple connection entries so the VPN Client software can be used to connect to multiple networks, although not at the same time.

Creating a New Connection Entry

Because the order of the following steps is important, we'll view this as if it were a lab exercise. While it isn't difficult, if the end user is going to do the initial configuration, the network administrator should furnish the information in written form with at least minimal instructions.

1. Start the VPN Client by choosing Start | Programs | Cisco Systems VPN Client | VPN Dialer from the Start button menu. Figure 12-2 shows the menu sequence.
2. The VPN Dialer application starts and displays the Cisco Systems VPN Client, as shown in Figure 12-3.

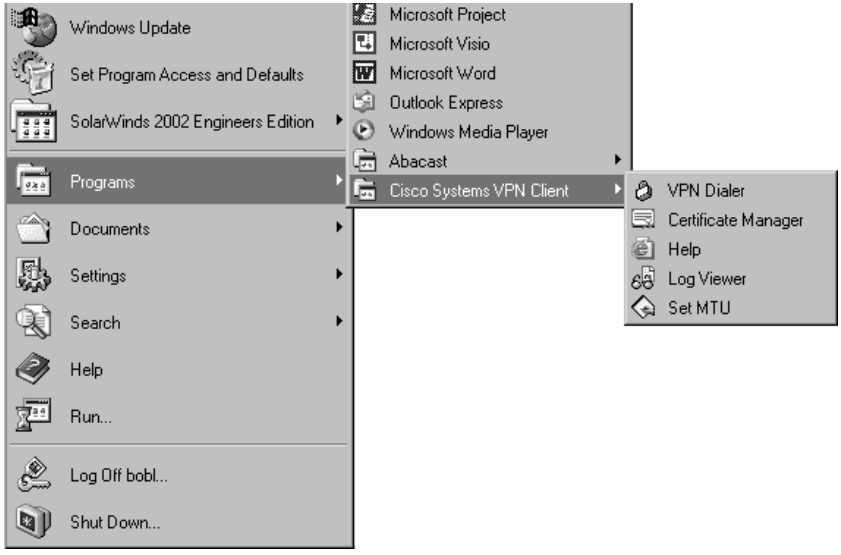

Figure 12-2 Starting the Cisco VPN Client software

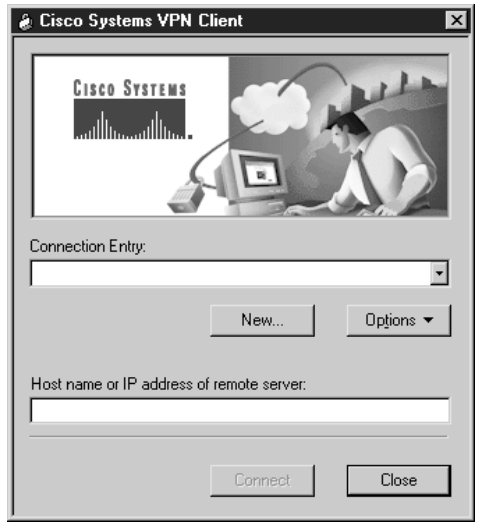

Figure 12-3 Cisco Systems VPN Client opening dialog box

3. On the Cisco Systems VPN Client main dialog box, click the NEW button. The first New Connection Entry Wizard dialog box appears, as shown in Figure 12-4.
4. Each connection on a host machine must have a unique name. The same name can be used for all users in a group who don't share a computer. If the users are configuring this unique name, the name and description might be defined by the network administrator to reduce any possible user doubt.
   Both the name and description can contain spaces and neither is case-sensitive. When the entries are complete, click the NEXT button.

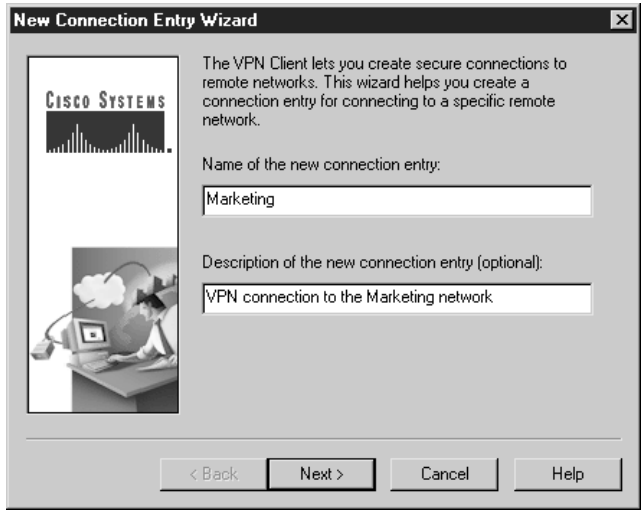

Figure 12-4 Define connection name and description

5. The next dialog box appears asking for host name or IP address of the server. This is the VPN head-end device to which the client will be connecting. Depending on the network, this could be either an IP address or a host name. The network administrator will furnish the proper form. This needs to be typed in exactly. Figure 12-5 shows an example of using an IP address to define a VPN device.
6. After typing the host name or IP address of the remote VPN device, click the NEXT button.
   The third New Connection Entry Wizard dialog box appears, asking for either the group information supplied by the network administrator or the digital certificate to be used to authenticate this client.
   Figure 12-6 shows an example of choosing group authentication. Notice the group name is displayed in Cleartext, but the passwords are masked for security. What isn't so obvious is that all three entries are case-sensitive. When the entries are complete, click the NEXT button.
7. Figure 12-7 shows the fourth and final New Connection Entry Wizard dialog box. The BACK button can be used to change earlier entries or the FINISH button can complete the process.
   When the final dialog box closes, the new connection entry now appears in the Connection Entry drop-down list on the VPN Client's main dialog box.

Trying Out the New Connection

Use the following steps to Start the VPN Dialer:

1. Start the VPN Client just as before.
   The VPN Dialer displays the Cisco Systems VPN Client main dialog box with the connection information from the earlier steps, as shown in Figure 12-8. If multiple connections were created, the Connection Entry: drop-down list would be used to choose the desired connection entry.

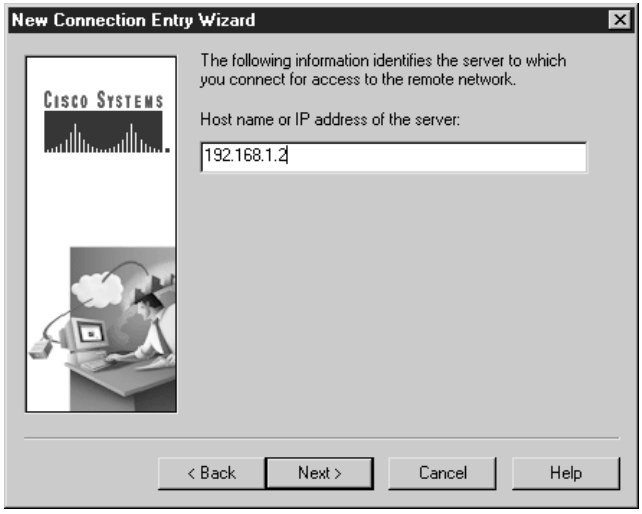

Figure 12-5 Using an IP address to define a VPN head-end device

Figure 12-6 Entering the IPSec group and password

2. Clicking the CONNECT button starts the connection. Figure 12-9 shows the connection progress dialog box.
3. The remaining screens will vary, depending on the authentication system used to verify the user. Figure 12-10 displays a Windows-type login screen.

Customizing the Connection

Figure 12-11 shows how to customize the connection and the available options. From copying (cloning) the connection settings, to deleting it, to launching third-party

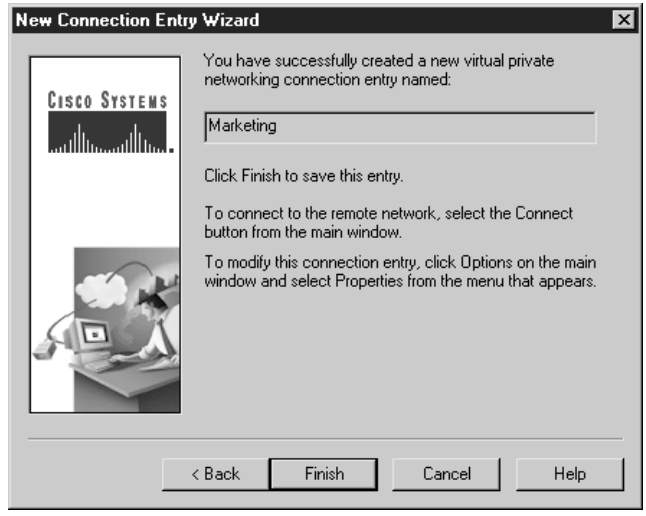

Figure 12-7 Final New Connection Entry Wizard dialog box

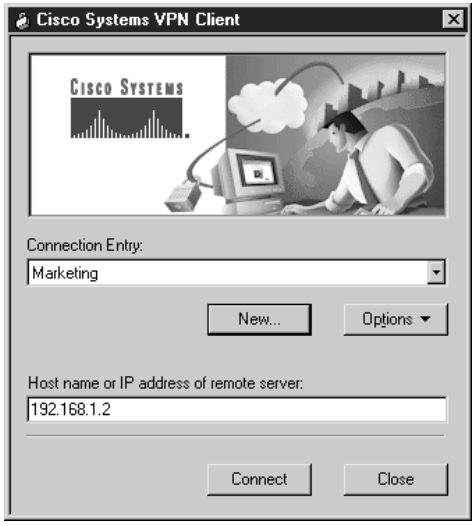

Figure 12-8 The Cisco Systems VPN Client main dialog box is ready to connect.

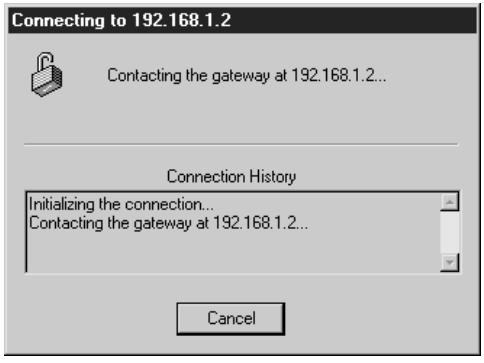

Figure 12-9 VPN connection progress dialog box

Figure 12-10 Windows user authentication box

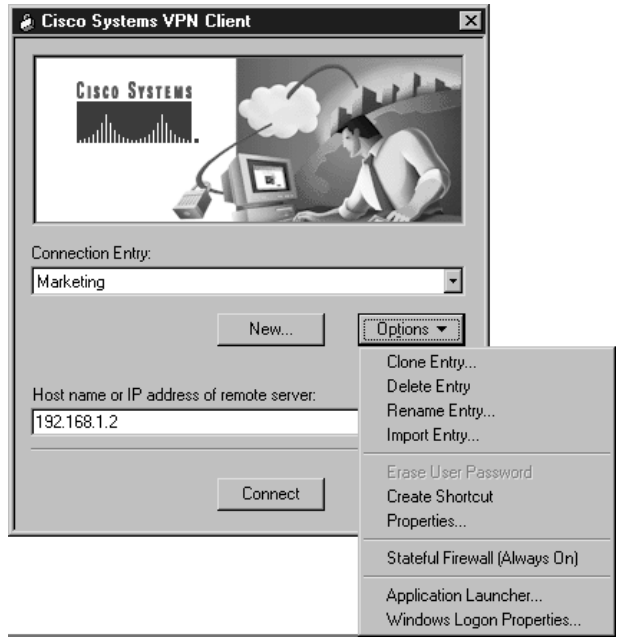

Figure 12-11 VPN Client customization options

applications, this tool offers some useful options. A new feature is support for the following personal firewalls (in addition to PIX support).

- Cisco Integrated Firewall (CIC)
- ZoneAlarmPro 2.6.3.57
- ZoneAlarm 2.6.3.57
- Zone Integrity
- BlackIce Agent and BlackIce Defender 2.5

Management Center for VPN Routers

Management Center for VPN Routers v1.1 (Router MC) is a web-based application designed for large-scale management of VPN and firewall configurations running on Cisco routers. Management Center for VPN Routers is one of the components of Cisco Works.

Router MC enables network administrators to configure and maintain VPN connections between multiple Cisco VPN routers using a hub-and-spoke topology. Router MC features support provisioning all critical connectivity, security, and performance parameters for a large-scale site-to-site VPN.

Router MC supports overlaying a VPN over a Frame Relay network for added security. It can be used to facilitate the migration from leased-line connections to Internet or intranet-based VPN connections.

Router MC-developed firewall and/or VPN configurations can be deployed to individual devices or groups of devices. The hierarchical device grouping and policy inheritance features of Router MC make it possible to configure multiple like devices simultaneously.

Features and Benefits

The Router Management Center offers features, such as smart rules hierarchy, resiliency support via IKE and generic routing encapsulation (GRE), import and deployment to files or devices, wizards-based support for setup of IKE and VPN tunnel policies, reusable building blocks, and more. The main features and benefits of the Router Management Center's are summarized as follows:

- Web browser interface—Makes it easy to define and deploy policies.
- Simplified policy definitions—Wizard-based interface that steps the user through the creation of IKE policies, tunnel policies, and transform sets.
- Support for a large number of devices—Uses device grouping, Smart Rules hierarchy, and reusable policy components to enable VPN configurations that scale to thousands of devices.
- Enhanced resiliency with IKE keepalives or GRE for failover—Supports IKE keepalives or GRE, with EIGRP or OSPF routing protocols, to provide failover options.
- Transparent translation of VPN policies to CLI commands—Web interface allows simplified and quick configuring and managing VPN policies. The application translates the VPN policies into CLI commands to be deployed to devices.
- Flexible deployment to files or devices—Choice of deploying configurations directly to devices as CLI commands or generates files containing CLI commands, which can be written to devices later.
- Configuration rollback—Possible to restore a device's previous configuration if not satisfied with a configuration after deploying VPN policies.
- Frame Relay network support—Deploy hub-and-spoke VPN configurations over a Frame Relay network.

Router MC v1.1 Firewall Features

Router MC v1.1 added support for the following firewall functionality features:

- Support for configuring ordered access rules to be assigned per interface.
- The ability to view a list of access rules per device or device group.
- Context Based Access Control (CBAC) feature support, including availability of the inspect action for access rules, alert and audit settings, fragmentation settings, DNS timeouts, protocol timeouts, and Denial of Service (DoS) prevention by monitoring half-open connections.

Router MC v1.1 Enhanced VPN Features

Router MC v1.1 added support for the following VPN features:

- Enhanced Certification Authority (CA) enrollment features, including support for trust-point and autoenrollment commands for devices running IOS 12.2(8)T and higher
- Advanced Encryption Standard (AES) encryption algorithm for use in IKE policies and transform sets
- High Availability (HA) groups of hubs, using Hot Standby Routing Protocol (HSRP)
- Multiple failover and routing policies within the Router MC device hierarchy
- Catalyst VPN Services Module support as a hub endpoint
- Network Address Translation (NAT) traversal
- Dialer interfaces

Router MC Server Requirements

Router MC is a component of Cisco's VPN/Security Management Solution (VMS) which integrates CiscoWorks, VPN Monitor, CiscoWorks Common Services, and other individual applications.

Requirement Type	Minimum Requirement
Hardware	Intel-based computer with 1 GHz or faster Pentium processor. Color monitor with video card capable of 16-bit color. CD-ROM drive. 10 Mbps (or faster) network connection.
Memory (RAM)	1GB.
Available disk space	9GB. 2GB virtual memory. NTFS file system (recommended).
Software	ODBC Driver Manager 3.510 or later. One of the following: Windows 2000 Professional. Windows 2000 Server. Service Pack 3.

Router MC Client Requirements

You can access all product features from any PC client meeting or exceeding the following hardware, software, and browser requirements.

Requirement Type	Minimum Requirement
Hardware	Intel-based computer with 300 MHz or faster Pentium processor.
Memory (RAM)	256MB.
Available disk space	400MB virtual memory.
Software	One of the following: Windows 98 Windows NT 4.0 Windows 2000 Server or Professional with Service Pack 2.
Web browser	Internet Explorer 6.0 or 5.5 with Service Pack 2. (Netscape Navigator not supported at this time.)

Router MC User Permissions

Router MC requires all users to log in with a user name and a password, which must be authenticated by the CiscoWorks server (default) or by a Cisco Secure Access Control Server (ACS) v3.1. Once authenticated, Router MC determines the user's defined role within the application. The role determines the set of tasks or operations the user is authorized to perform. Only those menu items, Table of Contents items, and buttons associated with authorized tasks are visible to the user. All others are hidden or disabled.

CiscoWorks Server Roles and Router MC Permissions

CiscoWorks supports the following five role types corresponding to typical functions within an organization:

Help desk	Read-only access for viewing devices, device groups, and the entire scope of a VPN.
Approver	Can review policy changes and can either approve or reject them. Can also approve or reject deployment jobs.
Network operator	Can make policy changes (except device inventory changes), as well as create and deploy jobs. Activities and jobs must be approved by an Approver.
System administrator	Can perform CiscoWorks server tasks and can make changes to the device hierarchy, such as move or delete devices. Can change administrative settings.
Network administrator	Can perform all CiscoWorks server and Router MC tasks. Can add users to the system with CiscoWorks or ACS. Can set user passwords and assign user roles and privileges.

If ACS authentication is used, the roles are different because of the group and task orientation of ACS. Each role is made up of a set of permissions that determine the role's level of access to Router MC tasks. User groups are assigned a role and each user in the group can perform all the Router MC tasks in the role. Use the Cisco online feature to get the most current configuration and implementation information if ACS will be used.

Easy VPN Remote Phase Two

Cisco Easy VPN Remote Phase Two features introduced with Cisco IOS v 12.2(8)YJ provide greater flexibility and feature options to Easy VPN Remote client devices, including Cisco 806, 826, 827, 828, and 1700 series routers, plus the Cisco uBR905 and uBR925 cable access routers. The goal is to provide the greatest flexibility in high-performance connections to the Internet, combined with the security of VPN, while minimizing the client-side configuration and support.

Easy VPN Remote Phase Two features still rely on most VPN parameters being defined at a VPN remote access server, such as a Cisco VPN 3000 concentrator, PIX Firewall, or a Cisco IOS router. When the IPSec client then initiates the VPN tunnel connection, the VPN remote access server pushes the IPSec policies to the IPSec client and creates the corresponding a VPN tunnel connection.

Supported VPN Servers

VPN Remote Phase Two requires that the destination server support the new features. The following device and OS combinations currently meet the requirements.

- Cisco uBR905 and uBR925 cable access routers running Cisco IOS v12.2(8)YJ or later
- Cisco 806, 826, 827, 828, 1700 series, 2600 series, 3600 series, 7100 series, 7200 series, and 7500 series routers running Cisco IOS v12.2(8)YJ or later
- Cisco PIX 500 series running software release 6.2 or later
- Cisco VPN 3000 series devices running software release 3.11 or later

The Cisco 806, 826, 827, 828, 1700 series routers, and uBR905 and uBR925 cable access routers support simultaneous remote client and remote server roles, allowing them to provide easy-to-implement VPN connections from small branch locations to corporate centers. The device is a client to the corporate center and a server to their VPN remote clients.

Phase Two Features

Cisco Easy VPN Remote Phase Two provides automatic management of the following features:

- Manual Tunnel Control
- Multiple Inside Interface Enhancements
- Multiple Outside Interfaces Support
- NAT Interoperability Support
- Local Address Support for Easy VPN Remote
- Cable DHCP Proxy Enhancement

- Peer Hostname Enhancement
- Proxy DNS Server Support
- PIX Interoperability Support
- Cisco IOS Firewall Support
- Simultaneous Easy VPN Client and Server Support
- Cisco Easy VPN Remote Web Manager

Manual Tunnel Control

With the original Cisco Easy VPN Remote, the VPN tunnel connects automatically on configuration. If the tunnel times out or fails, it automatically reconnects or retries indefinitely. Phase Two implements manual control over IPSec VPN tunnels, making it possible to establish and terminate the tunnel on demand. The manual feature is implemented with a new subcommand under the **crypto ipsec client ezvpn** command. The syntax is as follows:

```
Rtr1(config)#crypto ipsec client ezvpn name
Rtr1(config-crypto-ezvpn)#connect [auto | manual]
```

Automatic is the default setting, compatible with Phase One functionality. As such, the subcommand with the autoparameter is only needed to reverse the manual option.

With the manual option, the Easy VPN Client waits for the following command to attempt to establish the connection or to reestablish a timed out or failed session:

```
Rtr1#crypto ipsec client ezvpn connect name
```

Use the following **clear** command to disconnect an established tunnel.

```
Rtr1#clear crypto ipsec client ezvpn [name]
```

Multiple Inside Interface Enhancements

Phase One supports only one inside interface on the remote client router. The **crypto ipsec client ezvpn** *name* **inside** command option allows designating up to three inside interfaces. Each inside interface supports only one tunnel. The syntax is as follows:

```
Rtr1(config-if)#interface interface-id
Rtr1(config-if)#crypto ipsec client ezvpn name [outside | inside]
```

The following example shows configuring an outside interface using the default designation outside:

```
Rtr1(config)#interface ethernet 0
Rtr1(config-if)#crypto ipsec client ezvpn vpn1 inside
```

Multiple Outside Interfaces Support

Phase One supports only one outside interface on the remote client router. The **crypto ipsec client ezvpn *name* outside** command option allows designating up to four outside interfaces. The default option is outside. The syntax is as follows:

```
Rtr1(config-if)#interface interface-id
Rtr1(config-if)#crypto ipsec client ezvpn name [outside | inside]
```

This feature is applicable only to platforms, such as the Cisco 1700 series routers, that support multiple outside interfaces.

While each inside or outside interface supports only one tunnel, multiple inside interfaces can be mapped to one outside interface.

The following example shows configuring an outside interface using the designation outside. As the default, the outside didn't need to be included:

```
Rtr1(config)#interface serial 0/0
Rtr1(config-if)#crypto ipsec client ezvpn vpn1 outside
```

NAT Interoperability Support

With the Easy VPN Client, the features automatic NAT and access list configuration replaced any existing NAT and access list configuration. If a tunnel timed out or dropped its connection, the NAT and access configuration were removed automatically, preventing any Internet access even to nontunnel destinations.

Cisco Easy VPN Remote Phase Two supports interoperability with locally configured NAT. When the IPSec VPN tunnel is down, the router automatically restores the previous NAT configuration. Users can continue to access nontunnel Internet connections when the tunnel times out or disconnects.

Local Address Support for Easy VPN Remote

Easy VPN Remote Phase Two provides an interface configuration option, which makes it possible to specify the interface to use in determining the IP address as the source of VPN tunnel traffic. Typically, the loopback interface is the interface used to source tunnel traffic. The syntax is as follows:

```
Rtr1(config)#crypto ipsec client ezvpn name
Rtr1(config-crypto-ezvpn)#local-address interface-id
```

The following example shows the **local-address** subcommand used to specify the loopback0 interface for sourcing tunnel traffic:

```
Rtr1#config t
Rtr1(config)#crypto ipsec client ezvpn telecom-client
Rtr1(config-crypto-ezvpn)#local-address loopback0
Rtr1(config-crypto-ezvpn)#
```

Cable DHCP Proxy Enhancement

With Phase Two, cable providers can use the Cable DHCP Proxy feature to obtain a public IP address, assign it to the loopback interface, and then have the cable modem interface get its IP address from the loopback interface. The Phase Two feature enhancement applies to the existing **cable-modem dhcp-proxy interface** configuration command for the uBR905 and uBR925 cable access routers.

For the router to configure the loopback interface automatically with the public IP address obtained from the DHCP server, the loopback interface must be created before issuing the **cable-modem dhcp-proxy interface** command.

The following example shows a loopback interface created first, and then the loopback interface being specified, so the router automatically assigns it with the public IP address:

```
Rtr1#config t
Rtr1(config)#interface loopback 0
Rtr1(config-if)#interface cable-modem 0
Rtr1(config-if)#cable-modem dhcp-proxy interface loopback0
```

NOTE The **cable-modem dhcp-proxy interface** command is currently only supported for the Cisco uBR905 and uBR925 cable access routers.

Peer Host Name Enhancement

When defining the VPN peer, you can use either an IP address or a host name. When a host name is used, a DNS lookup is performed immediately to resolve the name to an IP address. Because Phase Two supports manual tunnel control, and to support DNS entry changes, the host name text string is stored and the DNS lookup is performed at the time of the tunnel connection. The host name enabling syntax is as follows:

```
Rtr1(config)#crypto ipsec client ezvpn name
Rtr1(config-crypto-ezvpn)#peer [ip-address | hostname]
```

The following shows the crypto entries:

```
Rtr1(config)#crypto ipsec client ezvpn vpn-client
Rtr1(config-crypto-ezvpn)#connect auto
Rtr1(config-crypto-ezvpn)#group vpn-client-grp key vpn-client-password
Rtr1(config-crypto-ezvpn)#local-address Loopback0
Rtr1(config-crypto-ezvpn)#mode client
Rtr1(config-crypto-ezvpn)#peer vpn-server
```

Proxy DNS Server Support

When the VPN connection to the corporate network is up, the enterprise DNS servers should resolve domain names to IP addresses. But when the VPN connection to the

enterprise is down, the ISP or cable provider DNS servers should be used to resolve DNS requests.

An Easy VPN Remote Phase Two router can be configured to act as a proxy DNS server. As a proxy DNS server for LAN, the router receives DNS queries from local users on behalf of the enterprise DNS server. The DHCP server sends out the router's LAN address as the DNS server IP address. When the VPN tunnel connection comes up, the router forwards the DNS queries to the enterprise DNS server. Otherwise, they're forwarded to the ISP DNS.

To enable the proxy DNS server functionality with the **ip dns server** command in Global Configuration mode, use the following commands beginning in Global Configuration mode.

PIX Interoperability Support

Cisco Easy VPN Remote Phase Two more fully supports the Cisco PIX Firewall v6.2 features than the original implementation of Easy VPN Remote.

Cisco IOS Firewall Support

Cisco Easy VPN Remote Phase Two more fully supports the Cisco IOS Firewall feature set than the original implementation of Easy VPN Remote.

Simultaneous Easy VPN Client and Server Support

Cisco Easy VPN Remote Phase Two more fully supports configuring simultaneous Easy VPN Client and Cisco Easy VPN Server support on the same Cisco 1700 series routers. You can configure one outside interface as a Cisco Easy VPN Server and another outside interface on the same router as a Cisco Easy VPN Client. Figure 12-12 shows an example of a router (Rtr2) acting as both an Easy VPN Client and a Server.

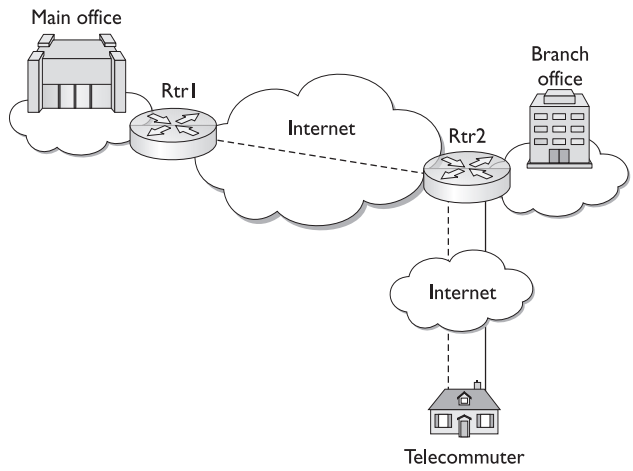

Figure 12-12 Router (Rtr2) acting as both an Easy VPN Client and a Server

The following example shows the configuration for the VPN client and server features on Rtr2. Some lines were eliminated to conserve space:

```
Rtr2#show run
version 12.2
!
hostname Rtr2
!
aaa new-model
aaa authorization network vpn-client-grp local
aaa session-id common
!
ip subnet-zero
no ip domain-lookup
ip ssh time-out 120
ip ssh authentication-retries 3
!
crypto isakmp policy 10
 authentication pre-share
 group 2
crypto isakmp client configuration address-pool local loc-pool
!
crypto isakmp client configuration group vpn-server-grp
 key vpn-grp-key
 dns 192.168.0.13 192.168.0.15
 wins 192.168.0.14 192.168.0.16
 domain vpn-test.com
 pool loc-pool
!
crypto ipsec transform-set trans-set-1 esp-3des esp-md5-hmac
!
crypto ipsec client ezvpn client1
 connect auto
 group vpn-client-grp key vpn-grp-key
 mode client
 peer 1.1.100.17                                    (Rtr1)
!
crypto dynamic-map dyn-map 1
 set transform-set trans-set-1
!
crypto map dyn-map isakmp authorization list vpn-server-grp
crypto map dyn-map client configuration address respond
crypto map dyn-map 1 ipsec-isakmp dynamic dyn-map
!
interface FastEthernet0/0
 description Connection to Branch Office - VPN Clients
 ip address 5.0.0.1 255.0.0.0
 crypto ipsec client ezvpn client1 inside
!
interface Serial0/0
 description Connection to Corporate Network - VPN Server
 ip address 1.0.0.1 255.0.0.0
 no fair-queue
 crypto ipsec client ezvpn client1
!
interface Serial0/1
 description Connection to telecommuters - VPN Clients
 ip address 1.2.0.1 255.255.0.0
```

```
 crypto map dyn-map                              (for server functionality)
 crypto ipsec client ezvpn client1 inside        (for client functionality)
!
ip local pool loc-pool 1.2.0.3 1.2.0.31
ip classless
!
line con 0
line aux 0
line vty 0 4
end
Rtr2#
```

Cisco Easy VPN Remote Web Manager

Cisco Easy VPN Remote Phase Two introduced using the Cisco Easy VPN Remote Web Manager to manage the Cisco uBR905 and Cisco uBR925 cable access routers. The Cisco Easy VPN Remote Web Manager is a built-in web-interface application resident on the uBR905 and uBR925 devices. The Web Manager enables the user to avoid the command-line interface (CLI) to perform the following functions:

- See the current status of any Easy VPN Remote Phase Two tunnels
- Connect or disconnect a tunnel configured for manual control
- Reset a tunnel configured for automatic connection
- Be prompted for Xauth information if Xauth information is needed

Cisco VPN Firewall Feature for VPN Client

The VPN Client software now includes an integrated stateful firewall feature set that provides protection to the client. The feature set protects the VPN Client PC from Internet attacks both from split-tunneling implementations and IPSec tunnel connections to a VPN Concentrator. This feature is called Stateful Firewall (Always On).

Overview of Software Client Firewall Feature

The built-in Stateful Firewall (Always On) service provides even tighter security by blocking all new inbound sessions from all networks, regardless of whether a VPN connection is active. The Stateful Firewall filtering applies to both encrypted and nonencrypted traffic. Outbound traffic creates entries in a state table, which allows returning packets to be allowed through. Any sessions originating on the outside interface are blocked by default, though, because no state table entries exist.

Two exceptions exist to this no unsolicited inbound traffic rule. The first involves supporting DHCP services: DHCP client requests to a DHCP server pass out on one port, but the resulting responses return through a different port. The Stateful Firewall feature is programmed to know this and allows that specific inbound traffic. The second exception is edge services processor (ESP) traffic through ESP modules from the secure gateway. The Stateful Firewall software recognizes ESP traffic as packet filters, and not as session-based filters, and allows it through.

To enable the Stateful Firewall, click Stateful Firewall (Always on) on the Options menu, as shown in Figure 12-13. The check in front of the option indicates the Stateful Firewall (Always On) feature is enabled. This feature is disabled by default. The feature can be enabled or disabled by clicking the entry in the VPN Client Options menu.

During a VPN connection, you can view the status of the firewall features by double-clicking the lock icon in the taskbar system tray or right-clicking the same icon and choose Status from the resulting menu. You can also enable or disable the feature from the same menu. The result is a three-tab window, as shown in Figure 12-14, with the firewall features on the third tab. The information displayed on the tab varies according to the configured firewall policy.

Defining a Client Firewall Policy

The VPN Concentrator network administrator can define and manage the firewall policy using the Configuration | User Management | Base Group or Group | Client FW tab. You can choose from three options:

- Are You There
- Centralized Protection Policy
- Client/Server

The Are You There Feature

Since v3.1, the Cisco VPN Client supports the Are You There (AYT) feature. When the AYT feature is enabled, the VPN Client polls the local firewall every 30 seconds to make

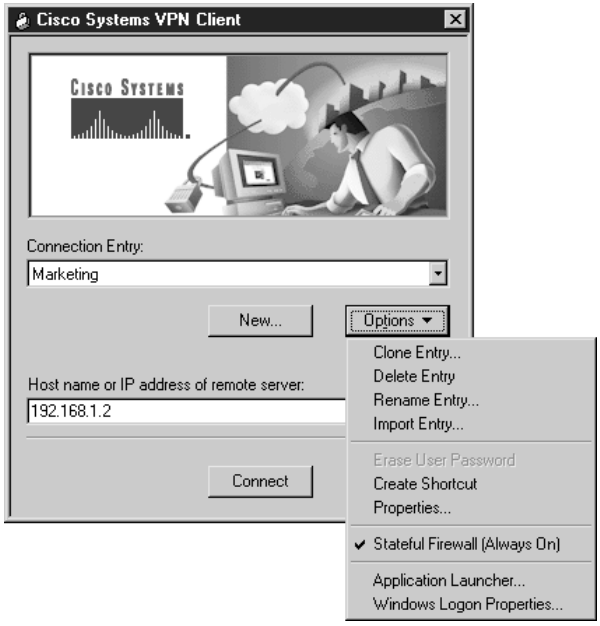

Figure 12-13 Stateful Firewall (Always on) on the Options menu

PART III

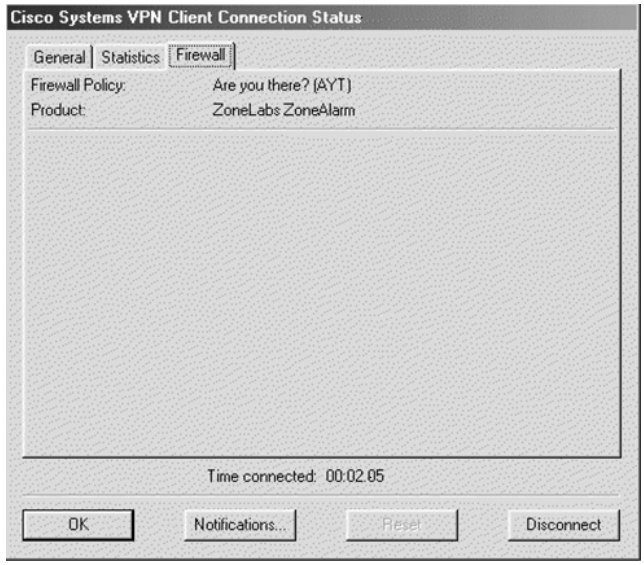

Figure 12-14 Cisco System VPN Client Connection Status information box

sure it's still running. While the VPN Client confirms the firewall is running, it doesn't confirm that a specific policy is enforced.

If the security policy requires that remote users have firewalls running on their PCs, the VPN Concentrator can allow these clients to connect only if they have the designated firewall installed and running. If the designated firewall isn't running, the connection attempts fails. Once the connection is established, the VPN Client uses the AYT feature to monitor the firewall to make sure it's running. If the firewall stops for any reason, the VPN Client immediately drops the connection to the VPN Concentrator.

The Cisco System VPN Client Connection Status information box Firewall tab shows only the firewall policy (AYT) and the name of the firewall product, as shown in Figure 12-14 earlier.

The Central Policy Protection Feature

Central Policy Protection (CPP) is a stateful firewall policy that leverages the Cisco Integrated Client feature by letting the VPN Concentrator manage the client firewall policies. The specific policy rules are defined by the administrator on the VPN Concentrator, and then pushed down to the VPN Client every time a connection is attempted. The VPN Client then enforces these policy rules for all nontunneled (split-tunnel Internet) traffic while the tunnel is active.

Because CPP works only on out-of-tunnel Internet traffic, if the client is operating in a tunnel-everything mode, enabling CPP has no effect.

The Cisco System VPN Client Connection Status information box Firewall tab shows the firewall policy, the firewall in use, and firewall rules, as shown in Figure 12-15.

Firewall Rules

The Firewall Rules section of the Status box shows all the firewall rules currently implemented on the VPN Client. The rules are arranged in order of importance, with the highest importance at the top. All but the last two rules are defined by the VPN administrator to allow inbound and outbound traffic between the VPN Client and the secure gateway, as well as between the VPN Client and the private networks with which it communicates. Because the rules are implemented from the top down, the VPN Client enforces them before trying the two CPP default rules at the bottom. This approach lets the traffic flow to and from private networks.

The bottom two rules define the filter's default actions, which are to drop both inbound and outbound traffic. These rules are implemented only if the traffic doesn't match any of the preceding rules.

To see the full fields of a specific rule, click the first column in the top half of the Firewall Rules: window; the selected rule is displayed in the bottom half of the window.

Figure 12-15 Firewall tab for CPP

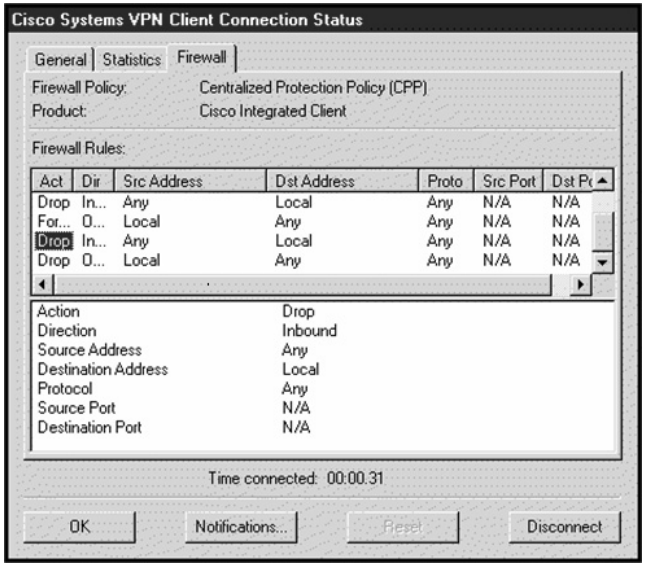

A firewall rule includes the following fields and options:

Action	Action to be taken if the data traffic matches the rule: **Drop**—Discard the session. **Forward**—Allow the session to go through.
Direction	Direction of traffic to be affected by the rule: **Inbound**—traffic coming in to the local machine. **Outbound**—traffic going out from the local machine.
Source Address	Source address of the traffic this rule affects: **Any**—all traffic, for example, drop any inbound traffic. **IP address and subnet mask**—A specific host address. **Local**—The local machine for outbound traffic.
Destination Address	Destination address this rule affects: **Any**—All traffic. **Local**—Local machine if the direction is inbound.
Protocol	The Internet Assigned Number Authority (IANA) number of the protocol covered by this rule concerns (6 for TCP, 17 for UDP).
Source Port	Source port used by TCP or UDP.
Destination Port	Destination port used by TCP or UDP.

The Stateful Firewall Process

In the stateful Cisco Integrated Client, firewall protocols TCP, UDP, and ICMP automatically allow inbound responses to outbound packets. To allow inbound responses to outbound packets for any other protocols, the network administrator needs to define specific filters on the VPN Concentrator. These are then passed down to the VPN Client the next time a session is established.

Client/Server Feature

The Client/Server policy supports the Zone Labs Integrity solution. *Zone Labs Integrity* is a Client/Server firewall solution in which the Integrity Server acts as the firewall server that pushes firewall policy to the Integrity Agent residing on the VPN Client PC. Because Integrity is a fully functional stateful personal firewall, it can intelligently decide on network traffic based on application layer information, as well as on traditional Layer 3 and 4 fields.

The Cisco System VPN Client Connection Status information box Firewall tab shows the firewall policy as Client/Server, the name of the product as ZoneLabs Integrity Agent, the user ID, session ID, and the addresses and port numbers of the firewall servers, as shown in Figure 12-16.

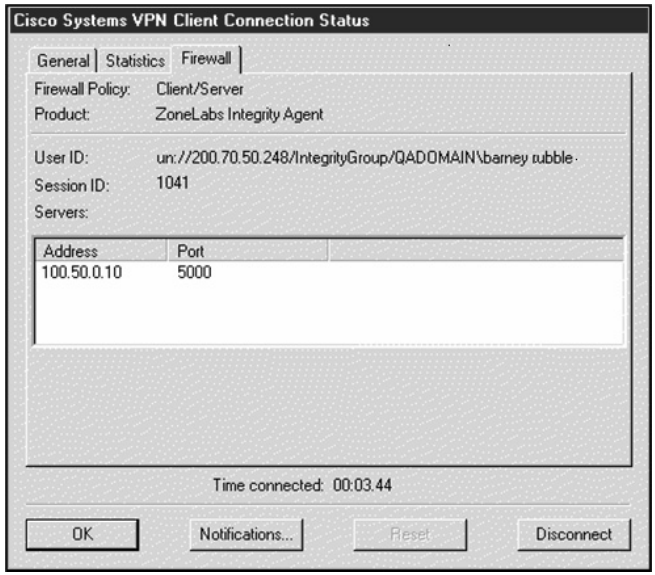

Figure 12-16 Client/Server Firewall tab

Client Firewall Statistics

The Statistics tab on the Cisco System VPN Client Connection Status dialog box shows statistics on the VPN Client data packets processed during the current session or since the statistics were reset. The data collected includes the following information. Figure 12-17 shows the Statistics tab information.

Bytes in	Total data received after the secure packets have been decrypted.
Bytes out	Total amount of encrypted data transmitted through the tunnel.
Packets decrypted	Total number of data packets received on the port.
Packets encrypted	Total number of secured data packets transmitted out the port.
Packets bypassed	Total number of packets not processed by the VPN Client because they didn't need to be encrypted, such as local ARPs and DHCP.
Packets discarded	Total number of packets the VPN Client rejected because they weren't from the VPN peer device.

Clicking the RESET button clears these counters, but doesn't impact the other two tabs. Clearing the counters often makes seeing the current activity easier, particularly if some time has passed since the last reset.

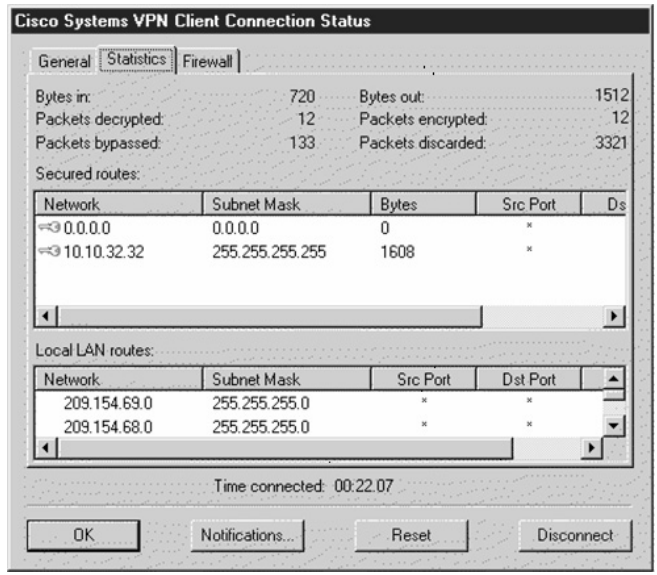

Figure 12-17 VPN Client connection statistics

Chapter Review

Cisco Easy VPN is part of Cisco's Unified Client Framework, in which VPN management is centralized across all Cisco VPN devices. The goal is to simplify VPN implementation, at least at the remote end. The Easy VPN strategy incorporates all VPN platforms, including VPN software client, VPN 3002 hardware client, Cisco routers, Cisco PIX Firewalls, and Cisco VPN 3000 Concentrators.

Cisco Easy VPN strategy incorporates centrally stored and managed configurations with dynamic configuration of end-user VPN devices. This should require less manual configuration by end users and field technicians, reducing errors and additional service calls.

The Cisco Easy VPN strategy has two main components: the Cisco Easy VPN Server and the Cisco Easy VPN Remote. The Easy VPN Server service allows a growing number of Cisco IOS routers, PIX Firewalls, and Cisco VPN 3000 Concentrators to act as VPN head-end devices in site-to-site or remote-access VPNs. In addition to providing VPN connections, these devices would store the client configurations and "push" those settings down to the Easy VPN Remote devices on their next connection.

Cisco VPN client software is available, for many common computer OSs bring those systems directly in contact with the VPN networks. While this isn't a full part of the Easy VPN Remote family, the software client shares similar features and ease of configuration.

Questions

1. Which of the following Cisco products is *not* a Cisco Easy VPN Server?
   - **A.** PIX Firewall
   - **B.** IOS routers
   - **C.** IDS Manager
   - **D.** VPN Series Concentrator
2. Which two of the following are modes of operation of Cisco Easy VPN Remote?
   - **A.** Server
   - **B.** Client
   - **C.** Network extension
   - **D.** Push
3. Split tunneling refers to which one of the following?
   - **A.** Running two simultaneous VPN connections
   - **B.** Allowing two-way simultaneous exchanges
   - **C.** Unsecured Internet connections in addition to secure VPN sessions
   - **D.** Load balancing VPN circuits
4. Which of the following is marketed as both an Easy VPN Client and an Easy VPN Server?
   - **A.** Cisco 800
   - **B.** Cisco 1700
   - **C.** PIX 501
   - **D.** Cisco VPN 3002
5. Which one of the following technologies is *not* typically used by VPN Clients?
   - **A.** DSL
   - **B.** Cable modem
   - **C.** Dial-up modems
   - **D.** T1 dedicated line
6. Which one is *not* a required Easy VPN Server configuration task?
   - **A.** Applying Mode Configuration and Xauth
   - **B.** Enabling Policy Lookup via AAA
   - **C.** Configuring RADIUS Server Support
   - **D.** Defining Group Policy Information for Mode Configuration Push

7. Which of the following is a new VPN keepalive scheme?
   A. Split Tunneling Control
   B. IKE Dead Peer Detection (DPD)
   C. Initial Contact
   D. Group-Based Policy Control
8. In setting up a Cisco VPN client connection to use preshared keys, what information is used to authenticate the connection?
   A. Digital certificates
   B. Local IP address
   C. Server IP address
   D. Group name and password
9. With which one of the following must the user configure their computer to support VPN?
   A. PIX Firewall with VPN service
   B. IOS Router with VPN service
   C. VPN 3002 Hardware client
   D. Cisco VPN Client
10. Which VPN implementation would *not* usually include NAT or PAT?
    A. PIX Firewall with VPN service
    B. IOS Router with VPN service
    C. Cisco VPN Client
    D. VPN 3002 Hardware client
11. Which of the following is *not* true about Management Center for VPN Routers v1.1 (Router MC)?
    A. It's a component of Cisco Works
    B. It uses a web interface
    C. It's a router-based Java application
    D. Developed configurations can be deployed directly to a device
12. Which of the following is *not* a supported feature of Management Center for VPN Routers v1.1 (Router MC)?
    A. Configuration rollback
    B. Transparent translation of VPN policies to CLI commands
    C. Simplified policy definitions
    D. Command-line interface for easy configuration

13. Which of the following Cisco Easy VPN Remote Phase Two features is only supported on the Cisco uBR905 and Cisco uBR925 cable access routers?
    A. Multiple Inside Interface Enhancements
    B. Simultaneous Easy VPN Client and Server Support
    C. Cisco Easy VPN Remote Web Manager
    D. NAT Interoperability Support
14. Which two of the following can be used for the Router Management Center server?
    A. Windows NT server
    B. Windows 2000 Professional
    C. Windows 2000 Server
    D. Sun Solaris 8.0 or above
15. Which of the following Cisco Easy VPN Remote Phase Two features allows a loopback address to be defined as the VPN tunnel source?
    A. Manual Tunnel Control
    B. Peer Host Name Enhancement
    C. Local Address Support
    D. NAT Interoperability Support

Answers

1. **C.** IDS Manager
2. **B.** Client and **C.** Network extension
3. **C.** Unsecured Internet connections in addition to secure VPN sessions
4. **B.** Cisco 1700
5. **D.** T1 dedicated line
6. **C.** Configuring RADIUS Server Support
7. **B.** IKE Dead Peer Detection (DPD)
8. **D.** Group name and password
9. **D** Cisco VPN Client
10. **C.** Cisco VPN Client
11. **C** It is a router-based Java application—it's a Windows 2000 application.
12. **D.** Command-line interface for easy configuration—Web interface

13. **C.** Cisco Easy VPN Remote Web Manager
14. **B.** Windows 2000 Professional and **C.** Windows 2000 Server
15. **C.** Local Address Support

CHAPTER 13

Cisco VPN Hardware Overview

In this chapter, you will learn to:

- Use Cisco products to enable a secure VPN
- Know about Cisco VPN 3002 Client devices
- Recognize Cisco VPN 3000 Concentrator devices

In this chapter, you learn about VPN hardware technologies supported by Cisco. Specific configuration is covered in the relevant chapters. Router-based VPNs are covered in Chapters 10 and 11. VPN client software is covered in Chapter 12. VPN connections to PIX Firewalls are covered in Chapter 21, and Cisco VPN 3000 Series devices are covered in Chapters 14 and 15.

Cisco Products Enable a Secure VPN

A wide variety of Cisco products actively support VPN connections throughout the network. Cisco's Unified Client Framework strategy enables all adhering VPN client technologies to connect to any Cisco central-site VPN concentrators, including the Cisco 3000 series VPN concentrators, PIX Firewalls, and Cisco IOS-based routers.

The Cisco VPN software client is typically provided at no charge with central site concentrators, so the software client is inexpensive to deploy and scale to large numbers. These VPN clients have specific limited operating system (OS) support and might require some level of support. These software clients could be impractical in those extranet environments where the organization doesn't own or control the remote PC, such as in a franchise or customer location. The organization might not want to absorb the expense associated with maintaining and supporting non-company workstations.

Cisco small router and firewall devices can be an inexpensive way to provide security features like stateful firewall and VPN capabilities to remote sites. Unfortunately, these devices don't scale beyond a few hundred devices. Deployment, monitoring, and ongoing support can be inconvenient, time-consuming, and expensive because of the variety of features and options that must often be manually configured at both ends of the connection.

The Cisco VPN 3000 Series Concentrator, combined with the Cisco VPN 3002 Hardware Client device, scales to support customers with 100 or fewer remote access users to

large organizations with up to 10,000 simultaneous connections. The Cisco VPN 3000 Series Concentrator is available in both nonredundant and redundant configurations, providing the customer with multiple options to design and build robust, reliable, and cost-effective networks.

What's New?

The latest Cisco VPN 3000 Series Concentrator OS, release 3.6, is the most recent version at press time. This release includes simplified VPN Software Client—Easy VPN—which allows minimal client configuration and "push" technology to download the configuration from the concentrator. Additional release 3.6 features include AES support, expanded support for Microsoft's Integrated VPN Client in Windows 2000/XP, and security for wireless environments.

The new Cisco VPN 3002 Hardware Client, covered in the next section, combines easy implementation and rollout, plus high scalability, with the performance, reliability, and stability of a hardware device.

The Cisco VPN 3000 Series Concentrator Wireless Client Support for Elliptic Curve Cryptography (ECC) provides faster information processing by personal digital assistants (PDAs) and smart phones.

The Cisco VPN 3000 Series Concentrator, Cisco VPN 3002 Hardware Client, and the Cisco VPN Client work together with the Cisco Internet Mobile Office to provide mobile professionals with secure, high-speed broadband connectivity to their networks in airports, convention centers, hotels, and a growing number of other public spaces.

Cisco VPN 3002 Client Devices

The Cisco *VPN 3002 Hardware Client* was specifically designed for those organizations with many remote users and sites that need to operate as secure clients in a VPN environment. The 3002 combines the ease of configuration and installation, plus the scalability of a software VPN client, with the stability and reliability of a hardware platform.

The VPN 3002 Hardware Client eliminates the need to install and configure VPN client software on the local workstation(s), plus it supports workstations running any OS, including Windows, Sun Solaris UNIX, Mac, and Linux. Each 3002 client appliance can connect one or more devices, including workstations, servers, hubs, cash registers, printers, and IP telephones to a company's central network. Figure 13-1 shows the book- sized (2 × 8.6 × 6.5 inches) 3002 device.

Figure 13-1 Cisco VPN 3002 Client device (front view)

The VPN 3002 is designed with simplicity and reliability of installation. It has few local setup parameters that must be configured. Additional parameters and policy are "pushed" to the device from the central site (head-end) device with the first connection. In Chapter 15, you see how little installation configuration must be done to connect to the corporate network.

The user simply plugs the minimally configured VPN 3002 device into a DSL/cable connection, router, or other wide area network (WAN) access devices at the remote site. The central site VPN concentrator takes over, using push policy features to centrally set policy, manage, and upgrade the device. This central control and management approach minimizes the need to rely on remote users to deploy or maintain the unit. Tunnel setup and policy configuration is automated, so companies needn't dedicate IT staff to configure individual devices manually. Troubleshooting aids and centralized monitoring features are built into the 3002 software to ensure proper operation after the unit is set up.

This ease of installation and because the 3002 can coexist with other types of VPN clients on the network make it an ideal choice for the organization that needs to set up and support dozens, or even thousands, of remote end users who need secure network connections from geographically dispersed branch or home office sites. Some examples that might fit this business model include fast food outlets, grocery store-based banking operations, retail chains, insurance or brokerage offices, employment offices, vehicle dealerships, drugstores, and the like.

In calculating return on investment (ROI), larger enterprises typically find the initial price of the hardware client is more than offset by the savings from reduced or eliminated service calls common with supporting software VPN clients. The same is true when compared to supporting expanding LAN-to-LAN networks with their complex configuration requirements at central and remote sites, particularly if on-site support is limited.

Cisco VPN 3002 Client Models

The *3002 Client*, a small-footprint, book-sized device designed for wall mount or table top operation, currently comes in two models. The *CPVN3002-K9* has two 10/100 Mbps RJ-45 Ethernet autosensing interfaces: one for the outside or public connection and one for the inside or private connection. This makes the device ideal for placement between a DSL or a cable modem, either connected directly to a workstation or to multiple hosts via a separate switch or hub.

The *CVPN3002-K9* sports a single 10/100 Mbps RJ-45 Ethernet autosensing interface for the public connection plus eight 10/100 Mbps Ethernet autosensing interfaces via an integrated switch. This device can provide simplified installation for the small office or home that needs to connect several computers to the network, or the switch can connect to additional switches, providing connections for up to 253 host devices. Figure 13-2 shows the back of the 8E unit.

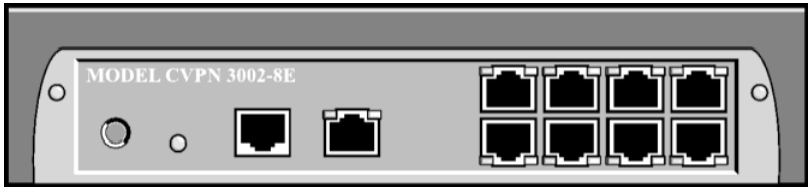

Figure 13-2 Cisco VPN 3002 Client with built-in 8-port switch

Client and Network Extension Modes

The Cisco VPN 3002 supports two modes of operation to offer implementation choices based on flexibility, security, and easy configuration. Those modes are

- Client mode
- Network Extension mode

A large VPN implementation might frequently have both types of operation.

Client Mode

In *Client mode*, the VPN 3002 emulates the VPN client software appearing to the main network like a remote user. The private hosts protected behind the VPN 3002 are a separate network that remains invisible and nonroutable to the central site. The local hosts are assigned their IP addresses from the VPN 3002 Dynamic Host Control Protocol (DHCP) server feature, while the public port can use the VPN 3002 DHCP client feature to acquire its IP address from an ISP. From a cost and address preservation standpoint, it would make sense for the local IP addresses to be private IP addresses.

To help secure the local network and to allow local hosts to travel out of the network in Client mode, the VPN 3002 uses Port Address Translation (PAT). Because all traffic to the central network will have the public interface IP address, PAT supplies and manages unique port number mappings to be used in combination with the IP address.

Split tunneling is a useful feature that provides the capability to have a secure tunnel to the central site, while simultaneously maintaining a clear text tunnel to the Internet through the ISP. The VPN 3002 uses PAT to protect the local workstations during split tunneling to the Internet. If the organization security policy prohibits split tunneling, it can be blocked by creating a policy on the central site device, which is then pushed down to the 3002 Client.

The VPN 3002 Client can only create outbound connections, so no way exists for an outside source to initiate a connection with the VPN 3002 or through it to the stations behind.

Network Extension mode

In *Network Extension* mode, the VPN 3002 establishes a secure site-to-site connection with the central site device. The local stations behind the VPN 3002 are fully routable and the local network is visible to the central site. As the name implies, the local network becomes part of the organization's intranet. VPN and 3002 configuration and security policies are pushed from the central site.

In Network Extension mode, the private addresses are assigned manually and permanently, allowing central site host and applications to reliably reach any local server, printer, POS terminal, IP phone, or other device critical to the business.

PAT provides security for local host traffic heading to the Internet through split tunneling. This outbound PAT on the VPN 3002 provides centralized security control because no configuration parameters exist for local users to adjust, which could possibly cause the central site to be compromised.

Standards Supported

To support fast, easy, and reliable deployment and scalability to thousands of sites, the Cisco VPN 3002 Hardware Client is a full-featured VPN client that incorporates IPSec and other industry standards. The 3002 support for the Cisco VPN Client Release 3.5 software, using the Unified Client Framework, enables it to connect to any Cisco central-site VPN Concentrator, including the Cisco 3000 Series VPN Concentrators, PIX Firewalls, and Cisco IOS–based routers.

The 3002 supports the following standards and protocols. The details of configuring these features are covered in Chapter 15.

- DHCP client and server services are supported. DHCP *client* implementation allows the public interface to be assigned an IP address from the head-end device on first connection. This is both easier and more reliable than end-to-end statically assigned addresses, which are typically required for LAN-to-LAN devices. DHCP *server* support allows the private interface(s) to assign IP addresses to up to 253 stations behind the Cisco VPN 3002.
- Three methods of NAT Transparent IPSEC, including the UDP method implemented in the original product release, IPSec/TCP method, and ratified IPSec/UDP NAT-T specification, which includes autodetection and fragmentation avoidance.
- PAT can be configured on the 3002 to hide the stations behind the Cisco VPN 3002 private interface(s) from external view and attack.
- IPSec encryption protocols, including 56-bit DES or 168-bit Triple DES for securing the data transmissions.
- MD5, SHA-1, HMAC with MD5, and HMAC with SHA-1 authentication algorithms.
- IPSec tunneling protocol with Internet Key Encryption (IKE) key management.
- AAA RADIUS accounting and security from the central site.
- H.323 support allows users to host and access NetMeeting sessions or to access other H.323 applications, such as voice-over IP (VoIP).
- Embedded web management interface accessible via local web browser, Secure Shell (SSH)/Secure Socket Layer (SSL), or conventional console port.
- SNMP MIB-II for monitoring, configuration, and event logging.

Cisco VPN 3002 Hardware Client Features

The following summarizes the features and benefits provided by the Cisco VPN 3002 Hardware Client devices. Those requiring configuration are addressed in Chapter 15, when configuring the client is covered or, because many features are "pushed" down from the central site, they're enabled and configured in Chapter 14.

Auto Upgrade

The client update feature was added in version 3.0 for the VPN 3002 Hardware Client and version 3.1 for the Cisco VPN software client. If the central device supports the feature (v3.0 for VPN Concentrators), the central device can be used to upgrade the software and configuration on the client. In the case of the VPN 3002 Hardware Client, firmware upgrades can also be pushed down to the client.

For VPN 3002 Hardware Clients, the *client update* allows administrators to update software and firmware automatically for the 3002 device. If an upgrade is needed, the unit upgrades automatically from an internal TFTP server specified on the central site VPN Concentrator. The process of maintaining security, managing the system, and upgrading it is transparent to the end user.

For Cisco VPN software clients the process is a little less automatic. This is more of a notification mechanism with an assisted upgrade. The client update for the Cisco VPN software clients allows central location administrators to notify the client users automatically when it's time to update. Then action is required on the part of users to retrieve and install the newer software.

Authentication Features

The VPN 3002 supports the following two levels of client authentication mechanism that supplies a high level of security for both the VPN 3002 and the users behind the VPN 3002:

- Interactive Unit Authentication
- Individual User Authentication

The VPN 3002 *Interactive Unit Authentication* technology uses Saved or One Time Passwords to reauthenticate itself to the head-end device. With Saved passwords, the 3002 client device needn't reauthenticate if the tunnel cycles. With One Time passwords, the device must be reauthenticated each time the tunnel cycles. The VPN 3002 supports preshared secrets, digital certificates, and tokens for this authentication.

The VPN 3002 *Individual User Authentication* feature can be set to require each user behind the VPN device to authenticate before traversing the tunnel. This feature can require the users behind the 3002 to use preshared secrets or tokens to authenticate. The individual authentication can be used by itself or in conjunction with Interactive Unit Authentication to maximize security.

To simplify the process and make it as transparent as possible to the end users, this technology automatically intercepts any user attempting to traverse the VPN tunnel and redirects them to a browser page to authenticate. The user needn't initiate or remember to initiate the security authentication because this is done automatically. If a user is only attempting to access the Internet via split tunneling, that user isn't prompted to authenticate.

Load Balancing and Failover

The VPN 3002 hardware device (release 3.5) and the Cisco VPN software client (v3.0) both support Cisco's VPN 3000 load-balancing strategy. To implement load balancing, multiple concentrators are grouped together logically on the same private LAN-to-LAN

network in a virtual cluster. These VPN Concentrators can be configured to direct session traffic transparently to the least-loaded device, thus distributing the load among all devices. In addition to increasing efficient use of system resources, this strategy provides increased performance, high availability, and reliability.

The VPN 3002 supports up to 10 back-up concentrators, in case the primary location is down or otherwise unavailable. The 3002 cycles through each backup concentrator in order until it makes a successful connection, maximizing network availability to the client.

PPPoE Support

Point-to-Point Protocol over Ethernet (PPPoE) is a specification for connecting Ethernet users to the Internet using a common broadband medium, such as a DSL line, a cable modem, or a wireless device. Many ISPs now require PPPoE authentication for DSL or other access to their networks. The VPN 3002 supports PPPoE Client mode to access these networks. Users need only to authenticate to the PPPoE server the first time and VPN 3002 then authenticates for all the user's subsequent attempts.

Cisco VPN 3000 Concentrator Devices

The Cisco VPN 3000 Series Concentrator is a growing family of VPN devices specifically designed and built to provide fast, reliable, and secure remote access to organization network resources. These devices combine with Cisco VPN client software and hardware to incorporate high availability, high performance, and scalability, plus advanced encryption and authentication technologies to allow customers to implement the latest VPN technology, while protecting and persevering in their communications investments.

The VPN 3000 platform offers customer-upgradeable and field-swappable components to increase capacity dramatically, while maintaining the original device, rack space, and power requirements. Scalable Encryption Processing (SEP) modules can be added to the 3015 to 3060 model case to enable users to add capacity and throughput easily.

The Cisco VPN 3000 Concentrator series comes in several models to meet organization capacity requirements and applications. The platform includes models to support customers with 100 or fewer remote access users to large organizations with up to 10,000 simultaneous remote connections. The latest Cisco VPN Software Client is provided at no additional charge with unlimited distribution licensing with all versions of the Cisco VPN 3000 Concentrator.

The Cisco VPN 3000 Concentrator platform offers models that can be implemented in both redundant and nonredundant configurations, offering the organization flexibility in design and budget management. In addition, advanced routing support, such as OSPF, RIP, and NAT, are available.

Cisco VPN 3000 Concentrator Models

The VPN 3000 Concentrators create virtual private networks (VPNs) by creating secure connection across a TCP/IP network, such as the Internet, that allows remote end users to connect to the corporate network. The VPN Concentrator can be used to create single-user-to-LAN connections for connecting traveling employees or telecommuters,

plus LAN-to-LAN connections for connecting remote sites. The VPN 3000 family of VPN concentrators currently come in the following five models, offering support for a few secure connections up to 10,000 sessions:

- Cisco VPN 3005
- Cisco VPN 3015
- Cisco VPN 3030
- Cisco VPN 3060
- Cisco VPN 3080

The *3005* is a unique 1U rack-mountable nonupgradeable device. The others share a common 2U rack-mountable chassis that allows choices in number of power supplies and SEP units. These units support after purchase upgrade options that allow the VPN security architecture to grow with the organization.

The following features are common to all Cisco VPN 3000 Concentrator Models:

- 10/100 Ethernet autosensing interfaces:
  - Model 3005: Two interfaces
  - Models 3015-3080: Three interfaces
- Motorola PowerPC CPU
- SDRAM memory for normal operation
- Nonvolatile memory (NVRAM) for critical system parameters
- Flash memory for file management

NOTE While each 3000 series device still specifies a T1/E1 optional interface, a Cisco Systems announcement is of the end of sale of Cisco VPN 3000 WAN Interface products, including all T1 and E1 WAN modules for VPN 3000 Concentrator Series products.

Cisco VPN 3005 Concentrator

The *3005* is a fixed-configuration VPN platform designed for small-to-medium networks with bandwidth requirements up to full-duplex T1/E1 (4 Mbps maximum performance) and up to 100 simultaneous sessions. Figure 13-3 shows the front and rear views of a VPN 3005. The basic configuration and expansion capabilities include the following:

- Software-based encryption
- Single power supply
- Expansion capabilities:
  - Optional WAN interface module with dual T1/E1 ports

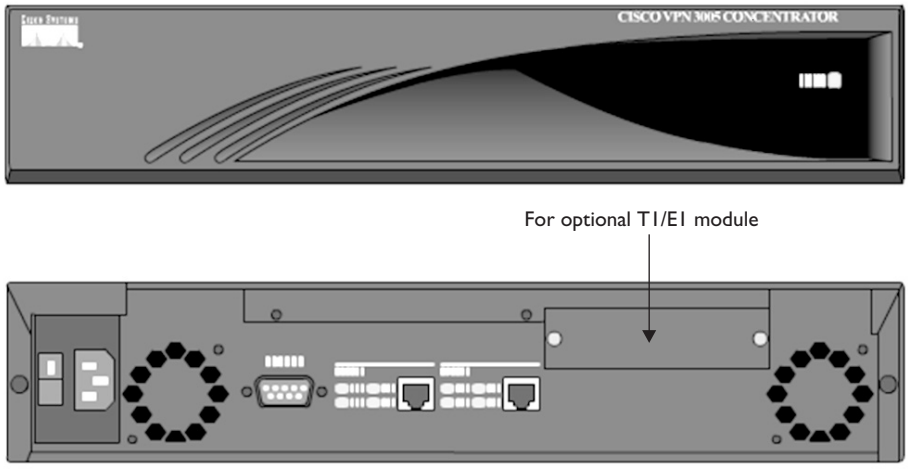

Figure 13-3 VPN 3005 front and rear views

Cisco VPN 3015 Concentrator

Like the Cisco VPN 3005, the *3015* is a VPN platform designed for small-to-medium networks with bandwidth requirements up to full-duplex T1/E1 (4 Mbps maximum performance) and up to 100 simultaneous sessions. Also, like the 3005, the default encryption processing is performed in software. Unlike the 3005, the 3015 is field-upgradeable to the VPN 3030 and 3060 models by adding memory and SEP modules. Figure 13-4 shows the front and rear view of a VPN 3015. The larger units look the same except that SEP modules and the redundant power supply replace the blank covers. The basic configuration and expansion capabilities include the following:

- Software-based encryption
- Single power supply
- Expansion capabilities:
  - Optional WAN interface module with dual T1/E1 ports
  - Up to four Cisco SEP hardware encryption modules
  - Optional redundant power supply

Cisco VPN 3030 Concentrator

The *3030* is a VPN platform designed for medium-to-large networks with bandwidth requirements from full T1/E1 through fractional T3, up to 50 Mbps maximum performance.

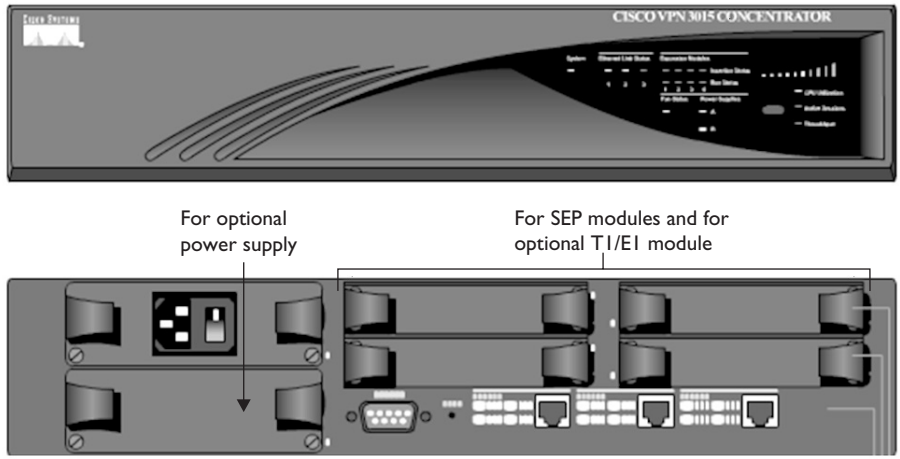

Figure 13-4 VPN 3015 front and rear views

The 3030 can support up to 1,500 simultaneous sessions. The basic configuration and expansion capabilities include the following:

- One SEP module for hardware-based encryption
- Single power supply
- Expansion capabilities:
  - Optional WAN interface module with dual T1/E1 ports
  - Up to three additional SEP hardware-based encryption modules
  - Optional redundant power supply

Cisco VPN 3060 Concentrator

The *3060* is a VPN platform designed for large networks requiring the highest level of performance and reliability, with high-bandwidth requirements from fractional T3 through full T3/E3 or greater connections. The 3060 can support up to 5,000 simultaneous sessions. The basic configuration and expansion capabilities include the following:

- Two SEP modules for hardware-based encryption
- Optional dual redundant power supplies (hot swappable)
- Expansion capabilities:
  - Optional WAN interface module with dual T1/E1 ports
  - Up to two additional SEP hardware-based encryption modules

Cisco VPN 3080 Concentrator

The *3080* is a top-of-the-line platform fully optimized to support large enterprise networks requiring the highest level of performance with support for up to 10,000 simultaneous remote access sessions. The basic configuration and expansion capabilities include the following:

- Four SEP modules for hardware-based encryption
- Dual redundant power supplies (hot swappable)
- Expansion capabilities:
  - Optional WAN interface module with dual T1/E1 ports

Side-by-Side Model Comparison

The following table summarizes the key features of the Cisco VPN 3000 Concentrator series of devices.

	3005	3015	3030	3060	3080
Simultaneous Users	100	100	1,500	5,000	10,000
Encryption Throughput	4 Mbps	4 Mbps	50 Mbps	100 Mbps	100 Mbps
Encryption Method	Software	Software	Hardware	Hardware	Hardware
Encryption (SEP) Module	0	0	1	2	4
Expansion Slots Available	0	4	3	2	0
Redundant SEP	N/A	N/A	Option	Option	Yes
System Memory	32MB fixed	64MB	128MB	256MB	256MB
T1 WAN Module	Fixed Option	Option	Option	Option	Option
Dual Power Supply	Single	Option	Option	Option	Yes
Client License	Unlimited	Unlimited	Unlimited	Unlimited	Unlimited
Hardware	1U, Fixed	2U, Scalable	2U, Scalable	2U, Scalable	2U

Standards Supported

To support fast, easy, and reliable deployment and scalability to thousands of remote users and sites, the Cisco VPN 3000 Concentrators are full-featured VPN devices that incorporate IPSec and other industry standards. The 3000 series support the following standards and protocols. Any details for configuring these features are covered in Chapter 14.

- Tunneling Protocols: IPSec, PPTP, L2TP, and L2TP/IPSec
- Encryption/Authentication: IPSec Encapsulating Security Payload (ESP) using DES/3DES (56/168 bit) or AES (128, 192, 256 bit) with MD5 or SHA, MPPE using 40/128 bit RC4
- Key Management: Internet Key Exchange (IKE) and Diffie–Hellman (DH) Groups 1, 2, 5, 7 (ECDH)

- NAT: NAT Transparent IPSec, IPSec/TCP, Ratified IPSec/UDP (with autodetection and fragmentation avoidance). Ratified IPSec/UDP support for NAT-T provides autodetection behind a NAT/PAT device, such as a small or home office router, and adds multivendor interoperability
- Routing: RIP, RIP2, OSPF, reverse route injection (RRI), static, automatic endpoint discovery, classless interdomain routing (CIDR)
- Release 3.6 includes DHCP relay/proxy for customers using the Cisco VPN 3000 Concentrator as an edge device in wireless configurations because it removes the need for an additional DHCP server
- Dynamic Domain Naming System (DNS) population (DDNS/DHCP) allows administrators to associate a remote access computer with its current IP address
- IPSec fragmentation policy control, including support for Path MTU Discovery (PMTUD)
- PPPoE Automatic Maximum Transmission Unit (MTU) adjustment for the network driver interface specification WAN (NDISWAN) during install improves remote access client operation in PPPoE DSL environments
- MovianVPN (Certicom) Handheld VPN Client with ECC

Cisco VPN 3000 Concentrator Features

The following summarizes the features and benefits provided by the Cisco VPN 3000 Concentrator devices. Chapter 14 addresses those that require configuration.

Modular Design (Models 3015 to 3080)

The Cisco SEP modules provide hardware-based encryption, ensuring consistent performance throughout the rated capacity for models 3030 through 3080. With multiple SEP modules, the devices became distributed-processing architecture, providing enhanced performance and increased reliability through redundancy. This modular design provides investment protection, redundancy, and a simple upgrade path, plus it minimizes the impact on rack space and power supply allocation.

Digital Design

The all-digital design of the VPN 3000 device provides high-degree reliability with solid, long-term performance, while providing 24-hour continuous operation. Incorporated into each unit is a robust instrumentation package for real-time monitoring and alerts.

Windows Compatibility

The VPN 3000 series' close support for Microsoft hosts, including Windows 2000/XP clients, makes large-scale client deployment and seamless integration with related network systems. The VPN 3000 supports the following Microsoft protocols:

- Microsoft PPTP/MPPE/MPPC, MSCHAPv1/v2, and EAP/RADIUS pass-through for EAP/TLS and EAP/GTC support

- Microsoft L2TP/IPsec for Windows 2000/XP (including XP DHCP option for route population)
- Microsoft L2TP/IPsec for Windows 98, Windows Millennium (Me), and Windows NT Workstation 4.0

VPN 3000 Release 3.6 added three improvements to support Microsoft's Integrated VPN Client including

- Microsoft L2TP/IPSec Extensible Authentication Protocol (EAP) pass-through support (TLS and GTX/SDI) for working from behind a PAT/NAT device with the VPN Client
- DHCP—XP route list population (split tunneling)
- IPSec/User Datagram Protocol (UDP) NAT-T compatibility (expected release by Microsoft in 2003)
- Support for Windows Installer (MSI) installation (Windows NT/20000XP only), providing the system administrator with the capability to customize installation packages and track system changes made during client installation

Security

The VPN 3000 Series support for current and emerging security standards, including RADIUS, NT Domain Authentication, RSA SecurID, one-time passwords (OTP), and digital certificates offering large-scale client deployment and seamless integration with external authentication systems, as well as interoperability with third-party products.

VPN 3000 release 3.6 offers two notable enhancements to concentrator encryption and security, including

- Advanced Encryption Standard (AES) addition to the concentrator offers a stronger encryption option and provides performance benefits for both the Cisco VPN 3002 Hardware Client and the Cisco VPN Client.
- RSA SecurID (SDI) Version 5.0 support. Users can now take advantage of the load balancing and resiliency features found in the RSA SecurID Version 5.0.

Advanced packet-filtering capabilities provide additional network security. Filtering options include source and destination IP address (Layer 3), port and protocol type (Layer 4), fragment protection, time and day access control, and FTP session filtering.

User and group-level policy management can be implemented for maximum flexibility and granularity in controlling network and feature access control.

High Availability

The VPN 3000 Series' redundant subsystems and multichassis failover capabilities help to ensure maximum system uptime and remote user connectivity. Redundant SEP and power supply options within individual devices promote reliability in a single or multidevice configuration. Multiple concentrators can be configured for both load-balancing and failover redundancy, providing protection and capacity to high-volume critical systems.

Extensive instrumentation and monitoring capabilities, as well as support for Cisco network management software applications, provide network managers with real-time system status and early-warning alerts.

A new feature in Release 3.6 is improved bandwidth limiting and traffic-shaping capabilities on the Cisco VPN 3000 Concentrators. This allows network administrators to assign minimum and maximum bandwidth parameters on a per-user basis. The administrator can establish limits on users with high-bandwidth use.

Robust Management

The Cisco VPN 3000 Concentrator can be managed using web-based applications from any standard web browser using HTTP or HTTPS. The VPN 3000s also support CLI commands using Telnet, Secure HTTP, SSH, and via a console port.

The VPN concentrator devices support configuration and monitoring capabilities for both the enterprise user and the service provider.

VPN concentrator device access levels can be configured per user and/or per group allowing configuration and maintenance control consistent with the organization security policies.

Monitoring and Logging

The Cisco VPN 3000 Concentrators support the following technologies for providing monitoring and logging services:

- Syslog output
- Configurable SNMP traps
- Event logging and notification via e-mail (SMTP) and, therefore, pager
- Automatic FTP backup of event logs
- SNMP MIB-II support
  General Statistics
  System Status
  Session Data (including Client Assigned IP, Encrypted Type Connection Duration, Client OS, Version, and so forth)

VPN 3000 Concentrator Client Support

Remote access VPN clients use the three following common connectivity techniques to reach the central site:

- VPN client software installed on PCs or workstations
- Hardware VPN routers
- Firewalls and hardware clients

Cisco Easy VPN could be the perfect solution for all three techniques, particularly with large implementations and limited local support.

Cisco Easy VPN

The *Cisco Easy VPN* is a software enhancement for existing Cisco routers and security appliances that can simplify VPN deployment for remote offices and telecommuters. Easy VPN is based on the Cisco Unified Client Framework using a centrally located Easy VPN Server, which is configured with all parameters required of remote device. The remote Easy VPN client can be preconfigured for mass deployments and initial logins require little user intervention. The full client configuration is "pushed" down from the Easy VPN Server when the client connects.

Cisco Easy VPN centralizes VPN configuration and management, thereby reducing the complexity of VPN deployments. The Cisco Easy VPN strategy incorporates all Cisco VPN client implementations into a single deployment, including IOS routers, PIX Firewalls, the VPN 3002 Hardware Client, and software VPN clients. This system offers consistent policy and key management methods, thus simplifying remote side setup and administration.

Using this feature, security policies defined and updated at the head-end are pushed to the remote VPN client, ensuring those connections have current policies in place before any connection is established. In addition, a Cisco Easy VPN Server-enabled device can terminate VPN tunnels initiated by mobile remote workers running Cisco VPN client software on PCs. This flexibility makes it possible for traveling workers or telecommuters to access the corporate intranet for critical data and applications.

The Cisco Easy VPN client on a hardware device supports both the VPN Client and Network Extension modes discussed earlier.

The Cisco Easy VPN client feature currently supports the following hardware platforms. It might be necessary to upgrade the IOS on older devices to have the feature. Be sure to check the Cisco web site to see if other models have been added.

- Cisco 806, Cisco 826, Cisco 827, and Cisco 828 routers
- Cisco uBR905 and Cisco uBR925 cable access routers
- Cisco 1700 series routers
- PIX 500 Series Firewalls
- Cisco VPN 3002 Hardware Client

The Cisco Easy VPN Server feature supports the following hardware platforms. As with the client, it might be necessary to upgrade the IOS have the feature. Be sure to check the Cisco web site to see if other models have been added.

- Cisco VPN 3000 Concentrators
- PIX 500 Series Firewalls
- Most Cisco IOS routers

The Cisco Easy VPN Remote feature provides for automatic management of the following details:

- Negotiating tunnel parameters—Addresses, algorithms, lifetimes, and so on
- Establishing tunnels according to the defined parameters

- Automatically creating the NAT/PAT translation and any associated access lists
- Authenticating users—Making sure users are who they say they are, by way of user names, group names, and passwords
- Managing security keys for encryption and decryption
- Authenticating, encrypting, and decrypting data through the tunnel
- Client software upgrades for the VPN 3002

The Cisco VPN Client provides support for Windows 95, 98, Me, NT 4.0, 2000, XP, Linux (Intel), Solaris (UltraSparc-32bit), and MAC OS X 10.1, including centralized split-tunneling control and data compression. VPN client configuration was covered in Chapter 12.

Cisco VPN 3002 Hardware Client

The Cisco VPN 3002 Hardware Client was designed for organizations with many remote office environments. The *3002* combines the ease of use and scalability of a software client with the reliability and stability of a hardware platform. The 3002 client supports Easy VPN Remote, allowing it to connect to any Easy VPN server site concentrator. The VPN 3002 Hardware Client works invisibly with any OS supporting IP, including Solaris, Mac, and Linux.

The VPN 3002 is available with or without a built-in 8-port switch and allows up to 253 station connections in a single network.

Release 3.6 included two significant feature enhancements for the VPN 3002 Hardware Client device:

- Software-based AES providing an enhanced security option through stronger encryption capabilities. As with the Cisco VPN Client, enhanced remote access performance also exists on the Cisco VPN 3002 Hardware Client.
- H.323 Fixup feature allows remote access users—in Client mode—behind the Cisco VPN 3002 Hardware Client, to use NetMeeting or other H.323 applications. H.323 requires no configuration on either the VPN Concentrator or the VPN 3002.

Wireless Client Support

With release 3.0, all Cisco VPN 3000 Concentrators support ECC. This is a new Diffie–Hellman Group, which allows for much faster processing of keying information by devices with limited processing power, such as PDAs and smart phones. Cisco VPN 3000 Concentrators can now securely terminate tunnels from IP-enabled wireless devices, allowing a whole new class of users to access enterprise information securely, while preserving the investment in VPN termination equipment in the enterprise data center.

Cisco Internet Mobile Office

The Cisco VPN 3000 Series Concentrator, Cisco VPN 3002 Hardware Client, and the Cisco VPN Client work together with the Cisco Internet Mobile Office to provide mobile professionals with secure, high-speed broadband connectivity to their networks in airports, convention centers, hotels, and a growing number of other public spaces.

Chapter Review

This chapter looked at the various hardware implementations for Cisco VPN technologies and focused mainly on the VPN 3002 Hardware Client and the VPN 3000 Series Concentrators. The VPN 3002 is typically implemented at remote sites in larger organizations. They can be augmented by Cisco IOS routers, PIX Firewalls, and VPN client software running directly on the host PC.

The 3002 comes in two models: a two-port unit that can support a single client or pass through to a hub or switch and a model with a built-in 8-port 10/100 switch. Both devices can support up to 253 users on the LAN interface and a maximum of 100 simultaneous secure connections back to the central network.

The Cisco VPN 3002 supports two modes of operation to offer implementation choices based on flexibility, security, and easy configuration: Client mode and Network Extension mode. A large VPN implementation might frequently have both types of operation. In Client mode, the VPN 3002 emulates the VPN client software appearing to the main network like a single remote user. The hosts protected behind the VPN 3002 are a separate network that remains invisible and nonroutable to the central site. In Network Extension mode, the VPN 3002 establishes a secure site-to-site connection with the central site device. The local stations behind the VPN 3002 are fully routable and the local network is visible to the central site.

The VPN 3002 supports a growing range of VPN standards and technologies, plus some implementation features to simplify large VPN implementation and support. These features include Easy VPN Client, which allows a thin installation on the 3002 and the final configuration pushed down to the 3002 on first connection to the VPN Concentrator. The 3002 also supports Auto Upgrade, which allows the VPN Concentrator to push any needed software or firmware upgrades down to the client.

The VPN 3002 supports the following two levels of client authentication mechanism that supplies a high-level of security for both the VPN 3002 and the users behind the VPN 3002: Interactive Unit Authentication authenticates the VPN device during the VPN setup, while Individual User Authentication requires each user behind the 3002 to authenticate before using the VPN tunnel.

The VPN 3002 hardware device (release 3.5) and the Cisco VPN software client (v3.0) both support Cisco's VPN 3000 load-balancing and failover strategies that allow for more efficient use of the Concentrators and provide backup alternatives if a Concentrator fails.

The Cisco VPN 3000 Series Concentrator is a growing family of VPN devices designed and built to provide fast, reliable, and secure remote access to organization network resources. These devices work with the Cisco VPN client software and various Cisco VPN hardware devices to incorporate high availability, high performance, and scalability, plus advanced encryption and authentication technologies to the network.

The VPN 3000 Concentrator platform offers customer-upgradeable and field-swappable components to increase capacity dramatically, while maintaining the original device, rack space, and power requirements. Scalable Encryption Processing (SEP) modules can be added to the 3015 to 3060 model case to enable users to add capacity and throughput easily.

The Cisco VPN 3000 Concentrator series comes in several models to meet organization capacity requirements and applications. The platform includes models to support customers with 100 or fewer remote access users to large organizations with up to 10,000 simultaneous remote connections. The latest Cisco VPN Software Client is provided at no additional charge with unlimited distribution licensing with all versions of the Cisco VPN 3000 Concentrator.

Questions

1. What is the protocol the Cisco VPN 3000 Series Concentrators use to provide Wireless Client Support for personal digital assistants (PDAs) and smart phones?
   - **A.** H.323
   - **B.** Elliptic Curve Cryptography (ECC)
   - **C.** IPSec/UDP NAT-T
   - **D.** PPPoE
2. The VPN 3002 Hardware Client supports how many simultaneous VPN connections?
   - **A.** 10
   - **B.** 50
   - **C.** 100
   - **D.** 253
3. The Cisco VPN 3002 supports which two modes of operation?
   - **A.** Split Tunneling mode
   - **B.** Client mode
   - **C.** DHCP mode
   - **D.** Network Extension mode
4. The VPN 3002 DHCP "client" service is implemented on which interface?
   - **A.** Private
   - **B.** DMZ
   - **C.** Public
   - **D.** It is user configurable on any interface
5. The VPN 3002 DHCP "server" service is implemented on which interface?
   - **A.** Private
   - **B.** DMZ
   - **C.** Public
   - **D.** It is user configurable on any interface

6. Which one of the following is *not* a VPN 3002–supported feature?
   - **A.** Auto Upgrade
   - **B.** Interactive Unit Authentication
   - **C.** PPPoE
   - **D.** Individual User Authentication
   - **E.** Wireless client
7. Which is *not* a model of VPN 3000 Concentrator?
   - **A.** 3005
   - **B.** 3010
   - **C.** 3030
   - **D.** 3060
   - **E.** 3080
8. Which model of VPN 3000 Concentrator supports only two 10/100 interfaces?
   - **A.** 3005
   - **B.** 3010
   - **C.** 3030
   - **D.** 3060
   - **E.** 3080
9. What is the maximum number of simultaneous VPN connections supported by any VPN 3000 Concentrator?
   - **A.** 1,000
   - **B.** 5,000
   - **C.** 10,000
   - **D.** 100,000
   - **E.** No limit
10. What is the maximum number of SEP modules that can be installed in a VPN 3000 Concentrator?
    - **A.** 1
    - **B.** 2
    - **C.** 3
    - **D.** 4
11. Which one is *not* a security protocol supported by the VPN 3000 series devices?
    - **A.** Diffie–Hellman (DH) Groups 1, 2, 5, 7 (ECDH)
    - **B.** RRI (reverse route injection)

C. AES (128, 192, 256 bit)

D. DES/3DES (56/168 bit)

12. Which one is *not* a common remote access VPN client connectivity technique to reach the central site?

A. VPN switch

B. Firewalls and hardware clients

C. Hardware VPN routers

D. VPN client software installed on PCs or workstations

13. With the Cisco Easy VPN strategy, where is the security policy configured?

A. Cisco Easy VPN Client

B. Cisco Easy VPN Server

C. Cisco VPN Client

D. NAS

14. The Cisco Internet Mobile Office provides mobile professionals with what service?

A. VPN connectivity while on airplanes

B. PDA connectivity within the network

C. VPN connectivity from home

D. Secure, high-speed broadband connectivity in public places

15. Which of the following can *not* be a Cisco Easy VPN server?

A. PIX Firewall

B. IOS Router

C. VPN 3000 Concentrator

D. VPN 3002 hardware device

Answers

1. **B.** Elliptic Curve Cryptography (ECC)
2. **C.** 100
3. **B.** Client mode. and D. Network Extension mode
4. **C.** Public
5. **A.** Private
6. **E.** Wireless client
7. **B.** 3010

8. **A.** 3005
9. **C.** 10,000
10. **D.** 4
11. **B.** RRI (reverse route injection)
12. **A.** VPN switch
13. **B.** Cisco Easy VPN Server
14. **D.** Secure, high-speed broadband connectivity in public places
15. **D.** VPN 3002 hardware device

CHAPTER 14

Cisco VPN 3000 Remote Access Networks

In this chapter, you will learn to:

- Describe VPN Concentrator user interfaces and startup
- Discuss VPN Concentrators in IPSec VPN implementations
- Configure VPN remote access with preshared keys
- Configure VPN remote access with digital certificates
- Configure VPN users and groups
- Configure Cisco VPN 3000 client support
- Configure the Cisco VPN client autoinitiation feature
- Monitor and administer Cisco VPN 3000 remote access networks

This chapter introduces working with the Cisco VPN 3000 Concentrators for basic operations, as well as for advanced features and options. The 3000 series devices were introduced and described in Chapter 13. That information won't be repeated here, but the features introduced are explored and the configuration steps defined.

The VPN 3000 devices, the concentrator series, and the remote client device all support the following three main activities:

- Configuration
- Administration
- Monitoring

This chapter looks at all three activities on the 3000 series Concentrators.

Basic VPN concepts, such as preshared keys and using Certificate Authorities to implement digital certificates, were covered in Chapters 9 through 11. While the features are configured in this chapter, the underlying technologies aren't addressed beyond an explanation for the processes being discussed. The menu-driven and web-based interfaces

used with the VPN 3000 devices change the implementation steps, but three primary tasks are still used to ensure a successful installation:

1. Prepare for IKE and IPSec
2. Configure the features
3. Verify and monitor configuration

The VPN Concentrator creates a virtual private network (VPN) by creating a secure connection across a public Transmission Control Protocol/Internet Protocol (TCP/IP) network, such as the Internet. It can create single user to local area network (LAN) (remote user) connections and LAN-to-LAN connections.

The VPN Concentrator functions as a bidirectional tunnel endpoint using various standard protocols to perform the following tasks:

- Establish tunnels
- Negotiate tunnel parameters
- Authenticate users
- Assign user addresses
- Encrypt and decrypt data
- Manage security keys
- Manage data transfer across the tunnel

Network configurations and VPN placement can vary widely. Chapter 27 introduces the Cisco SAFE strategy and how VPN devices fit into the secure network design. The *VPN Concentrator* is a flexible and functional device that can satisfy most applications. This chapter discusses configuring the VPN 3000 Series Concentrator to support VPN remote access implications. Chapter 15 reviews the VPN 3002 remote access client device. Chapter 16 covers the VPN Series Concentrator to support VPN LAN-to-LAN networks.

VPN Concentrator User Interfaces and Startup

The VPN 3000 Concentrator and the VPN 3002 Hardware Client support both a specialized command-line interface (CLI) and a web-based interface (Concentrator or Client Manager). You can do exactly the same tasks with either interface. The choice ultimately boils down to personal preference but, in general, the Manager web-based interface might be easier to learn and navigate if you're familiar with the Windows interface and browsers. The Help feature on the Concentrator Manager is much better than the CLI.

To manage the device with something other than the console cable, the CLI must be used to configure at least the private interface to be part of a network accessible by one or more hosts with either a web browser or Telnet capabilities.

Quick Configuration

The Quick Configuration Wizard that appears at initial startup allows configuring the minimal parameters needed to make the VPN Concentrator functional. While optional, this wizard and a console connection provide a relatively quick and easy way to add enough configuration to allow access from a network PC. Quick Configuration can be performed entirely from the CLI or switched to the Concentrator Manager any time after the first two steps. The Quick Configuration options, as shown in the following output, appear only once and can only be used once, unless the device is rebooted with the Reboot with Factory/Default option.

```
Login: admin
Password:                                (doesn't display)

               Welcome to
              Cisco Systems
       VPN 3000 Concentrator Series
          Command Line Interface
Copyright (C) 1998-2003 Cisco Systems, Inc.
-- : Set the time on your device. ...
> Time
Quick -> [ 10:13:37 ]
```


CLI Quick Configuration Steps

The CLI Quick Configuration Wizard prompts guide you through the following configurations steps:

1. Set the system time, date, time zone, and daylight saving time values.
2. Configure the VPN Concentrator private network interface (Ethernet 1) by responding to the following prompts:

```
This table shows current IP addresses.
    Interface             IP Address/Subnet Mask    MAC Address
-------------------------------------------------------------------
| Ethernet 1 - Private  |    0.0.0.0/0.0.0.0     |
| Ethernet 2 - Public   |    0.0.0.0/0.0.0.0     |
| Ethernet 3 - External |    0.0.0.0/0.0.0.0     |
-------------------------------------------------------------------
** An address is required for the private interface. **
> Enter IP Address
Quick Ethernet 1 -> [ 0.0.0.0 ] 192.168.1.1

> Enter Subnet Mask
Quick Ethernet 1 -> [ 255.0.0.0 ] 255.255.255.0

1) Ethernet Speed 10 Mbps
2) Ethernet Speed 100 Mbps
3) Ethernet Speed 10/100 Mbps Auto Detect
Quick -> [ 3 ]

1) Enter Duplex - Half/Full/Auto
2) Enter Duplex - Full Duplex
3) Enter Duplex - Half Duplex
Quick -> [ 1 ]
```

Once these settings are made, the remainder of the Quick Configuration can be completed with the VPN Concentrator Manager (Cisco's recommended method). If the CLI method is used, similar prompts would guide the configuration through the following features:

1. Configure any other interfaces. At a minimum, the Ethernet 2 public interface must be configured. An additional Ethernet 3 interface could be on models 3015-3080 and/or optional WAN interfaces.
2. Define information that identifies your VPN Concentrator on the network, such as the system name, IP address of the DNS, registered Internet domain name, and default gateway to which the VPN will forward unknown packets.
3. Define which tunneling protocols and encryption options are to be used.
4. Define the method(s) for assigning IP addresses to protected clients on the private interface as the defined tunnel is established.
5. Specify one of five types of servers to authenticate users: the concentrator's internal server, the external RADIUS server, the external NT Domain server, external server, or the external SDI (RSA Security Inc. SecurID) server, or Kerberos/Active Directory server.
6. When using the VPN Concentrator internal authentication server, populate the internal user database with at least one user, each with a user name and password, and, if *per-user address assignment* is specified, an IP address and subnet mask.

NOTE The maximum number of entries (groups and users combined) varies by model:
Model 3005/3015–100
Model 3030–500
Model 3060/3080–1,000

7. When using IPSec tunneling protocol, the remote-access client connects to the VPN Concentrator via a group name and password, which needs to be configured.
8. Change the admin password to improve system security.
9. Save the configuration file (menu option) to complete quick configuration.

Concentrator Manager Quick Configuration

The Concentrator Manager Quick Configuration assumes the first two steps of the CLI Quick Configuration were performed, so the private interface can be accessed by a web browser. As with the CLI version, you can only go through the quick configuration steps one time, unless the device is rebooted ignoring the Configuration File option.

The web browser session is opened using the admin/admin case-sensitive user name/password combination. The one-time screen, as shown in Figure 14-1, allows using the Quick Configuration Wizard to add additional settings.

The wizard then leads you through the same configuration options offered earlier in the CLI section. While they aren't identical, the Quick Configuration Wizard is similar to the one used by the VPN 3002 Hardware Client device, covered in detail in Chapter 15.

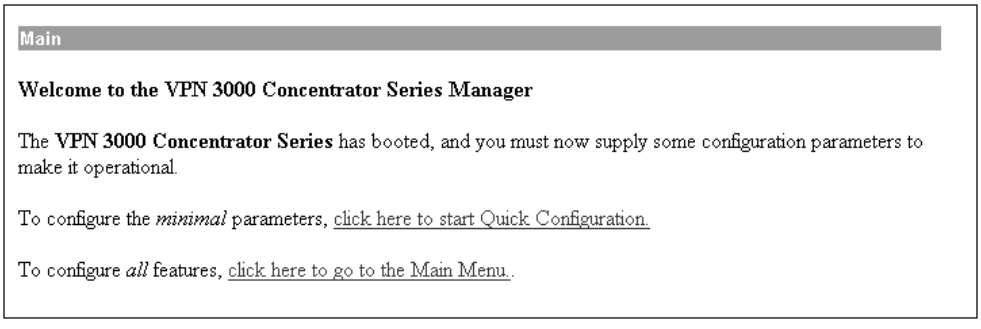

Figure 14-1 VPN Concentrator Manager Quick Configuration option screen

The example about preshared keys in the section "Remote Access VPNs with Preshared Keys" assumes the Quick Configuration wasn't done, and it covers the screens and settings at that time.

STUDY TIP VPN, PIX, and IDS technologies, features, and versions are quickly changing and improving. In some cases, three versions were released in the past year. Because the Cisco web site no longer specifies the versions covered, confirming as many processes as possible with the online documentation makes sense. For the latest information on the Quick Configuration Wizards, go to www.cisco.com and search for VPN Quick Configuration. Then look for the most recent version of the Getting Started Guide. This guide has step-by-step instructions with the most current screen features. Similar searches on any technologies that you can't work with hands-on would be good way to avoid being surprised in this rapidly changing environment.

Command-Line Interface (CLI) Basics

The VPN 3000 Concentrator CLI is a built-in, menu-driven configuration, administration, and monitoring system, which can be accessed via the device console port or a Telnet (or Telnet over SSL) session. Both Telnet options are enabled by default on the private network interface once the interface is configured with an IP address in the LAN. The CLI supports the same configuration options as the HTML-based VPN 3000 Concentrator Manager covered in the section "Concentrator Manager (Web Interface)."

NOTE The VPN 3000 concentrators use a straight-through serial connection. A serial cable with 9-pin adapter can be used on the concentrator end, and then whatever connection is required for the serial interface on the PC. To use a straight-through jumper cable, one of the Cisco 9-pin adapters (gray posts) from a Cisco console configuration kit is required on the concentrator. The VPN 3002 uses a standard Cisco console kit with a roll-over cable that plugs into an RJ-45 interface on the device.

Console port access is similar to the IOS routers using a terminal emulator program, such as HyperTerminal. Pressing ENTER few times might be necessary to get the login prompt to appear. Login user names and passwords for both console and Telnet access are the same. The factory-supplied default is configured and enabled for administrators using admin for both the login and password. Entries are case-sensitive. The following output shows the initial login and the main menu:

```
Login: admin
Password:                              (doesn't display)

               Welcome to
              Cisco Systems
       VPN 3000 Concentrator Series
          Command Line Interface
Copyright (C) 1998-2003 Cisco Systems, Inc.

1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main ->
```

Making changes typically involves making one or more menu choices, and then answering any appropriate prompts. The following are some things to remember about the CLI interface:

- Password entries are case-sensitive.
- The interface displays current or default entries in brackets, for example, [192.168.1.1].
- Use the ENTER key on the console keyboard to complete a choice.
- Any configuration changes take effect as soon as they're entered. These changes are part of the active, or running, configuration. To make this the boot configuration used when the device is rebooted, you must save the configuration (Main menu option #4).

Help

The Help menu system is somewhat limited, displaying only the following information when 5 is entered at the Main menu. Context-sensitive Help is unavailable and the familiar question mark (?) doesn't activate help features. The Help feature in the Hardware Client Manager is much better and offers context-sensitive assistance like most Windows applications.

```
Main -> 5
Cisco Systems.  Help information for the Command Line Interface
```

```
From any menu except the Main menu.
-- 'B' or 'b' for Back to previous menu.
-- 'H' or 'h' for Home back to the main menu.

For Data entry
-- Current values are in '[ ]'s. Just hit 'Enter' to accept value.

1) View Help Again
2) Back

Help ->
```

The *B* and *H* options (and especially the *H* option) come in handy when you're navigating the device menus. Pressing *H* returns you to the Main menu.

Saving Configuration File Changes

Configuration and administration changes made using Options 1 and 2 on the Main menu take effect immediately and become a part of the active, or running, configuration. Like the Cisco routers, if the VPN 3000 is rebooted without saving the active configuration, any changes will be lost.

Saving changes to the system configuration (CONFIG) file is a one-step process from the Main menu. At the Main -> prompt, typing the numeral 4 will save changes without additional steps or confirmation.

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 4
```

The system writes the current (active) configuration to the CONFIG file and redisplays the main menu.

Configuring the Private Interface

To access the Concentrator using the Manager web interface, you must have the workstation in the same subnet at the private interface. If the Quick Configuration Wizard wasn't used for a new or unconfigured Concentrator, the CLI is used to configure the interface. The following output shows the steps from the Main menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 1
```

```
1) Interface Configuration
2) System Management
3) User Management
4) Policy Management
5) Back

Config -> 1                                (shows the current settings)

This table shows current IP addresses.
  Intf         Status       IP Address/Subnet Mask          MAC Address
--------------------------------------------------------------------------
Ether1-Pri|      UP      | 192.168.10.1/255.255.255.0 | 00.03.A0.88.CE.AC
Ether2-Pub|Not Configured|      0.0.0.0/0.0.0.0       |
--------------------------------------------------------------------------
DNS Server(s): DNS Server Not Configured
DNS Domain Name:
Default Gateway: Default Gateway Not Configured

1) Configure Ethernet #1 (Private)
2) Configure Ethernet #2 (Public)
3) Configure Power Supplies
4) Back

Interfaces -> 1                          (to set the Private interface)

1) Interface Setting (Disable, DHCP or Static IP)
2) Set Public Interface
3) Select IP Filter
4) Select Ethernet Speed
5) Select Duplex
6) Set MTU
7) Set Port Routing Config
8) Set Bandwidth Management
9) Set Public Interface IPSec Fragmentation Policy
10) Back

Ethernet Interface 1 -> 1                  (to set a Static IP address)

1) Disable
2) Enable using DHCP Client
3) Enable using Static IP Addressing

Ethernet Interface 1 -> [ 3 ] 3

> Enter IP Address                  (current value appears in brackets)
Ethernet Interface 1 -> [ 192.168.10.1 ] 192.168.1.10

> Enter Subnet Mask
Ethernet Interface 1 -> [ 255.255.255.0 ]
Ethernet Interface 1 -> h                          (return to Main menu)
```

Concentrator Manager (Web Interface)

The VPN 3000 Concentrator Manager (Manager) is an HTML-based interface application that makes configuring, administering, monitoring, and managing the VPN 3000 device possible with a web browser on a PC in the private network.

By default, the Manager uses HTTP, which is convenient, but messages are transmitted in Cleartext. If security requires it, the Manager supports secure, encrypted HTTP connection over Secure Sockets Layer (SSL) protocol, known as HTTPS.

Browser Requirements

The Manager application supports either Microsoft Internet Explorer (IE) version 4.0 or higher or Netscape Navigator version 4.5–4.7. For the best results, Cisco recommends Internet Explorer. JavaScript and Cookies need to be enabled in the browser. Another recommendation is to install any updates and patches.

Recommended Display Settings

Cisco recommends the following monitor display settings for the best viewing:

- Screen area—1024 × 768 pixels or greater. (Minimum 800 × 600 pixels)
- Colors—256 colors or higher.

Browser Navigation Toolbar

Earlier implementations of the Manager were the CLI converted simply to a web interface. Each new version includes better Windows function integration. In particular, Help, a Java-based applet, is getting friendlier and more useful.

Cisco doesn't recommend using the browser navigation toolbar buttons Back, Forward, or Refresh/Reload with the Manager, unless you're specifically instructed to do so. To maintain access security, clicking the Refresh/Reload button automatically logs out the Manager session and returns to the login screen. Using the Back or Forward buttons could possibly display old Manager displays with incorrect data or settings. If you're concerned, the IE View | Full screen (F11) feature will eliminate the temptation.

Connecting to the Concentrator Manager

To access the VPN 3000 Concentrator Manager application using HTTP over a web browser, type the VPN 3000 private interface IP address (such as 192.168.1.10) in the browser Address or Location field. The browser automatically supplies the http:// prefix.

The browser displays the login screen shown in Figure 14-2.

Logging in to the Manager application is the same for Cleartext HTTP or secure HTTPS. At this point, a valid user name/password combination can be entered to gain access. Both entries are case-sensitive. Internet Explorer users can use the TAB key to move from field to field. The Clear button can be used to start over.

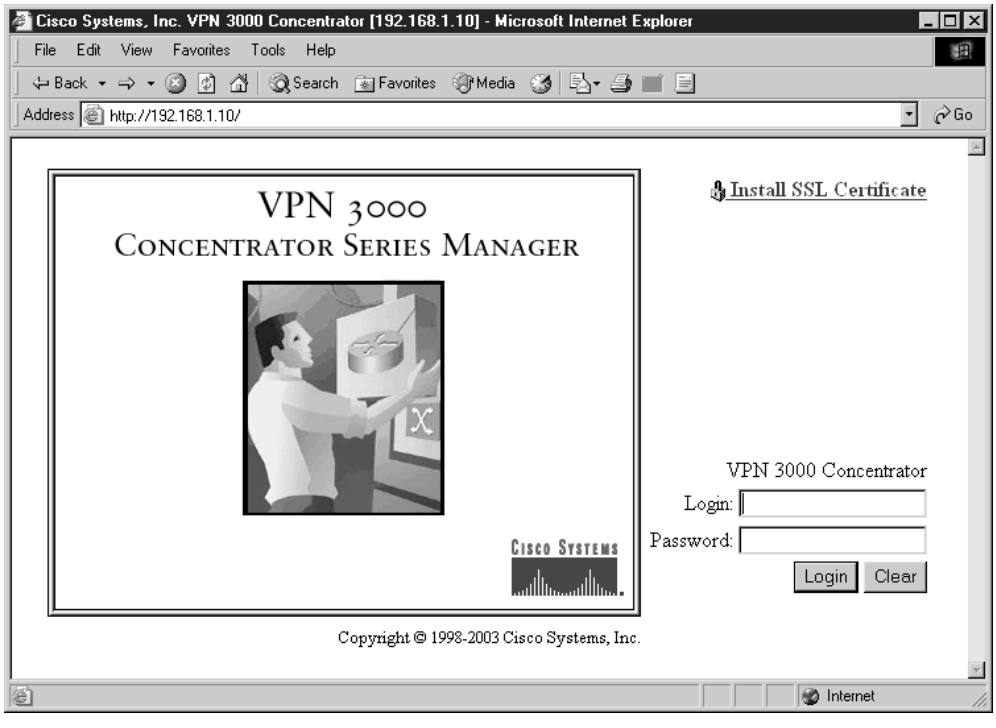

Figure 14-2 VPN 3000 Concentrator Manager login screen

An unconfigured unit will have the default Administrator combination admin/admin for initial entry.

Figure 14-3 shows the opening screen that appears, offering access to the three main application modules: Configuration, Administration, and Monitoring. The application tree on the left-hand side offers explorer-like navigation capabilities to move quickly from feature to feature. The menu tree can be expanded and contracted as needed. The figure shows the Configuration section expanded to show the second-tier choices. The same three choices are displayed in the upper-right corner for times when the left-side panel isn't displayed. The upper-right corner also offers quick access to the Main menu and the Help system, as well as support phone numbers and web pages.

The main body of the window provides a good overview of the various screen components and options of the application, including explanations for service icons that will appear in the upper-right corner.

The VPN 3000 Concentrator Reference, available online or on the documentation CD-ROM that came with the device, covers how to set up the device for installing an SSL Certificate in the browser for HTTPS connectivity.

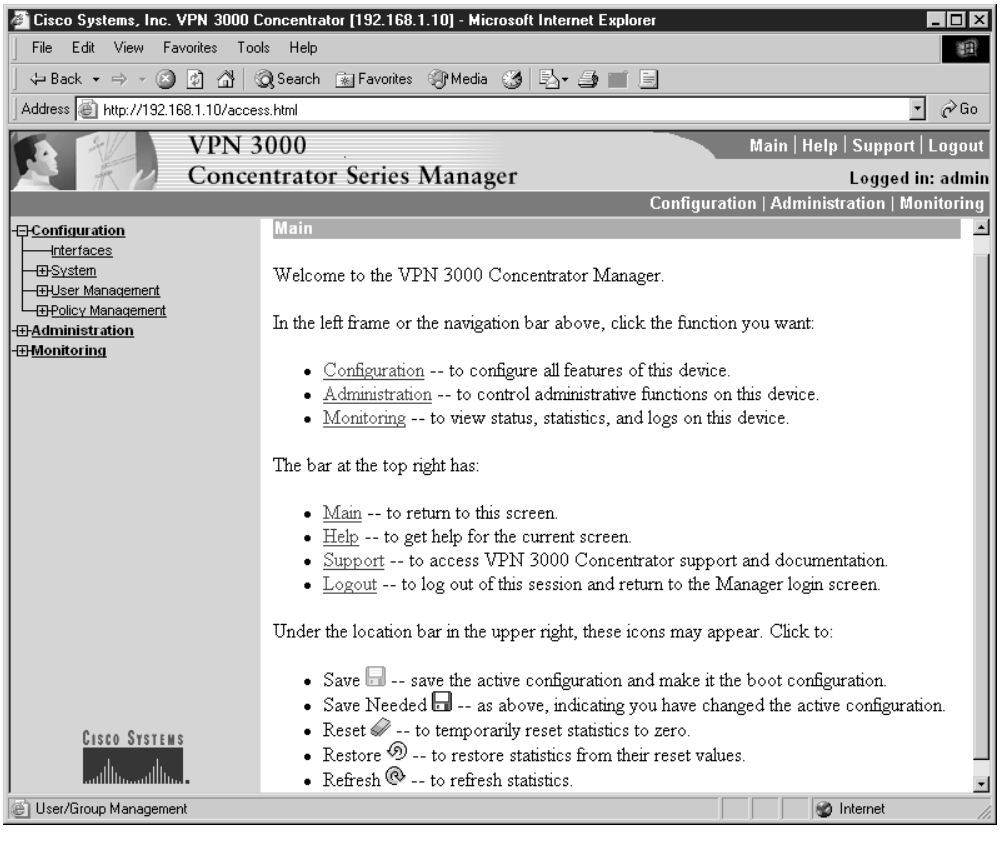

Figure 14-3 Manager opening screen

Manager Organization

The Manager, exactly like the CLI, is made up of three major sections and many second- and third-level subsections:

- **Configuration**
- **Administration**
- **Monitoring**

Configuration The Configuration menu is used to set all parameters that govern the unit's use and functionality as a VPN device. Cisco supplies default parameters that cover typical installations and uses. Figure 14-4 shows the Configuration menu fully expanded. The insert graphic shows the four second-tier options.

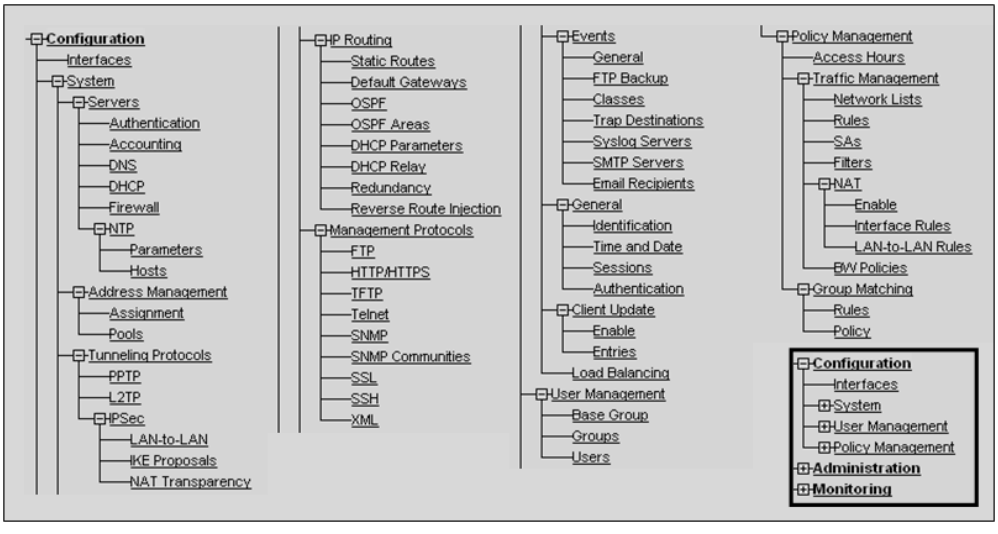

Figure 14-4 Configuration menu fully expanded

The Configuration section provides access to configure all VPN 3000 Concentrator features:

- **Interfaces**—Ethernet interfaces, DNS servers, domain name, and power supplies
- **System**—system-wide parameters: servers, address assignment, tunneling protocols, IP routing, IPSec, management protocols, events, identification, and the Client autoupdate feature
- **User Management**—create and modify groups and users
- **Policy Management**—access hours, network lists, rules, security associations, filters, and NAT and group matching

Administration The Administration menu manages the higher level functions that keep the 3000 unit operational and secure, such as who is allowed to configure the system, what software runs on it, and managing its configuration files and digital certificates. Only the administrator account can use the VPN Concentrator Manager. Figure 14-5 shows the Administration menu fully expanded out.

The Administration section provides access to control VPN 3000 Concentrator administrative functions:

- **Administer Sessions**—statistics and logout capability for all sessions
- **Software Update**—update concentrator and VPN client software
- **System Reboot**—system reboot options, including save and scheduling choices
- **Ping**—use ICMP ping to determine connectivity to an address

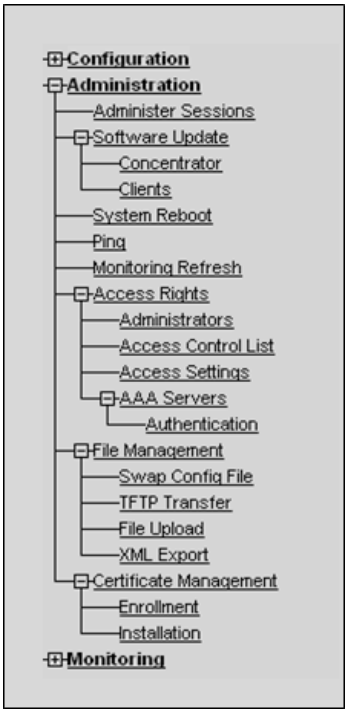

Figure 14-5 Administration menu fully expanded

- **Monitoring Refresh**—enable automatic refresh of Monitoring screens
- **Access Rights**—configure administrator profiles, access, sessions, and AAA
- **File Management**—view, save, delete, swap, and transfer files
- **Certificate Management**—install, enroll, and manage digital certificates

Monitoring The *Monitoring* menu is used to track many statistics and the status of many items essential to system administration and management. You can see the state of any LEDs that show the status of hardware subsystems in the device, as well as statistics stored and available in standard MIB-II data objects. Figure 14-6 shows the Monitoring menu fully expanded. The insert graphic shows the five second-tier options.

The Monitoring section of the Manager displays the VPN Concentrator status, sessions, statistics, and event logs. The Monitoring screens are read-only "snapshots" of data or status at the time the selection was made. These aren't real-time monitors. Most screens offer a Refresh button in the upper-right corner of the screen, which can be used to get a fresh image. The data on the screen can't be modified.

This section of the Manager lets you view VPN 3000 Concentrator status, sessions, statistics, and event logs.

- Routing Table—current valid routes and protocols
- Filterable Event Log—event logging with filtering capabilities

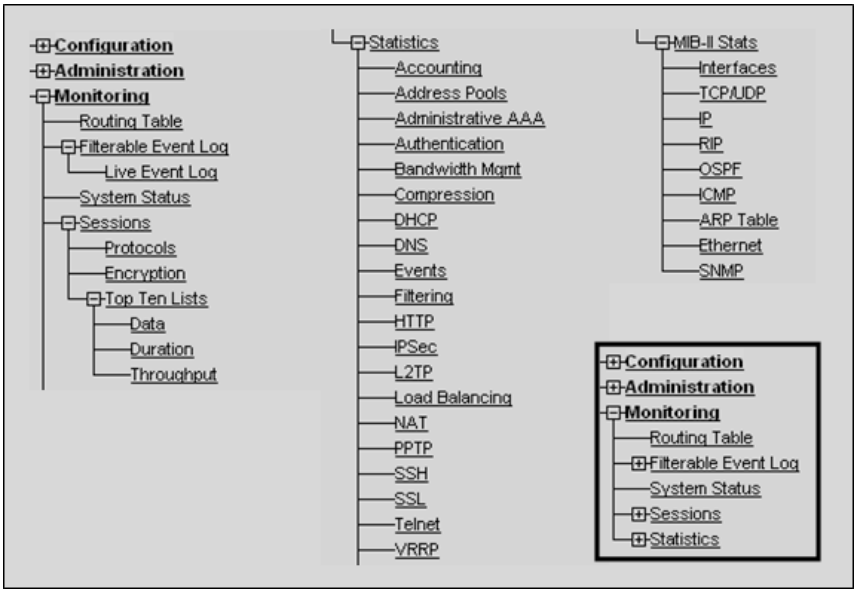

Figure 14-6 Monitoring menu fully expanded

- System Status—current software revisions, uptime, front-panel LEDs, network interfaces, SEP modules, and power supplies
- Sessions—all active sessions and "top ten" sessions, encryption and protocol data
- Statistics—accounting, address pools, administrative AAA, authentication, bandwidth management, compression, DHCP, DNS, events, filtering, HTTP, IPSec, L2TP, load balancing, NAT, PPTP, SSH, SSL, Telnet, and VRRP and MIB-II statistics

Figure 14-7 shows an example of a Monitoring screen using the Monitoring | System Status menu option. This screen shows the current version of the software, device serial number, the time the unit has been up, as well as CPU and Fan statistics. The device images show LED status and optional components, plus interface, modules, and power supplies with embedded links that display component statistics if clicked on.

Help

Figure 14-8 shows the result of selecting Configuration | Interfaces in the left panel, and then clicking the Help button in the upper-right corner. The context-sensitive Help window works like any Windows help document.

NOTE Help is a Java application, so Java must be enabled to see it and you might need to turn off any pop-up window protection software.

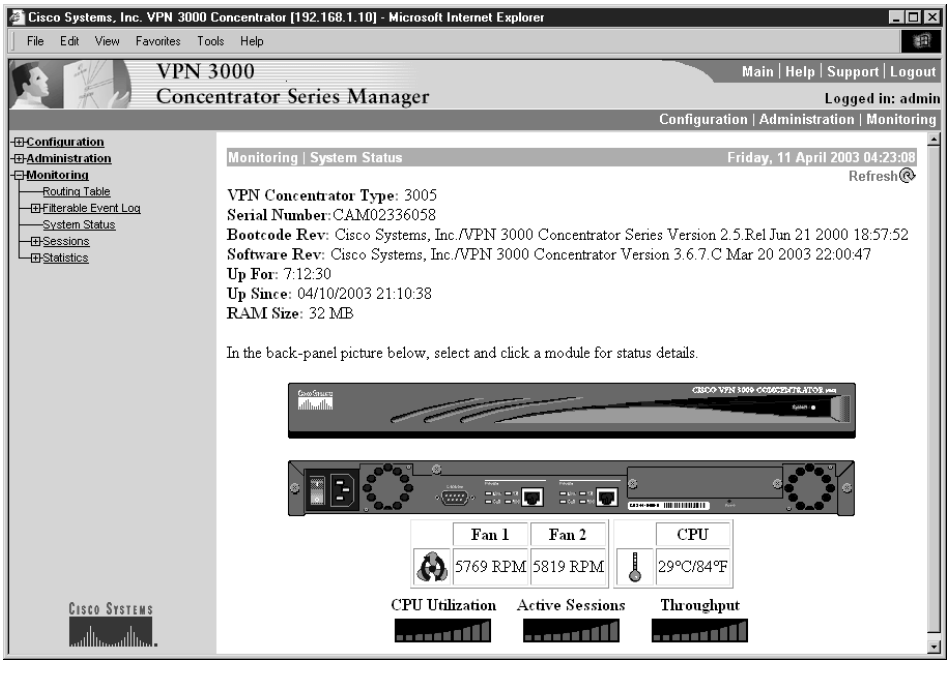

Figure 14-7 System monitoring screen

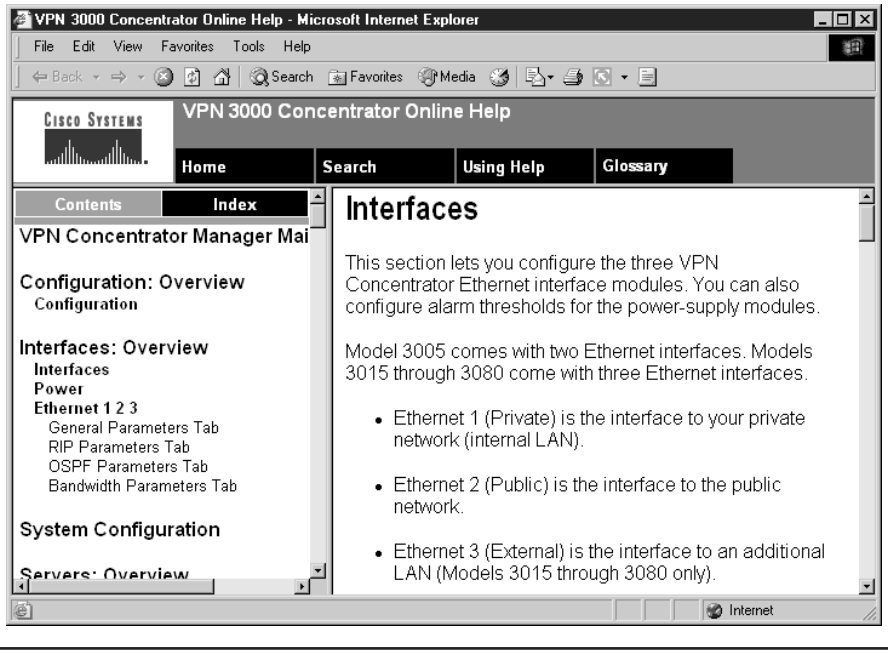

Figure 14-8 VPN 3000 Concentrator Manager Help system

PART III

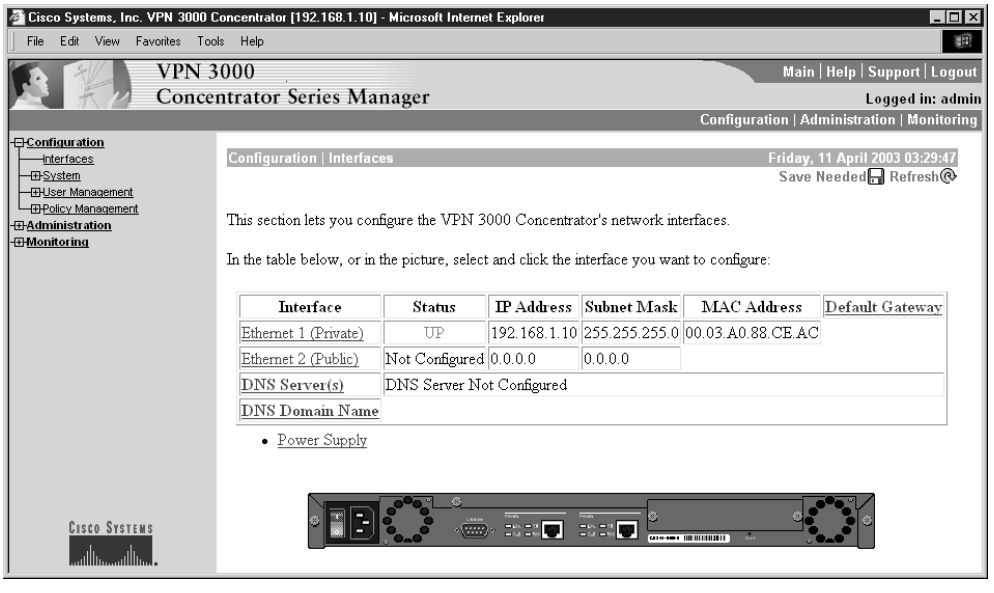

Figure 14-9 Typical Manager screen with reminder icons

Basic Operations

The Manager application takes advantage of the graphical interface to display as much information as feasible. Figure 14-9 shows the result of selecting Configuration | Interfaces from the menu. Notice the application map in the left-side panel shows where the screen is located in the structure. The figure might not show this, but the actual selection is highlighted.

In the upper-right corner, two icons remind you that changes must be saved and the screen needs to be refreshed—the data has timed out and is considered stale. This is a most useful reminder when counters and statistics are displayed. Clicking the icon activates the indicated service.

In the example figure, the back side of the device is displayed. In larger models, this would indicate any options that were added to the device.

VPN Concentrators in IPSec VPN Implementations

The Cisco VPN Concentrator series of devices has two basic VPN implementations:

- Remote access
- LAN-to-LAN

A large VPN implementation might frequently have variations of both types. It wouldn't be uncommon to have several LAN-to-LAN connections for branch office

links combined with many remote access connections for traveling executives, sales professionals, and telecommuters.

Remote Access Networks

Remote access involves connecting individual users to a LAN to provide secure, encrypted network access for telecommuters, traveling employees, and one-person offices of consultants, contractors, brokers, vendors, and so forth.

Today, many companies provide their own VPN connections through the Internet, allowing access to remote users running VPN client software over dial-up technologies through their Internet service providers (ISPs). The rapid expansion of cable and DSL markets makes it possible for telecommuters and other fixed location users to replace slower modem and ISDN services with fast connections at a fraction of the cost of dedicated lines.

Fast Internet connections offered in many hotels and the new wireless access facilities in many public places, such as airports, convention centers, and even fast-food restaurants, means traveling employees can also use fast, secure, remote VPN connections. Figure 14-10 shows two common types of remote access VPN examples.

Remote access VPN implementations with VPN Concentrators is covered in the upcoming section "Remote Access VPNs with Preshared Keys." Chapter 15 covers the VPN 3002 client device, which allows a single appliance to provide remote access VPN connectivity and protection to small groups at a single location.

LAN-to-LAN Networks

LAN-to-LAN (site-to-site) VPNs are an alternative WAN infrastructure used to create secure connections between two end devices, such as routers, firewalls, or VPN Concentrators. The hosts on each LAN connect to those end devices and can access the other LAN via the secure connection based on the organization security policy and the placement of shared resources.

Common examples of site-to-site VPN implementation could include connecting branch offices, vendor sites, dealer sites, or customer offices to the corporate network. Figure 14-11 shows the types of connections that might be VPN candidates.

LAN-to-LAN VPN implementations with VPN Concentrators are covered in Chapter 16.

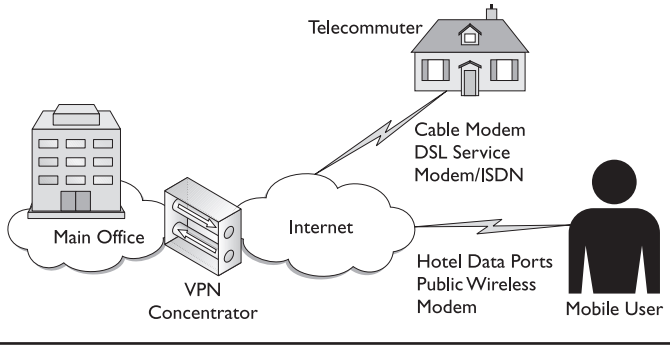

Figure 14-10 Remote access type VPN connections

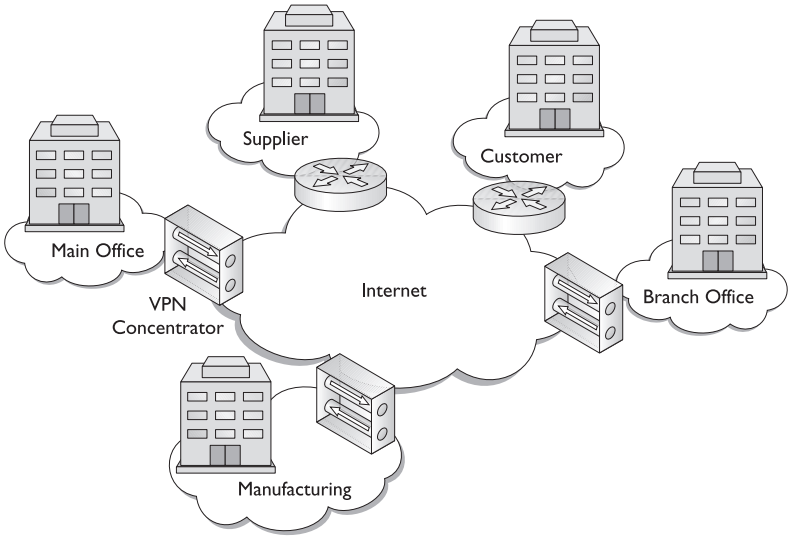

Figure 14-11 LAN-to-LAN VPN connection example

Remote Access VPNs with Preshared Keys

This section discusses configuring VPN 3000 Concentrators to support remote access implementations. Remote access VPN clients can use any of the four following common connectivity technologies:

- VPN client software installed on PCs or workstations
- VPN 3002 Hardware Client
- Router supporting VPN
- Firewall supporting VPN

The proliferation of low-cost small routers and firewall devices is making this an attractive and more secure option for many remote users. Cisco's small routers and firewall all support remote access options. For this example, we assume the remote users are using VPN client software. In any case, two levels of authentication need to occur. First, the device must authenticate itself to the concentrator or peer device, and then the user typically must authenticate on the network to gain access.

Cisco Easy VPN is a software enhancement that enables all three types of remote access clients to connect easily to the central site with minimal end user involvement. This is most important with large implementations and sites with limited local support. Chapter 12 covered the VPN client software, Chapter 15 covers the VPN 3002 Hardware Client, and Chapter 21 covers the PIX VPN connections, including Easy VPN. In this chapter, we focus on the VPN Concentrator that would connect to each of these technologies.

Figure 14-12 shows a scenario to use in the VPN remote-access configuration example. The Internet was simplified, in case someone wants to configure the scenario as a lab

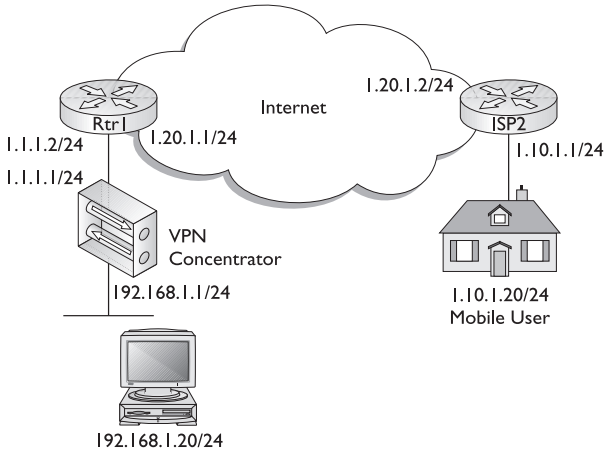

Figure 14-12 VPN Concentrator remote access scenario

exercise. The network behind Rtr1 could be the central site for a large number of remote clients and could contain multiple LANs.

NOTE The Concentrator Manager is used whenever possible in this example, but all steps could be accomplished using the CLI by following the same menu selections.

Preshared Keys

As Chapter 9 explained, VPN peer devices can authenticate each other using preshared keys or with digital certificates and certificate authorities. This section discusses using preshared keys and the section "Digital Certificates" looks at the changes required to use digital certificates. Three types of preshared keys exist.

- **Unique keys**—Each IP address has a unique key associated with it, providing a high level of security, but requiring much greater administrative overhead. While manageable for LAN-to-LAN implementations, it doesn't scale well for remote access networks with many users or a growing pool of users.
- **Group keys**—In this implementation, preshared keys are assigned to one or more groups. Managing new users becomes easier and more scalable by requiring only that the user be made a group member and given the appropriate passwords. This concept of groups has many other benefits, including the capability to separate different types of users based on access method (router vs. software client) for autoupdate features or even based on Security Policy access limitation differences. By default, the VPN Concentrator has one group: the Base Group. Additional groups can be added as needed, using the Base Group setting as defaults that can be modified as necessary.

- **Wildcard keys**—With a wildcard preshared key, it's possible to allow one or more clients to use a shared secret key to authenticate encrypted tunnels to the gateway. To be configured on the Concentrator, all devices must be using preshared keys. This option can be more prone to security problems, such as man-in-the-middle attacks, if a key is compromised. In that case, all keys must be changed.

Initial Configuration

The following example assumes either the Quick Configuration wasn't used or the related features should be changed. At a minimum, the private interface must be configured for the VPN 3000 Concentrator using the CLI. While the CLI can also be used to configure the Public Interface and any default routes, this exercise performs those tasks using the Manager interface. The following lists the correct addresses for the scenario:

VPN Concentrator: 192.168.1.1/24.

```
               Welcome to
              Cisco Systems
       VPN 3000 Concentrator Series
          Command Line Interface
Copyright (C) 1998-2003 Cisco Systems, Inc.

1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main ->1

1) Interface Configuration
2) System Management
3) User Management
4) Policy Management
5) Back

Config -> 1                              (shows the current settings)

This table shows current IP addresses.
  Intf         Status       IP Address/Subnet Mask          MAC Address
-------------------------------------------------------------------------
Ether1-Pri|      UP      | 192.168.10.1/255.255.255.0 | 00.03.A0.88.CE.AC
Ether2-Pub|Not Configured|      0.0.0.0/0.0.0.0       |
-------------------------------------------------------------------------
DNS Server(s): DNS Server Not Configured
DNS Domain Name:
Default Gateway: Default Gateway Not Configured

1) Configure Ethernet #1 (Private)
2) Configure Ethernet #2 (Public)
3) Configure Power Supplies
4) Back

Interfaces -> 1                         (to set the Private interface)
```

```
1) Interface Setting (Disable, DHCP or Static IP)
2) Set Public Interface
3) Select IP Filter
4) Select Ethernet Speed
5) Select Duplex
6) Set MTU
7) Set Port Routing Config
8) Set Bandwidth Management
9) Set Public Interface IPSec Fragmentation Policy
10) Back

Ethernet Interface 1 -> 1                  (to set a Static IP address)

1) Disable
2) Enable using DHCP Client
3) Enable using Static IP Addressing

Ethernet Interface 1 -> [ 3 ] 3

> Enter IP Address                  (current value appears in brackets)
Ethernet Interface 1 -> [ 192.168.10.1 ] 192.168.1.1

> Enter Subnet Mask
Ethernet Interface 1 -> [ 255.255.255.0 ]
Ethernet Interface 1 -> h                          (return to Main menu)
```


Basic Browser Configuration

Once the CLI is used to configure the private interface, you can use a browser on a PC on the same network to complete the configuration. The following steps assume the private interface was configured and a successful web session was established using a client on the LAN attached to that interface.

Setting the Public Interface

Use the Configuration | Interfaces | Ethernet 2 menu options to get to the screen, as shown in Figure 14-13. Select Static IP Addressing, and then enter the IP address and the appropriate mask. This same screen allows making choices on interface speed, full- or half-duplex, MTU size, fragmentation preferences, and whether the interface can be a DHCP client.

Check the DHCP Client check box if it's necessary to obtain the IP address, the subnet mask, and the default gateway for this interface via DHCP from an ISP. This would be more common at a branch location or any small client that doesn't require permanent static IP addresses for shared servers. If this check box is selected, don't make entries in the IP address and subnet mask fields that follow.

To make this interface a public interface, check the Public Interface check box designating the interface is part of a public network, such as the Internet. A public interface must be configured before you can configure NAT and IPSec LAN-to-LAN. Designate only one VPN Concentrator interface as a public interface.

Notice the tabs on this screen would allow configuring RIP, OSPF, and Bandwidth Management features. While static routes are common with Concentrators, some LAN-to-LAN features, such as Network Autodiscovery, require or can take advantage of this support.

Figure 14-13 Configuring the public interface

Clicking the Apply button applies the choices and returns to the Configuration | Interfaces screen, which confirms the new IP address and mask.

Defining the Default Gateway (Optional)

Default routes or *default gateways* are used to forward packets addressed to unknown networks to avoid having to discard the packet. The not-too-practical alternative is that all routers, firewalls, and VPN devices would have to know all the routes in the world.

The default route for the Concentrator to which you forward all outside bound packets can be defined either by choosing the Default Gateway link on the Configuration | Interfaces screen or by using the Configuration | System | IP Routing | Default Gateways menu to add these special forms of static routes. Figure 14-14 shows the resulting screen with the appropriate entries to use Rtr1 as the default gateway for the network.

In the example scenario, if no additional routers are in the LAN, the PCs in the LAN would use the Concentrator private address for their default gateways.

Verifying Configuration

Returning to the Configuration | Interfaces screen, as shown in Figure 14-15, confirms the default gateway was defined and provides a reminder icon to save the configuration.

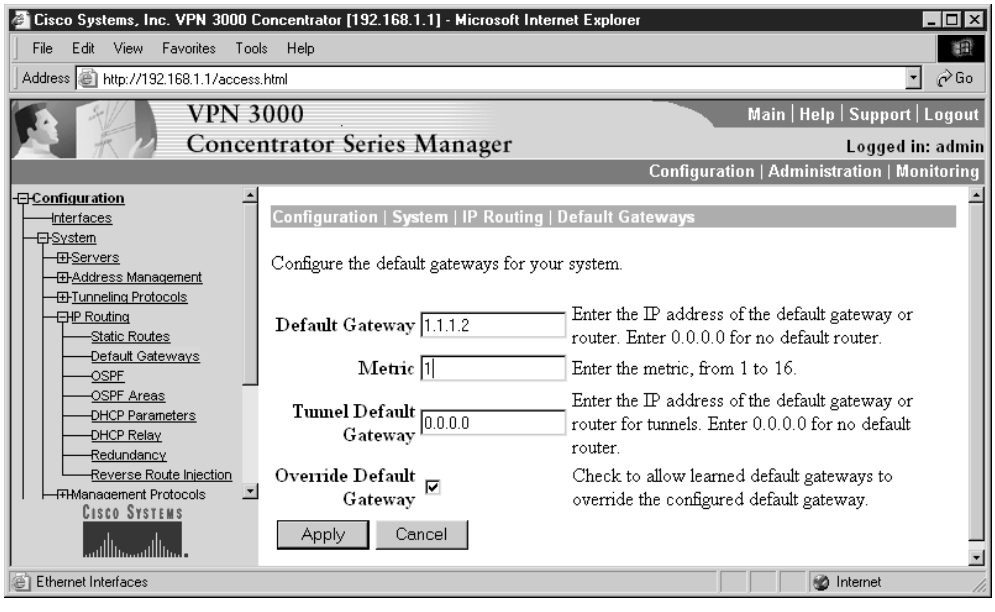

Figure 14-14 Setting the default gateway for the Concentrator

This same screen confirms the interface configurations and whether the two interfaces are up or down. If either interface is down, this could be caused by any of the following:

- You might need to click the Refresh icon.

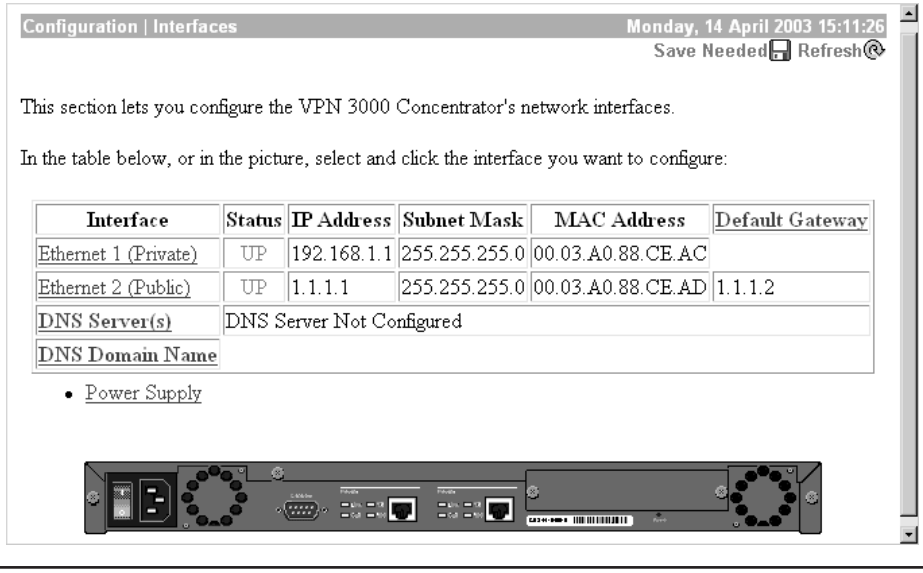

Figure 14-15 Using the Configuration | Interfaces screen to confirm configuration

PART III

- You might have used the wrong cable. Use straight through to the hub or the switch; otherwise, use a crossover.
- The device on the other end might be misconfigured or unconfigured.

At this point, you should be able to use the Administration | Ping screen to ping other devices on the network to verify connectivity. As with the IOS Routers and PIX Firewalls, a good idea is always to verify connectivity before configuring encryption and access control measures. While pinging the Concentrator public interface is possible, it won't be possible to ping through to the LAN or the private interface. Also, opening Telnet or browser sessions with the Concentrator from the outside network won't be possible.

Adding the Static Routes

Use the Configuration | System | IP Routing | Static Routes menu to add any necessary static routes. Unless configured earlier using the CLI, the Static Routes box should be empty or show only any configure default gateway. Use the Add button to bring up the screen, shown in Figure 14-16, to add any static routes required. The figure shows configuring a route specifically to the network the remote host is on. Because the default gateway would accomplish the same result, this static rout is included only as an example.

On the static route screen, you can specify the Destination Router Address (next-hop) or the Interface (Concentrator Public). Clicking the Add button returns to the Configuration | System | IP Routing | Static Routes screen. Figure 14-17 shows the Static Routes screen with a default gateway, a static route pointing to a next-hop router and another one pointing at the Concentrator public interface. The last entry shows a possible

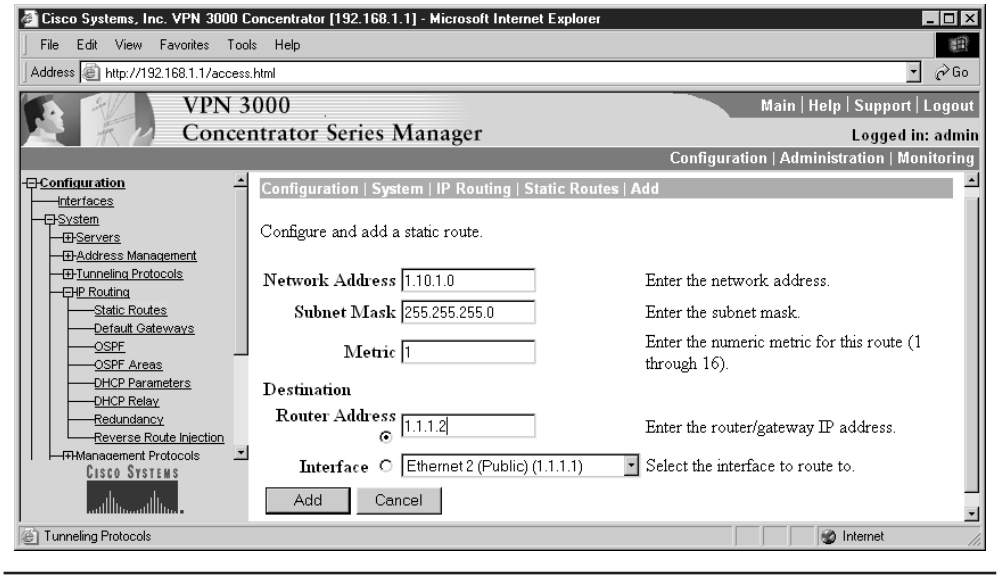

Figure 14-16 Creating a new static route

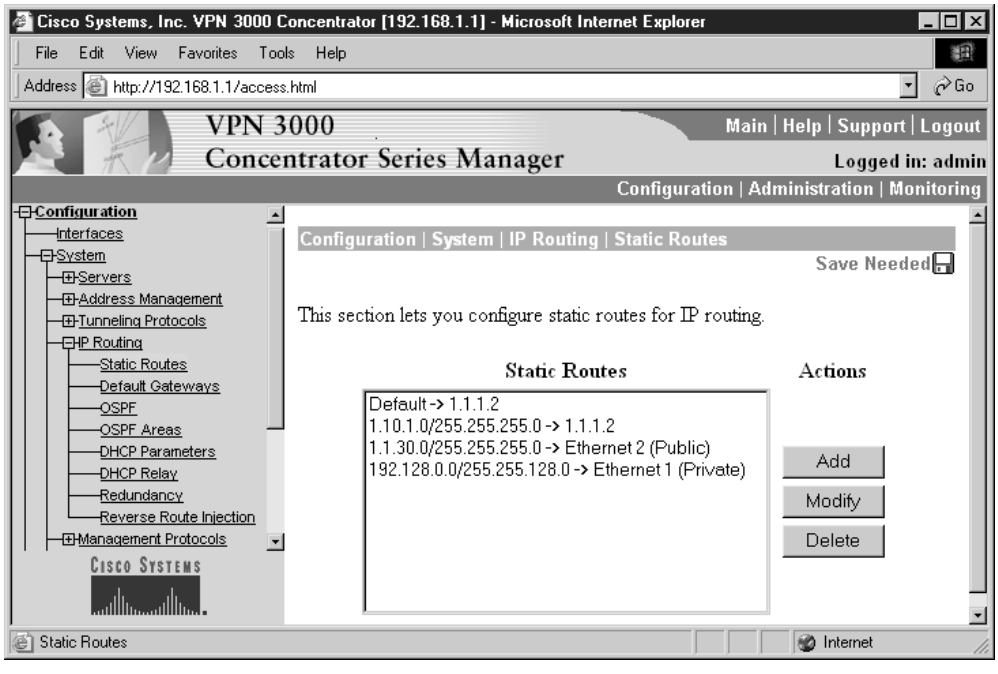

Figure 14-17 Defined default and static routes

configuration if the LAN protected by the Concentrator included the 192.168.0.0 to 192.168.127.0 networks built around additional routers and/or Layer 3 switches.

General System Information

Use the Configuration | System | General menu to add or modify any of the following. The time and date settings are critical to the proper operation of IPSec, especially when working with certificate authorities. Each related screen is self-explanatory.

- **Identification**—system name, contact person, device location
- **Time and Date**—device time and date
- **Sessions**—maximum simultaneous sessions
- **Authentication**—global authentication parameters

Define Inside Address Assignment Method

Remote users have network addresses associated with their local network or, more likely, their ISP network. To function within the private network (LAN), it's necessary to assign suitable "inside" addresses to these users. Conceptually, this is similar to a reverse NAT.

Use the Configuration | System | Address Management | Assignment menu to define how the remote users are assigned addresses within the private network. The screen shown in Figure 14-18 is used to select prioritized methods for assigning IP addresses to clients as a tunnel is established. The Concentrator tries the selected methods in the order listed, until it finds a valid IP address to assign. A minimum of one method must be selected or any combination of choices will be evaluated. There is no default method. The figure shows using an address pool defined on the VPN Concentrator.

Four possible methods exist for assigning addresses to the remote users:

- **Use Client Address**—enables the client to specify its own IP address. This isn't a good security strategy. Don't use this option for IPSec because IPSec doesn't allow client-specified IP addresses.
- **Use Address from Authentication Server**—used to assign IP addresses retrieved from an authentication server (AAA) on a per-user basis. This is the preferred and most secure method, if an authentication server (external or internal) is being used.
- **Use DHCP**—used to obtain IP addresses from a DHCP server. This is the most manageable and the most scalable of the two remaining options available to IPSec networks.
- **Use Address Pools**—used to have the VPN Concentrator assign IP addresses from an internally configured pool. Internally configured address pools are the easiest method of address pool assignment to configure. Use the Configuration | System | Address Management | Pools screens covered in the next section to define and prioritize the address pools.

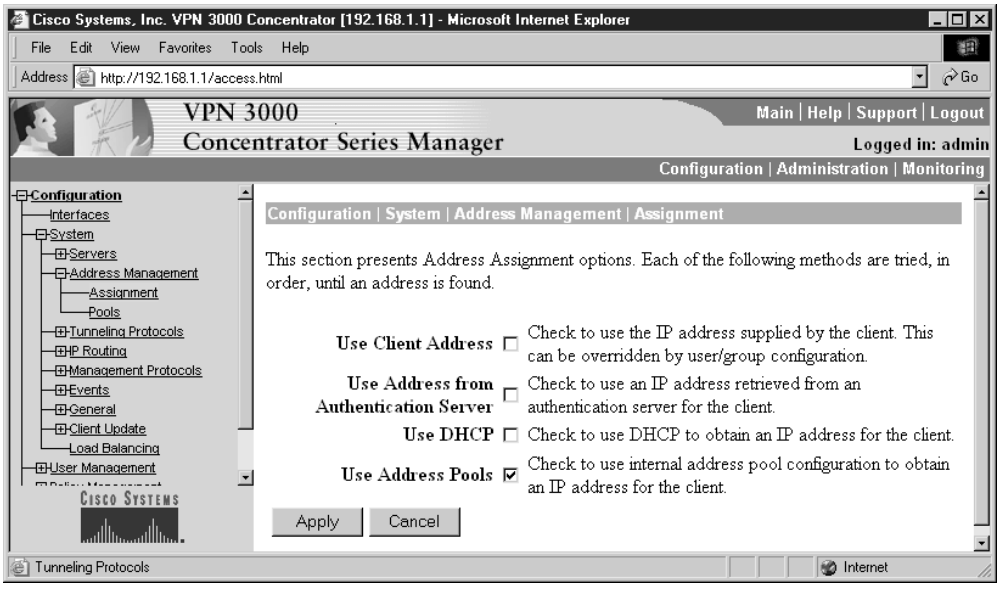

Figure 14-18 Selecting an inside address assignment method for remote users

If the assignment method uses addresses that aren't from the local subnet on the private interface, then it's necessary to add routing for those subnets.

Define Inside Address Pool for Remote Users

Use the Configuration | System | Address Management | Pools menu to add or modify the inside address pool to be used by remote users, making them part of the inside network. Initially, the IP Pool Entry box is empty. Click the Add button to bring up the screen shown in Figure 14-19, which shows it is using the last 54 addresses in the inside network. Note, these addresses shouldn't be assigned to inside hosts or included in any DHCP pools used in the network.

Clicking the Add button implements the pool and returns to the Configuration | System | Address Management | Pools screen, as shown in Figure 14-20, to confirm the pool was added. The Add, Modify, and Delete buttons can be used to create new, to edit existing, or to remove existing pools. The Move Up/Move Down buttons are used to change the priority of the address pools: the higher a pool is in the list, the more likely the addresses are to be used. No Undo option exists yet on any of the configuration options.

Configuring Groups and Users

VPN features and configurations are assigned and "pushed" down to clients on the basis of group and user assignment and configuration. As with basic network-resource sharing, this is always most scalable and easier to administer if features are assigned to defined groups. Individual users are then put in groups to assign features to them. Then, if a person leaves or changes his or her job role, that person only needs to be removed from the group.

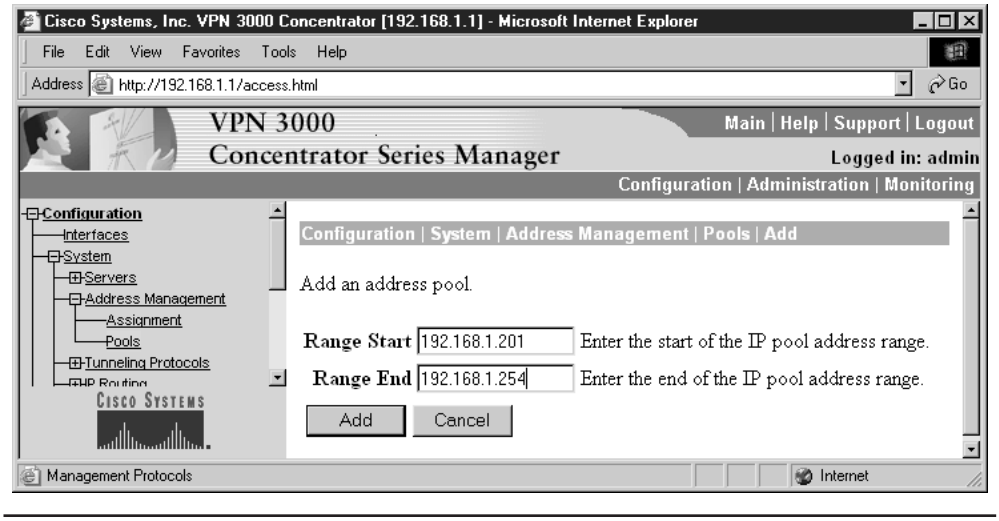

Figure 14-19 Defining a pool of addresses for remote users to use

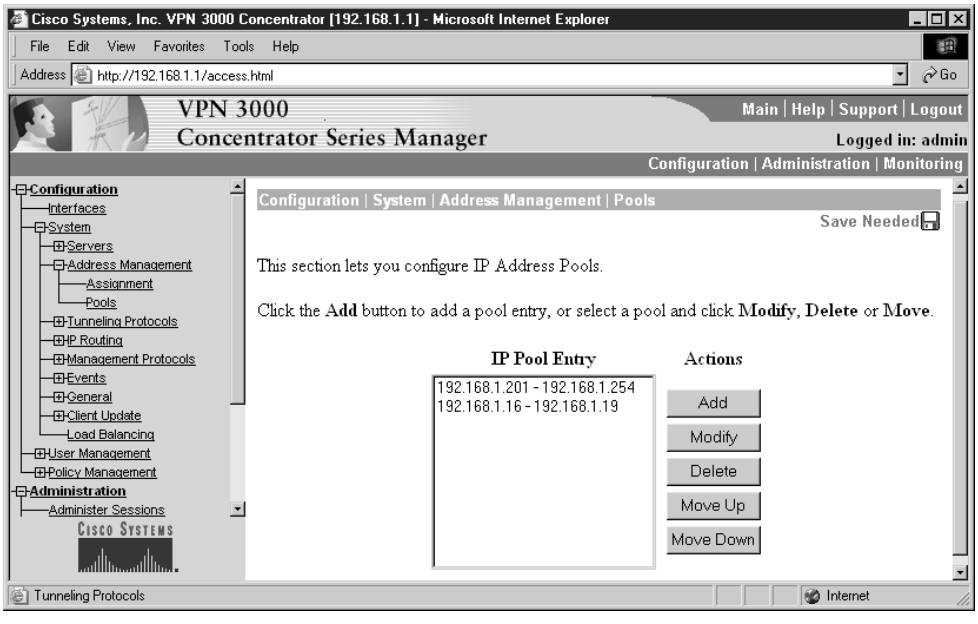

Figure 14-20 Managing IP address pools for remote users

The Configuration | User Management screen is used to configure VPN 3000 Concentrator group and user parameters for IPSec, PPTP, and L2TP. The following are the three configuration options:

- **Base Group**—default group and user parameters
- **Groups**—add and modify groups and group parameters
- **Users**—add and modify users and user parameters

Setting Group and User Defaults

While modifying the Base Group default settings isn't absolutely necessary, Cisco's default setting might not match the network or company security policy. The Configuration | User Management | Base Group screen shows a series of tabs—currently six—that can be used to define default settings for the various users (clients). Figure 14-21 shows the default settings for the General tab, which applies to all groups and users.

Most of the options are self-explanatory and the Help system provides enough information, including default values, in most cases to make intelligent choices. Remember, most of these options might be moot because of specific group and user settings, as well

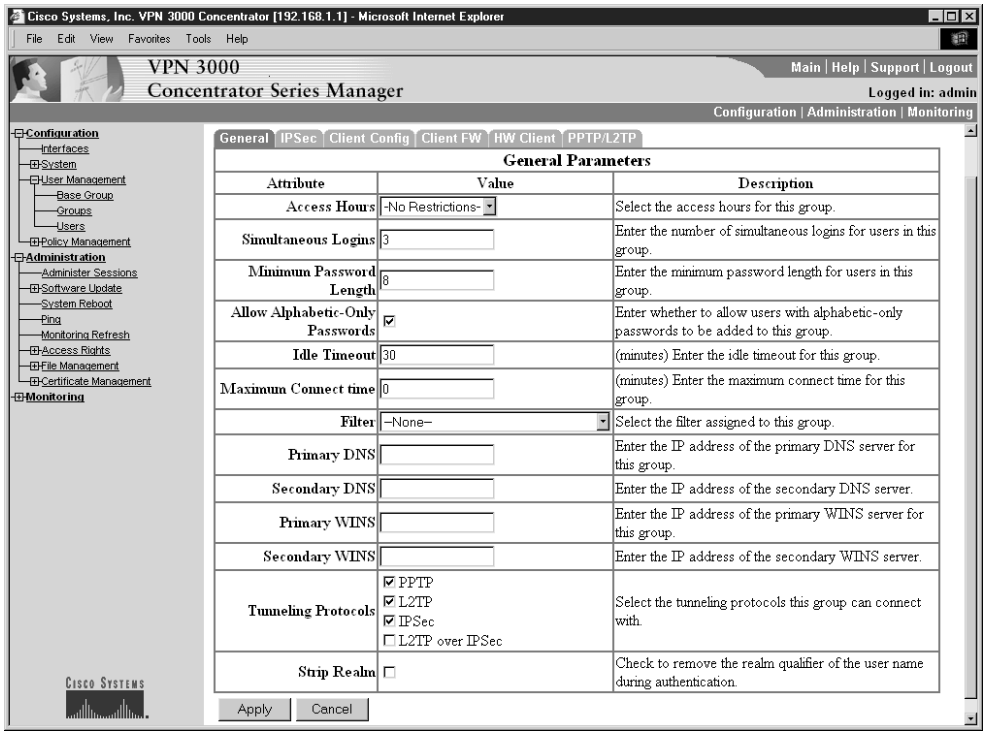

Figure 14-21 General Default Settings tab

as limitations imposed by the local client device. The following options are included as examples or to show interesting features:

- **Access Hours**—Drop-down menu button to select the hours remote access users can access the VPN Concentrator. The default entries are as follows:
  - **No Restrictions**—No restrictions on access hours.
  - **Never**—No access at any time.
  - **Business Hours—**Access 9 A.M. to 5 P.M., Monday through Friday. To configure access hours, use the Configuration | Policy Management | Access Hours screen. Any additional named access hours you create will appear on the list.
- **Simultaneous Logins**—For a single internal user. Default is 3. Minimum is 0, which disables login and prevents user access. No maximum limit.

- **Minimum Password Length**—Default is 8. Minimum is 1. Maximum is 32.
- **Idle Timeout**—Default is 30 minutes. Minimum is one minute. Maximum is 2,147,483,647 minutes (4,000+ years). Use 0 to disable timeout and allow an unlimited idle period.
- **Maximum Connect Time**—Default is 0 minutes, which allows unlimited connection time. Minimum is one minute. Maximum is 2,147,483,647 minutes (4,000+ years).
- **SEP Card Assignment**—The VPN Concentrator models 3015 and above can contain up to four SEP modules to handle encryption. This parameter allows configuring the load on each SEP module.
- **Tunneling Protocols**—Protocols to be supported. Clients can use only the selected protocols. The choices include the following:
  - **PPTP**—Point-to-Point Tunneling Protocol. Checked by default. A client-server protocol popular with Microsoft clients, especially older OS versions.
  - **L2TP**—Layer 2 Tunneling Protocol. Checked by default. A client-server protocol that combines many features from PPTP and L2F (Layer 2 Forwarding).
  - **IPSec**—IP Security Protocol. Checked by default. Used by both LAN-to-LAN (peer-to-peer) connections and client-to-LAN (remote access) connections. Cisco VPN Client is an IPSec client specifically designed to work with the VPN Concentrator.
  - **L2TP over IPSec**—L2TP using IPSec for security. Unchecked by default. A client-server protocol providing interoperability with Windows 2000 VPN clients. L2TP packets are encapsulated within IPSec, thus providing an additional authentication and encryption layer.

Setting IPSec Defaults

The Configuration | User Management | Base Group screen with the IPSec tab selected is used to configure IP Security Protocol parameters that apply to the base group. This section would apply if IPSec or L2TP over IPSec were selected on the General Parameters tab. Figure 14-22 shows the default settings. The options are self-explanatory and the Help system provides enough information, including default values, to make appropriate choices.

The IPSec SA option contains a drop-down menu button to select the default IPSec security association (SA) assigned to IPSec clients during tunnel negotiation. Remote-access clients must have a default SA defined, while LAN-to-LAN connections ignore this selection and use parameters from the Configuration | System | Tunneling Protocols | IPSec LAN-to-LAN screens.

The Default Preshared Key option (the third option from the bottom) is used to define the preshared secret key. Use a minimum of four options and a maximum of

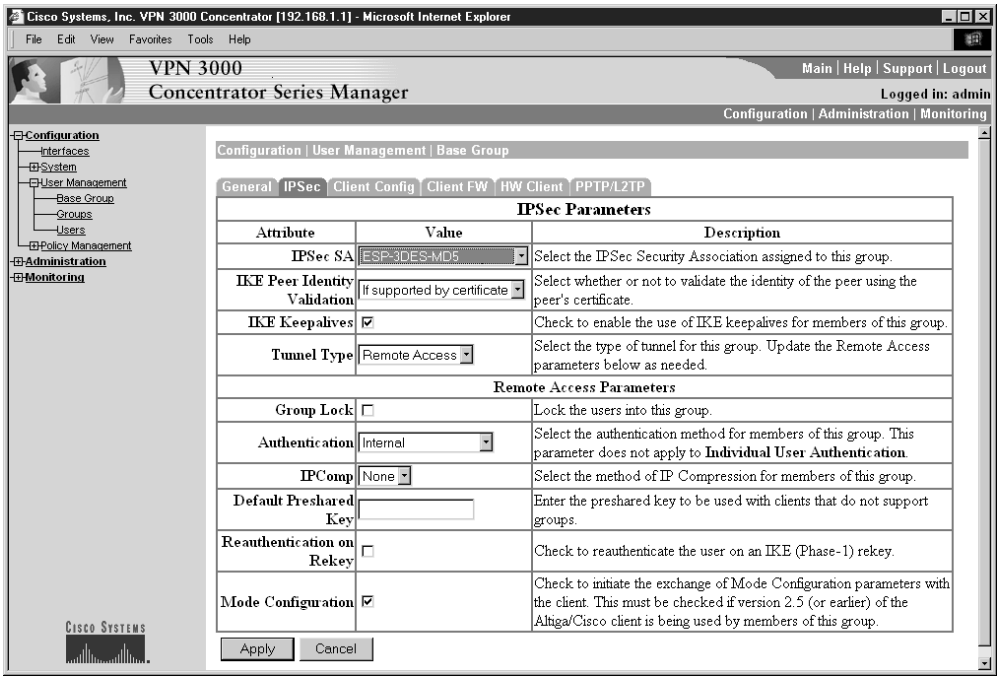

Figure 14-22 Setting IPSec defaults

32 alphanumeric characters. This option allows the following VPN clients to connect to the VPN Concentrator:

- VPN clients that use preshared secrets, but don't support "groups," such as the Microsoft Windows XP L2TP/IPSec client
- VPN router devices that are creating inbound connections from nonfixed IP addresses using preshared secrets

Setting Client Defaults

The remaining four tabs are used to configure client-specific default settings. These are the configuration features that are pushed down to the appropriate client the next time the device connects to the Concentrator. The *Client Config* tab is used to set general client features that might apply to all Cisco VPN Clients. The remaining three tabs are for setting those defaults features unique to that type of client.

Figure 14-23 shows the Cisco Client parameters such as login banners, choosing IPSec over UDP, and defining up to ten IPSec Backup Servers.

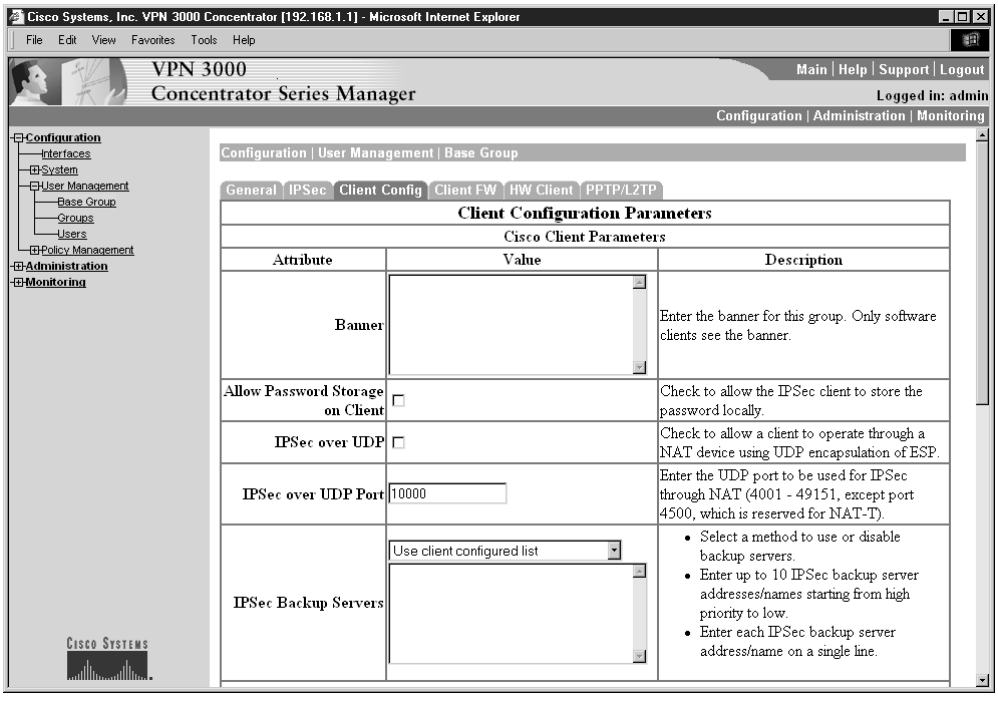

Figure 14-23 Setting Cisco Client default parameters

Figure 14-24 shows the rest of the previous screen, which includes setting two Microsoft Client parameters and three Common Client parameters. This is the screen where split tunneling is allowed by default for VPN Clients.

Split Tunneling Policy

Split tunneling allows the IPSec client to go directly to the Internet in Cleartext form for those destinations that don't require encryption. Split tunneling applies only to remote-access IPSec tunnels, not to LAN-to-LAN connections.

Split tunneling eases the device-processing load, simplifies traffic management, and speeds untunneled traffic. Split tunneling is a traffic management feature, not a security feature. In fact, for optimum security, split tunneling isn't recommended. Because only the VPN Concentrator, not the IPSec client, can enable split tunneling, it's possible to control implementation and minimize security vulnerabilities.

Split tunneling is disabled by default on both the VPN Concentrator and the client. To enable and configure the feature, all entries are made on the VPN Concentrator, and then pushed down to the IPSec client. The default split-tunneling policy is *Tunnel Everything*, which disables split tunneling. No traffic goes in clear text or to any destination

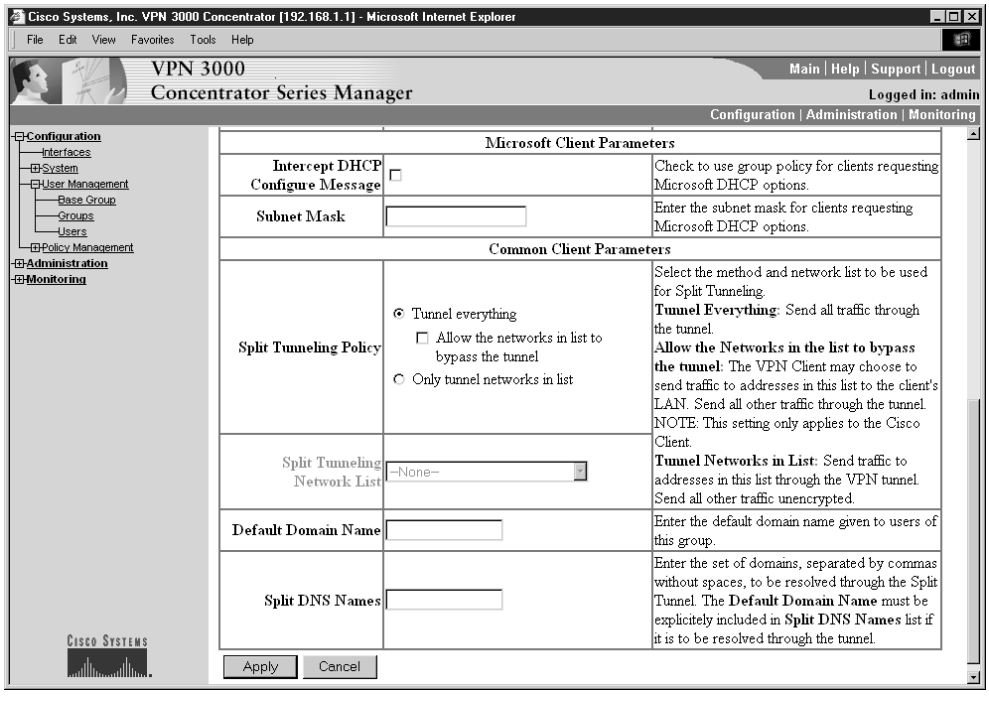

Figure 14-24 Setting Microsoft Client parameters and Common Client parameters

other than the VPN Concentrator. Remote users in this group reach Internet networks through the corporate network and don't have access to local networks.

The *Allow Networks in List to Bypass Tunnel* allows the administrator to define a list of networks to which traffic can go without passing through the tunnel. This allows remote users to access devices on their local networks, such as printers, while still connected to the corporate network through a tunnel.

Only Tunnel Networks in List allows remote users to access Internet networks without tunneling through the corporate network.

Client Firewall Requirements

The *Client FW Parameters* tab, as shown in Figure 14-25, is used to configure firewall parameters for VPN Clients running Microsoft Windows. The features are currently unavailable for hardware clients or other non-Windows software clients. If the Firewall Required option is selected, the client won't be allowed to connect to the protected network without meeting the defined requirements. The firewall feature protects the user PC and, thereby, the corporate network, from intrusions originating from the Internet or the user LAN.

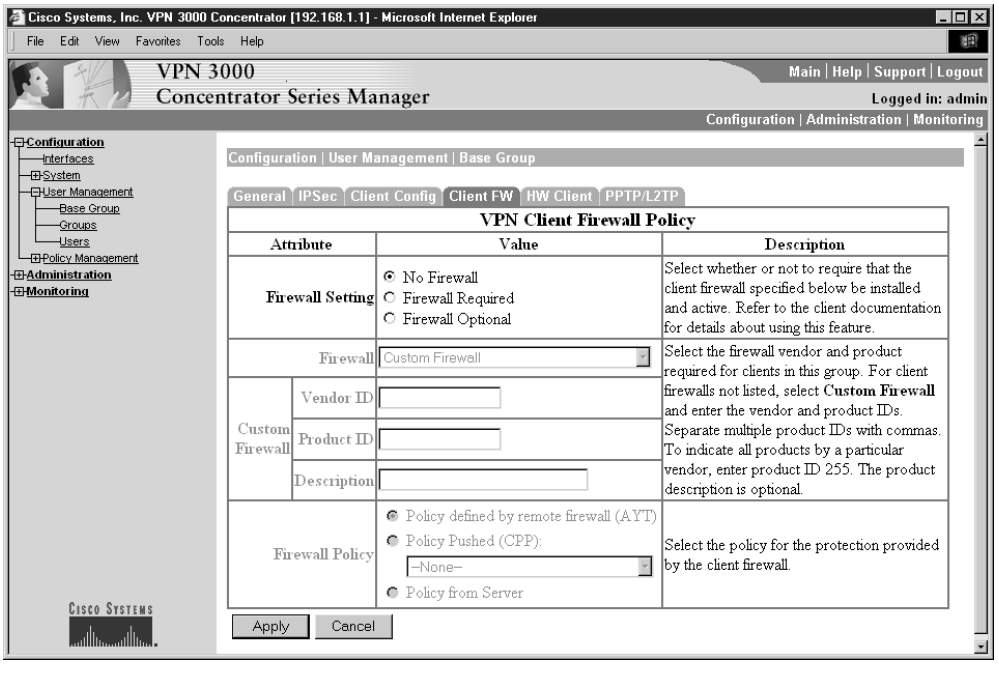

Figure 14-25 Defining client firewall default requirements

VPN 3002 Hardware Client Defaults

The *HW Client Parameters* tab, shown in Figure 14-26, is used to enable and configure interactive hardware client authentication and individual user authentication. These features are covered in the next chapter when the VPN 3002 Hardware Client is covered in greater detail.

Creating or Modifying Other Groups

The Base Group features set in the last sections automatically apply to all users accessing the VPN Concentrator. If the security policy requires defining multiple groups with unique features and/or requirements, the Configuration | User Management | Groups screen, shown in Figure 14-27, can be used to configure access and usage parameters. A *group* is a collection of users treated as a single entity. Groups inherit parameters from the Base Group. The figure includes a group created in the next section.

To create and use groups beyond the Base Group requires using an internal authentication server. The authentication server must be one of the following:

- RADIUS—An external RADIUS server is the default.
- NT Domain—An external Windows NT Domain server.
- SDI—An external RSA Security Inc. SecurID server.

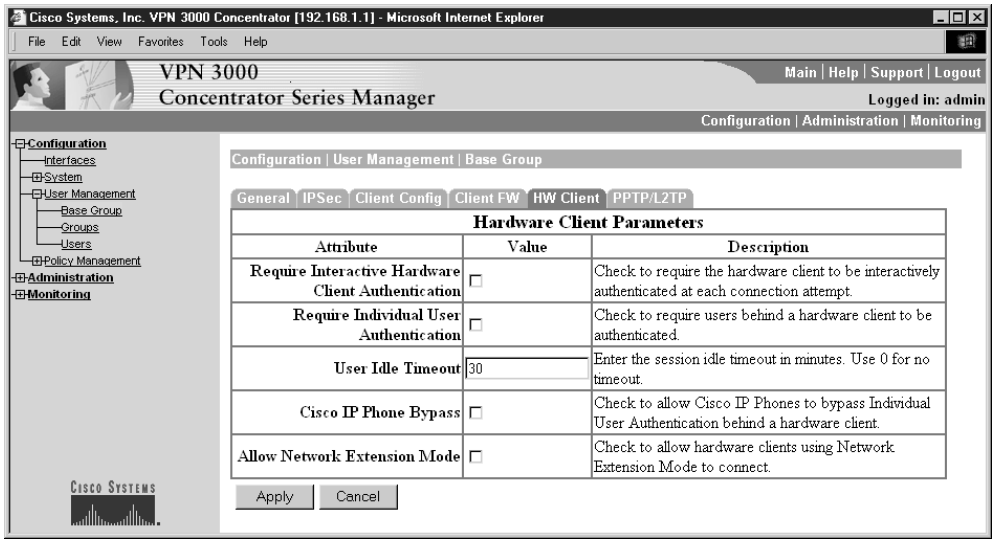

Figure 14-26 Defining VPN 3002 Hardware Client defaults

- Internal Server—The internal VPN Concentrator authentication server. With this server, you can configure a maximum of 100, 500, or 1000 groups and users (combined) in the internal database depending on the model number.

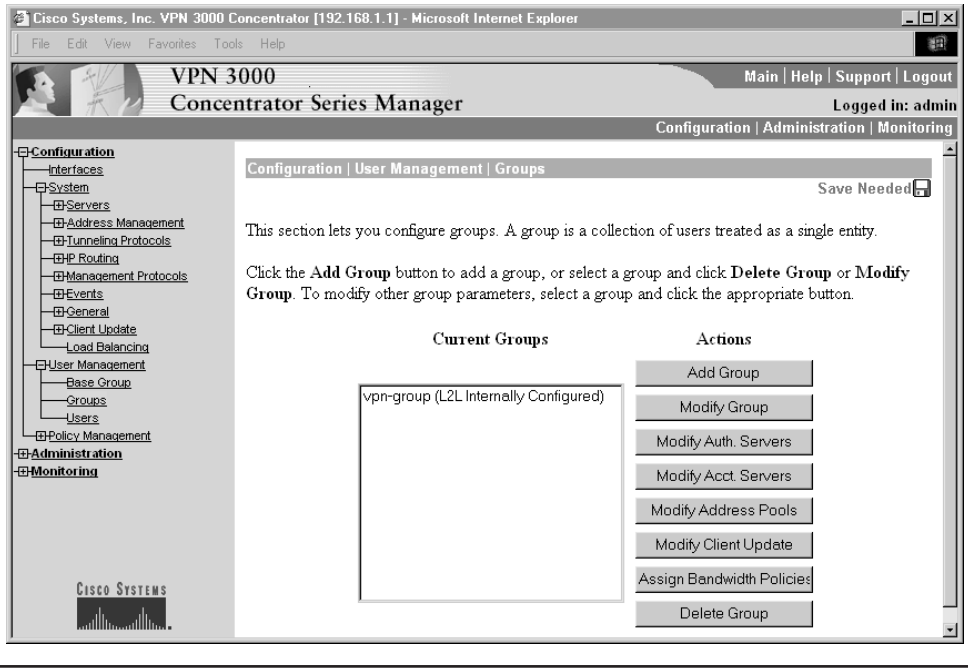

Figure 14-27 Screen to create and manage user groups

PART III

- Kerberos/Active Directory – Windows 2000/XP and Linux/Unix authentication server.

If no external server was defined, a link appears on the Groups or Users screen to create an internal server. Creating the internal server is simply a matter of clicking the link and choosing the Internal Server option.

Adding a Group The Configuration | User Management | Groups | Add screen, shown in Figure 14-28, shows the seven tabs used to define the configuration parameters for a new group. The *Identity Parameters* tab is used to configure the name, password, and authentication server type for this group.

Group Name	Unique case-sensitive name for this specific group. Maximum length is 64 characters. Changing a group name automatically updates the group name for all users in the group. For remote access users connecting with digital certificates, this name must match exactly the Organizational Unit (OU) field of the user's identity certificate.
Password	Unique case-sensitive password for this group. Minimum length is four characters. Maximum is 32 characters. The field displays only asterisks.
Verify	Reenter the group password to verify it. The field displays only asterisks.
Type	Use the Type drop-down menu button to select the authentication server type.

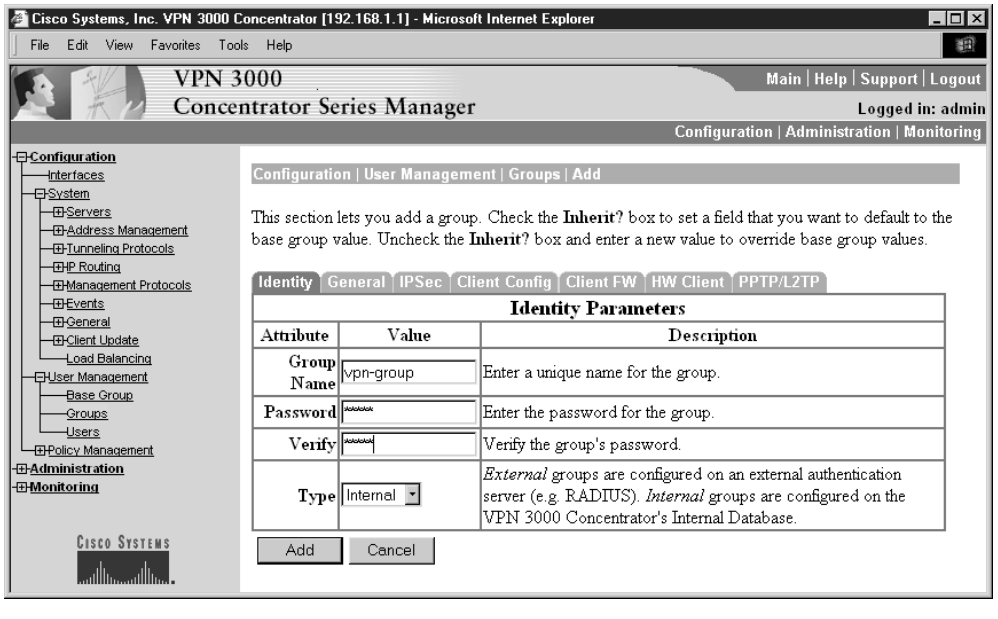

Figure 14-28 Creating a group and defining a password

The six remaining tabs are used to change any options that either weren't defined in the Base Group tabs with the same names or when the Base Group options need to be overridden. Figure 14-29 shows the parameter screens have a column to inherit the feature settings from the Base Group. A check mark in the Inherit? column indicates to inherit the default setting.

Creating or Modifying Users

The Configuration | User Management | Users screen, shown in Figure 14-30, allows configuring access, usage, and authentication parameters for users. Users inherit the feature parameters from the specific group to which they belong. Configuring users in this section requires configuring them in the VPN Concentrator internal authentication server. The authentication server is a properly configured RADIUS, NT Domain, SDI server, Internal Server, or Keberos/Active Directory server.

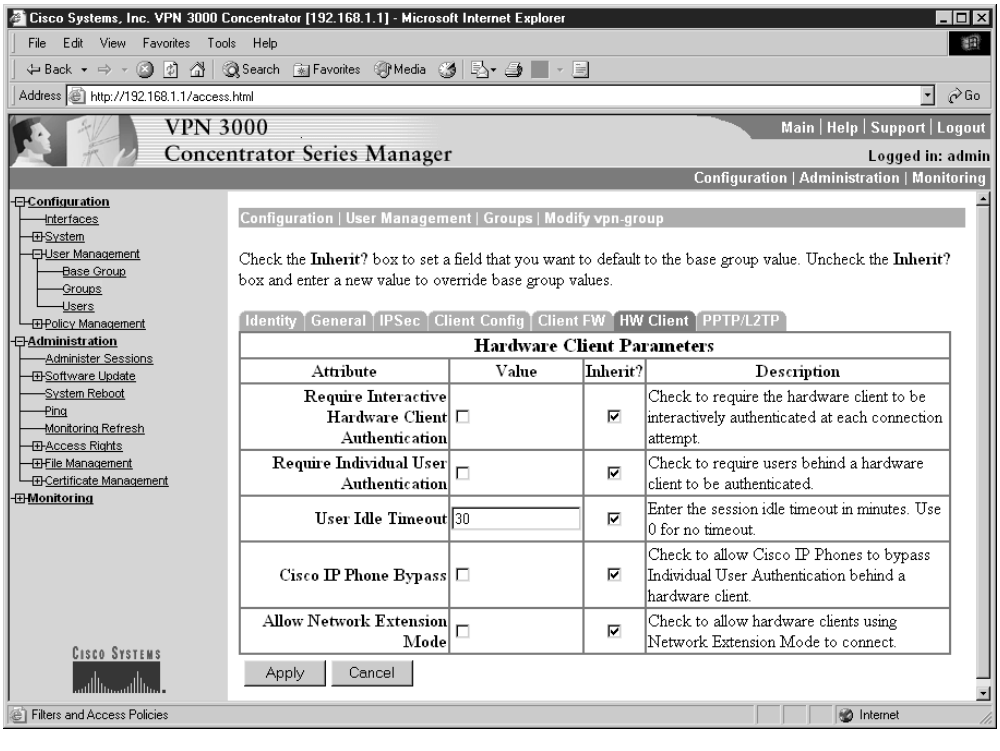

Figure 14-29 Group parameters screen showing the Inherit option

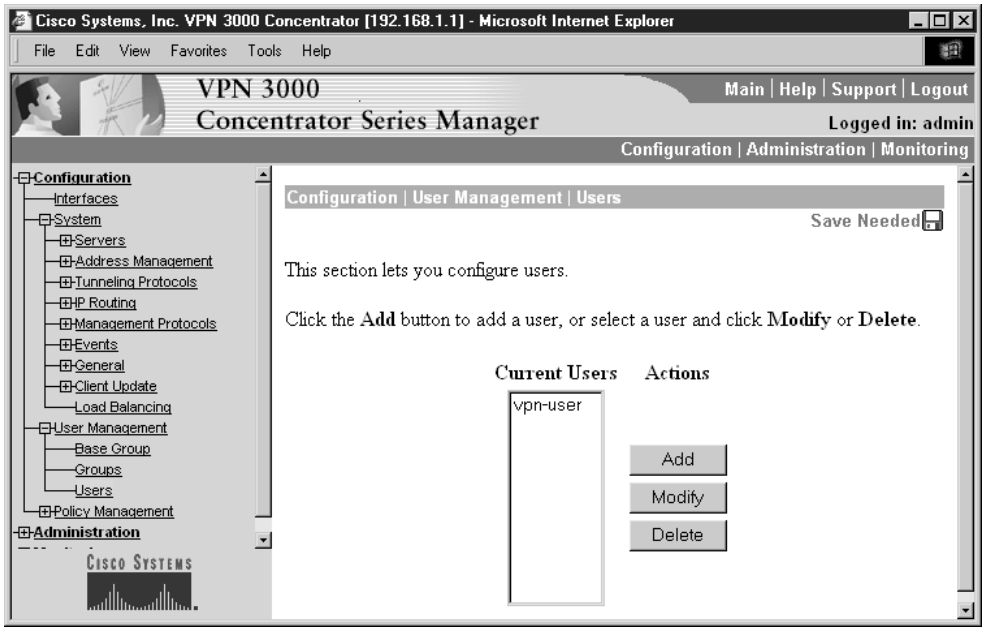

Figure 14-30 User management screen for adding users

Using the Add button brings up the Identity Parameters tab, as shown in Figure 14-31, used to configure the name, password, group, and IP address for this user.

User Name	Unique, case-sensitive name for this user. Maximum length is 64 characters. If the name is changed, this user profile replaces the existing profile.
Password	Unique, case-sensitive password for this user. Minimum length must satisfy the minimum for the group to which the user is assigned. Maximum length is 32 characters. The field displays only asterisks.
Verify	Reenter the user password to verify it. The field displays only asterisks.
Group	Drop-down menu used to select the group to assign this user. The list shows all defined groups plus the Base Group.
IP Address	Only if Use Address from Authentication Server on the Configuration | System | Address Management | Assignment screen. Otherwise, leave this field blank.
Subnet Mask	Same as IP address field.

The other three tabs are used to specify parameters for this user that vary from the group. Each feature has an Inherit check box, selected by default, which is used to override the group feature settings.

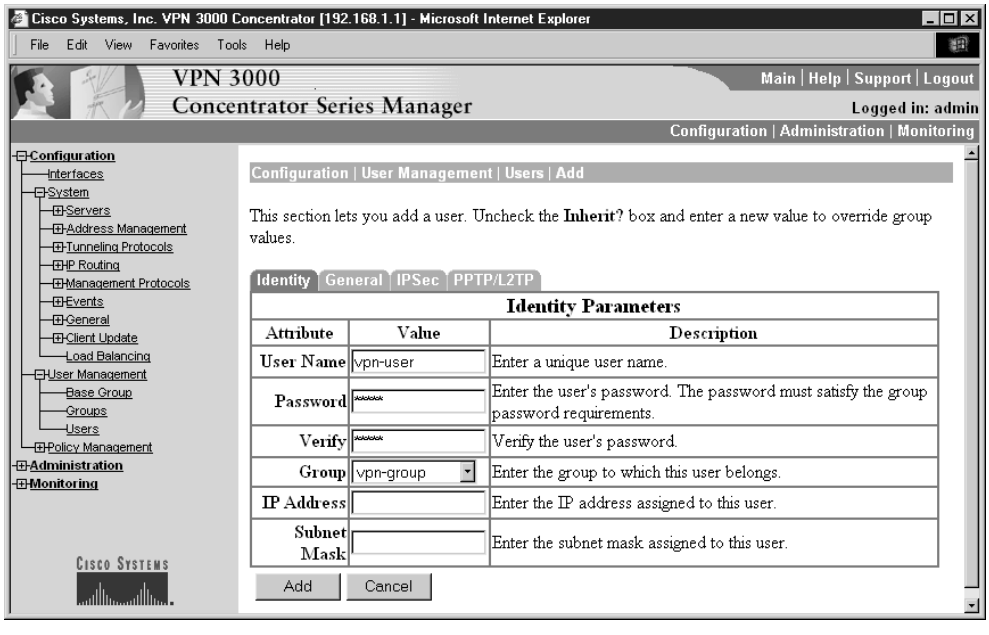

Figure 14-31 Defining a VPN user

Other Configuration Options

This section covers some additional and useful configuration options that are available. Note, other options related to backup servers, load balancing, client update features, and so forth for remote client devices are included at the end of Chapter 15.

These configuration options are accessed from the Configuration menu. Figure 14-32 shows the variety of the configurable settings. The following examples represent the process and the level of granularity attainable.

Configuration | Policy Management | Access Hours

This menu option makes it possible to configure the hours that remote-access groups and users can access the VPN Concentrator. These hours don't apply to LAN-to-LAN connections. Figure 14-33 shows the default screen with the default settings. Never is self-explanatory. Business Hours include Monday through Friday, 9 A.M. to 5 P.M.

Figure 14-34 shows the screen (Configuration | Policy Management | Access Hours | Modify) used either to modify the Business Hours settings or to create another definition based on these settings. To modify existing settings, select the choice on the initial screen

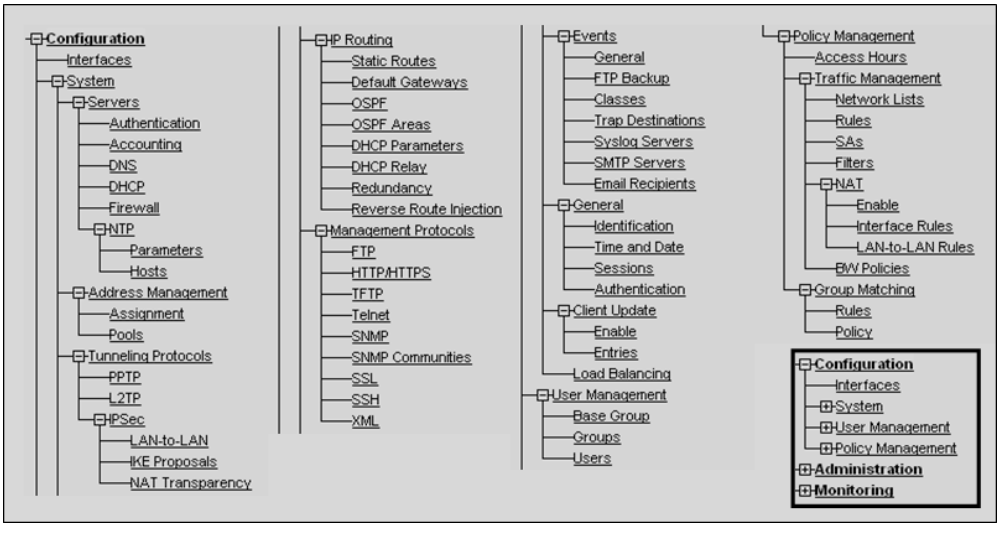

Figure 14-32 Configuration menu options

and click the Modify button. Note, the first column drop-downs allow defining whether to include (during) or exclude the row in the access hours. The Name box at the top could

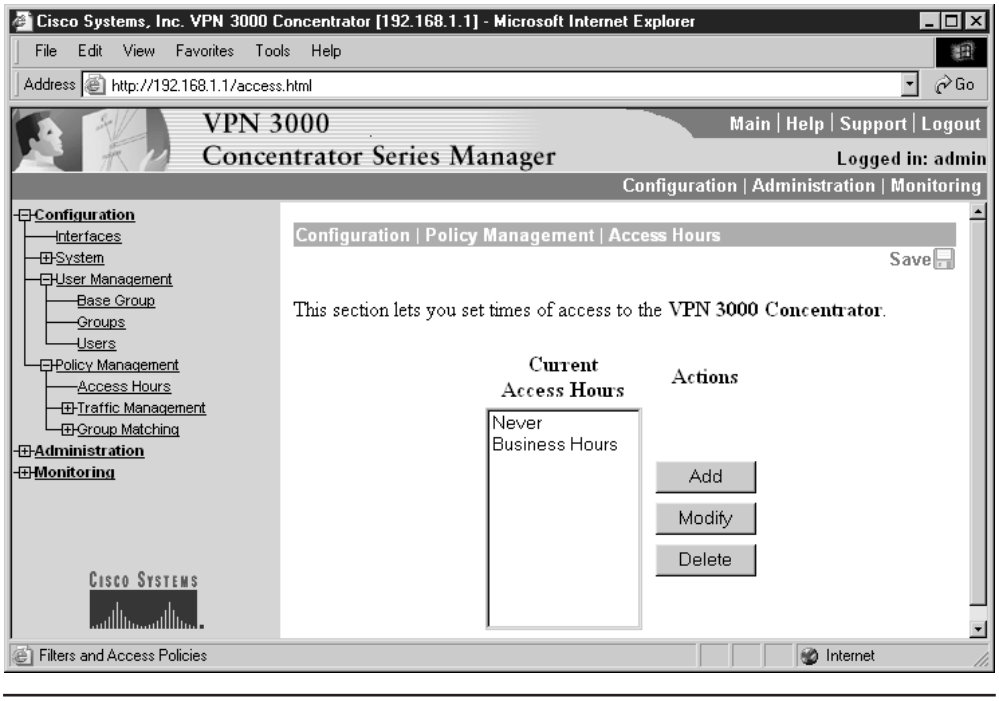

Figure 14-33 Default access definitions

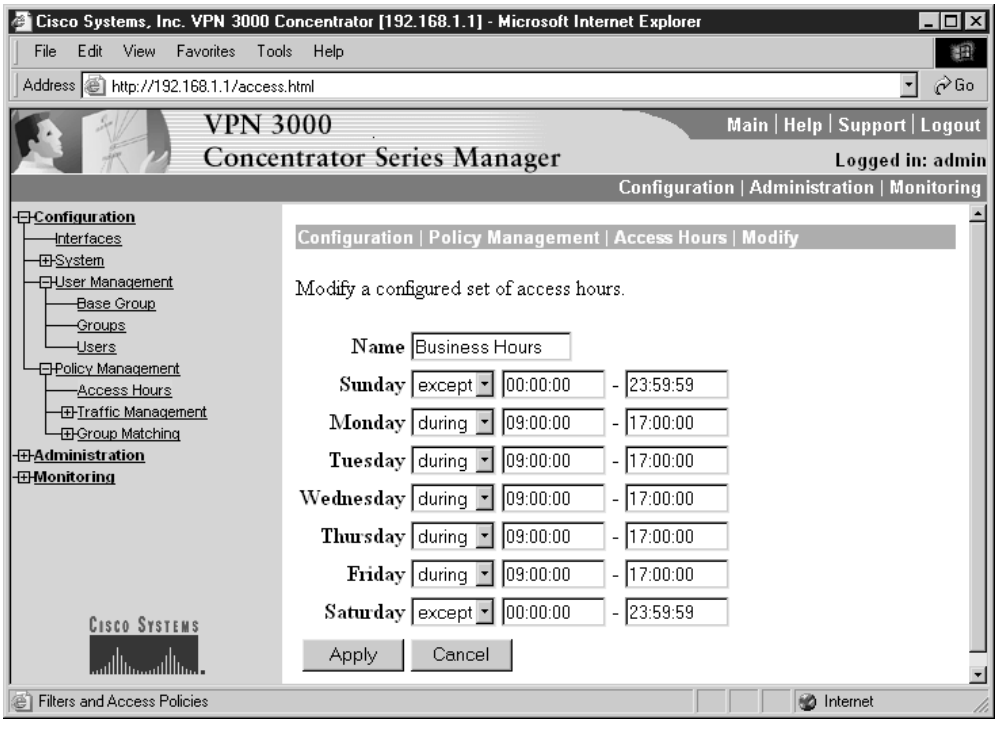

Figure 14-32 Modify or add a new access definition

be used to create a second range East Coast Hours for any groups in another time zone. The Apply button saves the changes, while the Cancel button abandons them.

Figure 14-35 shows the screen (Configuration | User Management | Groups | Modify *grp-name*) used to apply the previous access setting to a specific group. This is a common process in working with group settings where the feature needs to be configured, and then applied to the specific group. The first drop-down allows selecting any defined access time. Note, changing an existing access time definition will impact any groups or users who have that definition assigned to them.

Configuration | System | IP Routing

The VPN Concentrator can be configured to support RIP routing, plus the following routing features to communicate with other routers within the private network and to determine network connectivity, status, and optimum paths for sending data traffic.

- Static Routes—Manually configured route table entries.
- Default Gateways—Routes of last resort for otherwise unroutable traffic.
- OSPF—Open Shortest Path First routing protocol, including Area support.
- DHCP—Define global parameters for DHCP features.

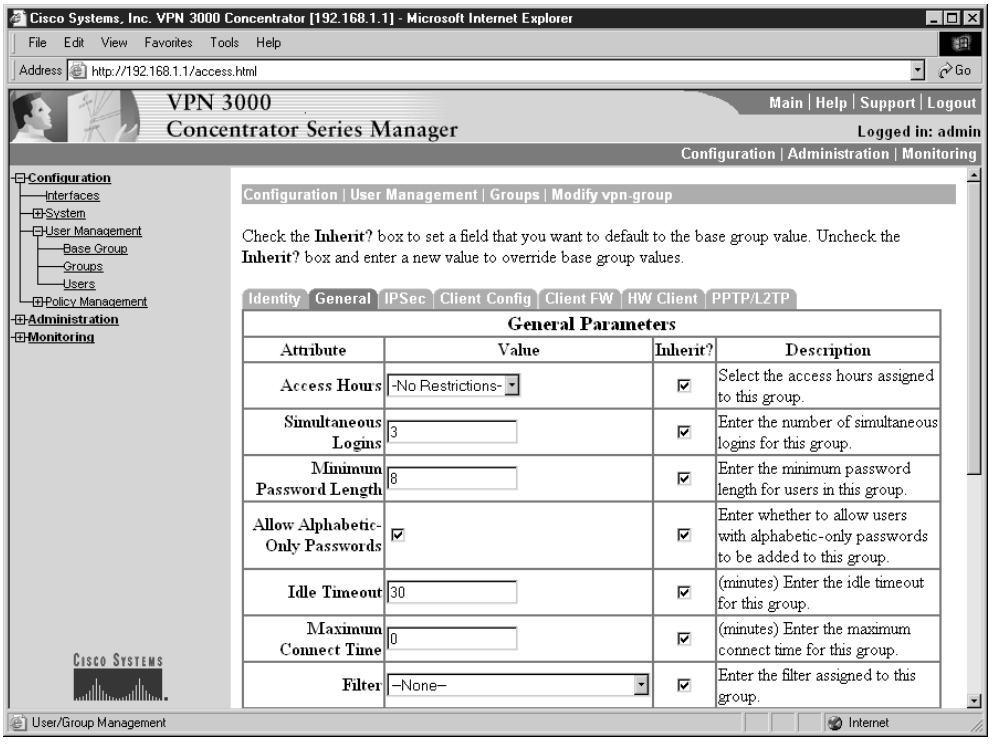

Figure 14-33 Change access hours for a specific group

- Redundancy—Virtual Router Redundancy Protocol parameters.
- Reverse Route Injection (RRI) —VPN Concentrator can add routes to its routing table for remote hardware or software clients.

Configuration | System | Management Protocols

The VPN 3000 Concentrator Series offers the following common built-in servers, supporting a variety of useful protocols used to manage devices and the network:

- FTP—File Transfer Protocol
- HTTP—Hypertext Transfer Protocol
- HTTPS—HTTP over SSL (Secure Sockets Layer) protocol
- SNMP—Simple Network Management Protocol
- SNMP Community Strings—Identifiers for valid SNMP clients
- SSL—Secure Sockets Layer protocol
- SSH—Secure Shell
- Telnet—Terminal emulation protocol and Telnet over SSL

- TFTP—Trivial File Transfer Protocol
- XML—Extensible Markup Language

Digital Certificates

As you learned in Chapters 9 and 11, a *digital certificate* is a form of credential, much like a driver's license or a passport in the paper-based world. Like its paper counterparts, the digital certificate has information on it that identifies the holder, plus the certificate authorities (CA), a trusted third-party that "signs" the certificates to confirm the holder's identity.

A digital certificate has additional information included with the holder's public key, which helps others to verify the key is genuine. This additional information, like a person's picture on a driver's license or a passport, can thwart attempts to substitute an unauthorized public key.

A digital certificate contains the following three items:

- Public key
- Certificate information—Identifying information about the holder, such as IP address, name, serial number, company, department, and so forth
- One or more digital signatures—Supplied by a CA

Certificate Types

CAs issue digital certificates for use in the Public Key Infrastructure (PKI). *PKI* uses public-key/private-key encryption methodology to ensure security. Some terms to remember include the following:

- **CA certificate**—Certificate used to sign (authenticate) other certificates.
- **Root certificate**—CA certificate that's self-signed.
- **Subordinate certificate**—Certificate issued by another CA certificate.
- **Identity certificates**—Certificates for specific systems or hosts.

VPN Concentrator and Certificates

To authenticate using digital certificates, at least one identity certificate and its root certificate must exist on the VPN Concentrator; there could be more. The VPN Concentrator model determines the maximum number of CA and identity certificates allowed.

- **Models 3015–3080—**Maximum of 20 root or subordinate CA certificates and 20 identity certificates.
- **Model 3005—**No more than six root or subordinate CA certificates and two identity certificates.

In both cases, CA certificate maximums include any supporting registration authority (RA) certificates.

All models of VPN Concentrator can have only one SSL certificate installed.

All digital certificates and private keys are automatically stored in the VPN Concentrator's Flash memory. Saving them is unnecessary. These stored items aren't listed and they can't be displayed using the Administration | File Management menu. All stored private keys are encrypted. Once installed on the VPN Concentrator, the identity certificate appears in the Digital Certificate list for configuring both IPSec LAN-to-LAN connections and IPSec SAs.

Certificate Revocation List (CRL)

The VPN Concentrator can be configured to enable CRL information caching in RAM to speed the process of verifying the revocation status of certificates. When the VPN Concentrator needs to check the revocation status of a certificate, it first checks to see if the CRL exists in cache and that it hasn't expired. If the CRL has expired, a new one is requested, but if it hasn't expired, the Concentrator searches the list of revoked serial numbers for the certificate serial number. If a match exists, the authentication fails.

Time Issues

Digital certificates have an expiration date beyond which they're of no value, much like the driver's license and passport examples in the paper-based world. Note, because of this expiration date, the VPN Concentrator time and date must be correct and synchronized with network time.

A second time issue is this: certificate enrollment and installation process must be completed within one week of generating the request. Otherwise, the request is deleted.

Enrolling and Installing Certificates

To use digital certificates for authentication, you must first enroll with a CA, and obtain and install the CA certificate on the VPN Concentrator. Then, you can enroll and install an identity certificate from the same CA. You can enroll and install digital certificates manually or automatically. The automatic method is a new feature that uses Simple Certificate Enrollment Protocol (SCEP), a secure messaging protocol that requires minimal user intervention to enroll and install certificates using only the VPN Concentrator Manager. SCEP was introduced in Chapter 11. SCEP is quicker than enrolling and installing digital certificates manually, but SCEP is available only if it meets the following two conditions:

- The CA must support SCEP.
- Enrolling must be done via the Internet.

If the CA doesn't support SCEP or if digital certificates are enrolled by other means, such as by e-mail or floppy disk, then they must be processed using the manual method, which requires more steps.

In either case, whichever method is used to install a CA certificate must also be used to request identity or SSL certificates from that CA.

Certificate Task Summary

Regardless of whether SCEP or the manual method is used, the following tasks must be completed to obtain and install certificates:

1. Request and install the required CA certificate(s).
2. Create an enrollment request for one or more identity certificates.
3. Request an identity certificate from the same CA that issued the CA certificate(s).
4. Install the identity certificate on the VPN Concentrator.
5. Enable CRL checking and caching.
6. Enable certificates.

Using SCEP to Manage Certificates

The following steps demonstrate using SCEP to enroll and install digital certificates. To use SCEP to enroll identity or SSL certificates, SCEP must also be used to obtain the associated CA certificate. The Manager doesn't allow enrolling a certificate from a CA unless that CA certificate was installed using SCEP. The certificate obtained using SCEP can issue other SCEP certificates and is, therefore, referred to as *SCEP-enabled*.

Using SCEP to Obtain and Install CA Certificates Automatically

Follow these steps for each CA Certificate you want to obtain:

1. Use the Concentrator Manager navigation system to display the Administration | Certificate Management screen, as shown in Figure 14-36.
2. Click the *Click here to install a CA certificate* option at the top of the screen. The Administration | Certificate Management | Install | CA Certificate screen appears, as shown in Figure 14-37.

   The previous link option is only available on this screen if no CA certificates have been installed on the Concentrator. If the link is missing, click the *Click here to install a certificate* option, the third link in the last figure. The Administration | Certificate Management | Install screen is displayed, from which you can choose Install CA Certificate.
3. Click the SCEP (Simple Certificate Enrollment Protocol) link to display the Administration | Certificate Management | Install | CA Certificate | SCEP screen, shown in Figure 14-38. Enter the following information in the two fields:

   URL—The URL of the CA's SCEP interface.

   CA Descriptor—Some CAs require and provide a descriptor to identify a certificate. If the CA doesn't use a descriptor, enter one of your own. Something must be entered in this field.

   Click Retrieve.

   Once complete, the CA certificate is installed on the Concentrator and appears in the Certificate Authorities box of the Administration | Certificate Management screen (as shown in the previous Figure 14-36).

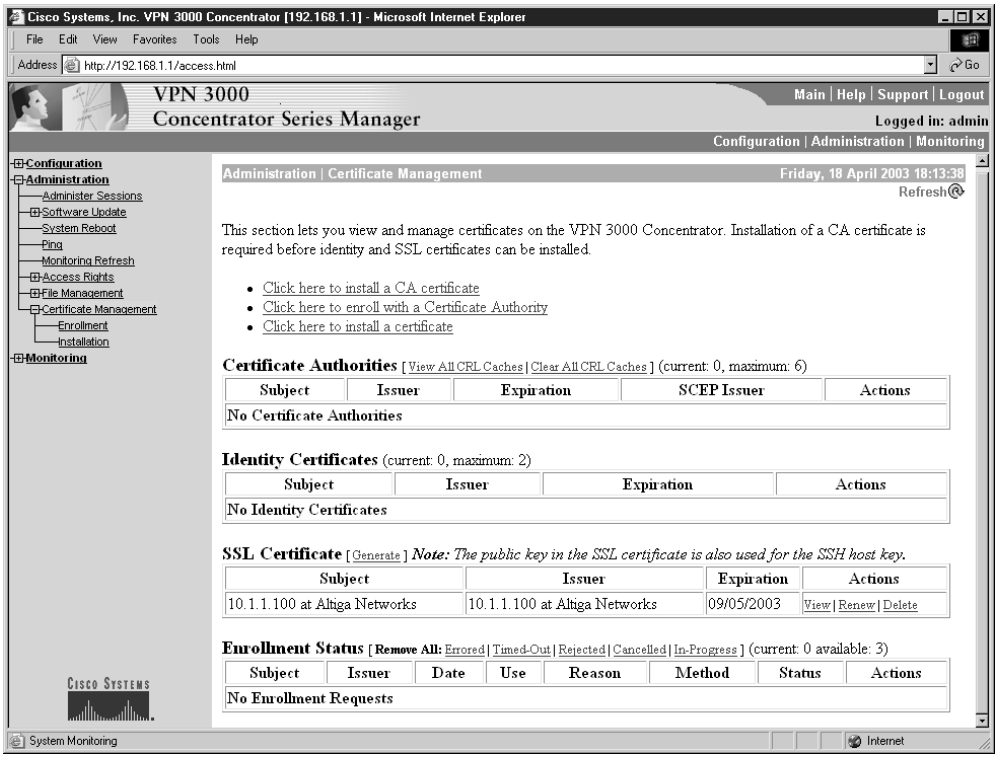

Figure 14-34 Certificate management screen

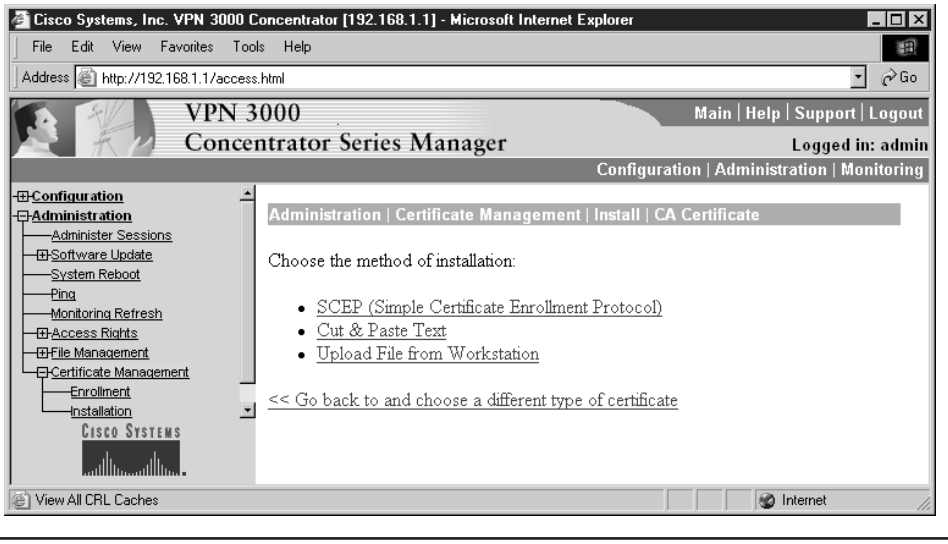

Figure 14-35 Install CA Certificate screen

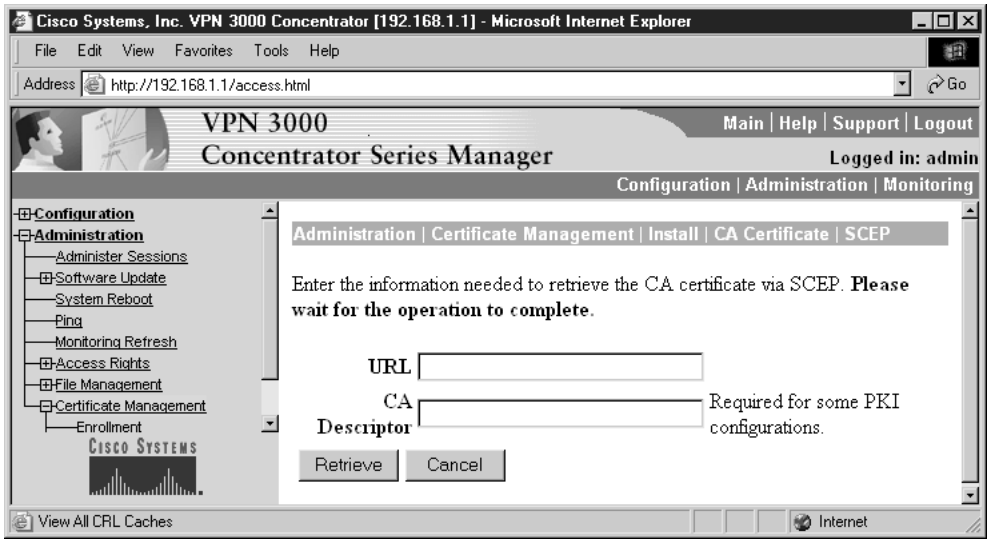

Figure 14-36 CA certificate request information

Using SCEP to Enroll and Install Identity Certificates Automatically

Follow these steps for each identity certificate you want to obtain:

1. Using the Administration | Certificate Management screen from the previous Figure 14-36, Click the *Click here to enroll with a Certificate Authority* link.The Administration | Certificate Management | Enroll screen displays, as shown in Figure 14-39.
2. Click the *Identity Certificate* link to display the Administration | Certificate Management | Enroll | Identity Certificate screen, as shown in Figure 14-40. If SCEP-enabled CA certificates were on the VPN Concentrator, they would be listed as links beneath the Enroll via PKCS10 Request (Manual) shown in the figure.

   The link title includes the name of the CA certificate in the following format: Enroll via SCEP at Certificate Name. So, a CA certificate on the Concentrator named "CA-Test" would look like the following:

   - Enroll via PKCS10 Request (Manual).
   - Enroll via SCEP at CA-Test.
3. Click the link to the SCEP certificate to be enrolled and the Administration | Certificate Management | Enroll | Identity Certificate | SCEP screen display, as shown in Figure 14-41.

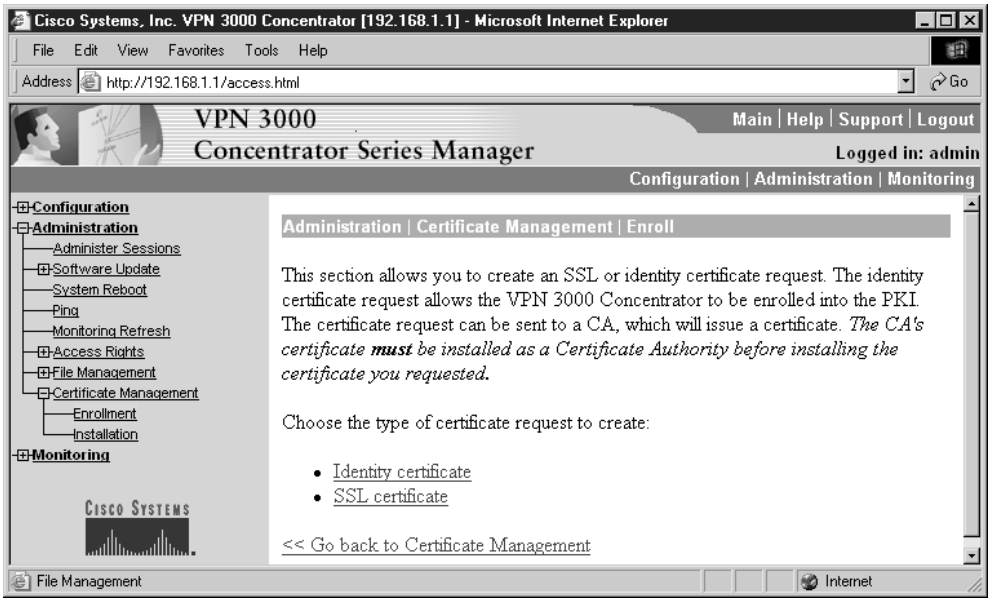

Figure 14-37 Certificate management enrollment screen

4. Complete the fields and click the Enroll button.

   Some CAs require manual verification of credentials and this can take some time—the certificate request could enter Polling mode. In this case, the

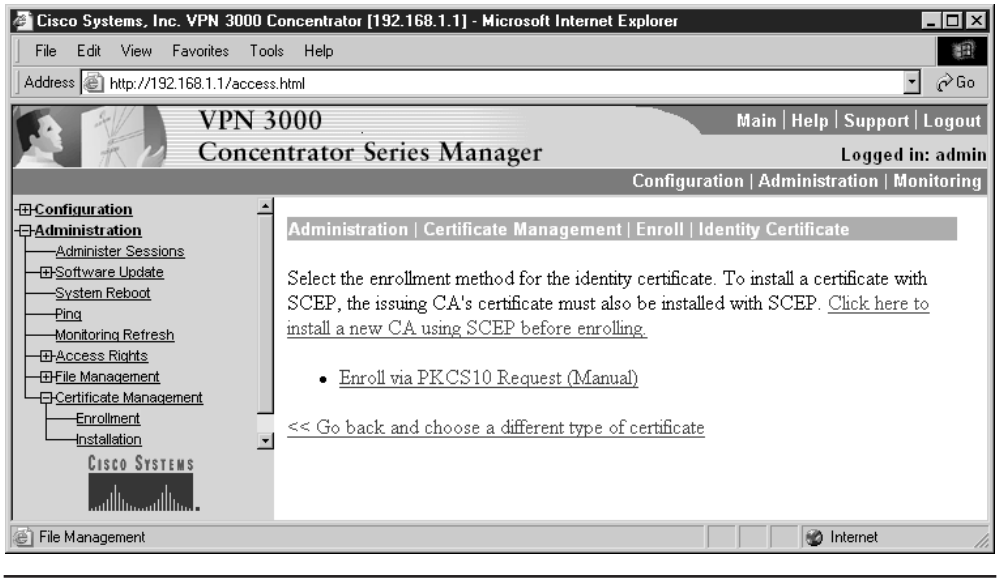

Figure 14-38 Enrollment Identity screen to select a certificate

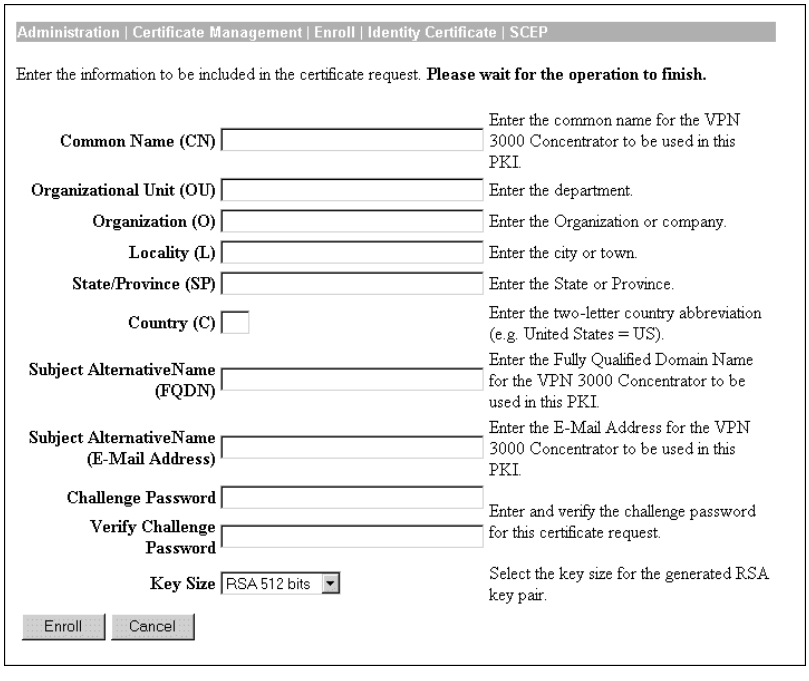

Figure 14-39 Screen to add certificate enrollment information

Concentrator will resend the request to the CA a defined number of times, until either the CA responds or the process times out.

Once the CA responds and issues the certificate, the VPN Concentrator installs it automatically and displays the Administration | Certificate Management | Enrollment | Request Generated screen, as shown in Figure 14-42.

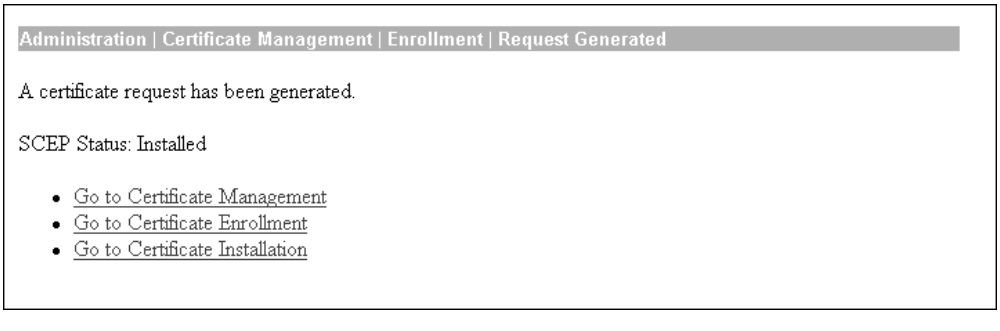

Figure 14-40 SCEP Status: Installed

Using the Certificates

Once the certificate is installed on the VPN concentrator, you must change settings for IKE negotiation. This requires two screen entries, the IKE transform to be used and the IPSec SA information.

IKE Configuration

Use the Manager navigation to locate the Configuration | System | Tunneling Protocols IPSec | IKE Proposals screen, shown in Figure 14-43. This screen displays both the Active and Inactive IKE options available on the Concentrator.

You can change an existing Active proposal from preshared keys to certificates or create a new one. Select an existing proposal, and then click the Modify button or click the Add button. Either way, a screen similar to the one shown in Figure 14-44 appears.

You only have one choice here. Use the Authentication mode drop-down list and select RSA Digital Certificate. Then click the Apply button.

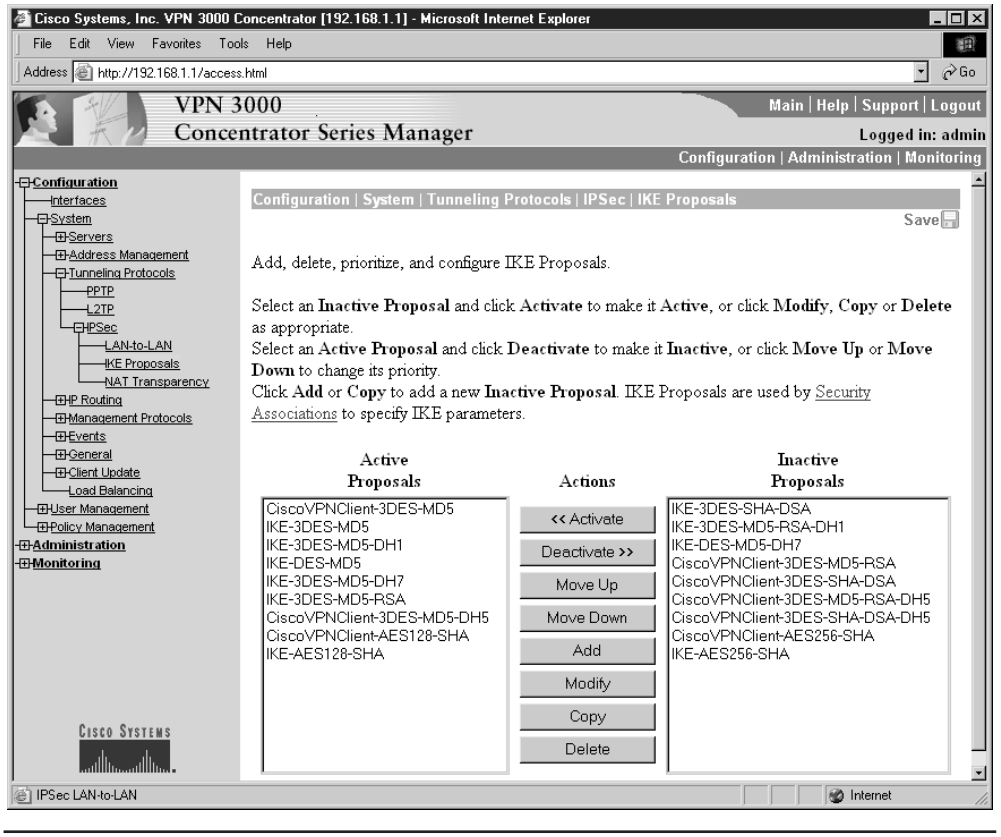

Figure 14-41 IKE Proposal options

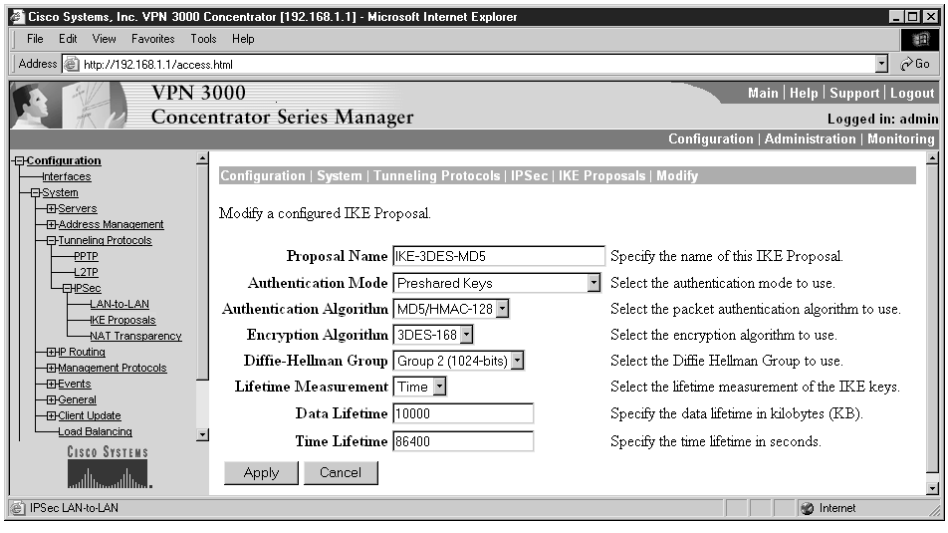

Figure 14-42 IKE Proposal option to be modified

IPSec Configuration

Use the Manager navigation to locate the Configuration | Policy Management | Traffic Management | Security Associations | Modify screen for the appropriate IPSec SA. The resulting screen is large, but the bottom panel, as shown in Figure 14-45, is all that must be changed.

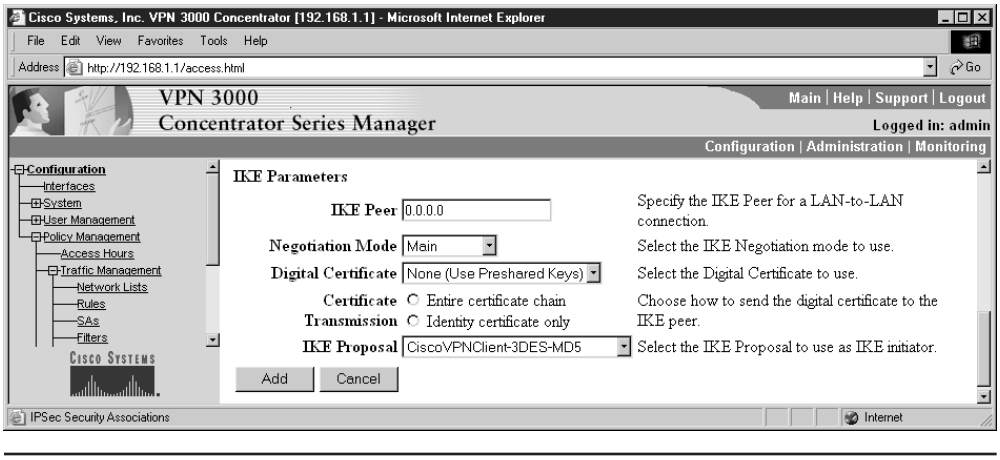

Figure 14-43 Defining the IKE parameters for the IPSec SA

Use the Digital Certificate drop-down list to select the appropriate certificate name.

If necessary, use the IKE Proposal drop-down list to select the IKE proposal defined in the last section, and then click Apply.

Configure Cisco VPN Client Support

Now, it's necessary to configure the Cisco VPN Client. While the full process was covered in Chapter 12, the following steps are used to contact the VPN Concentrator. Bring up the VPN Client on the PC using the Start | Programs | Cisco Systems VPN Client | VPN Dialer.

1. Click the New button.
2. Name the connection, as shown in Figure 14-46, and click the Next button.
3. Enter the Group name just created, and then click the Next button.
4. Enter the IP address of the public interface of the Concentrator.
5. Click the Finish button.
6. Click the Connect button, as shown in Figure 14-47, to establish the connection.
7. When prompted, enter the user name and the password, and then click OK.

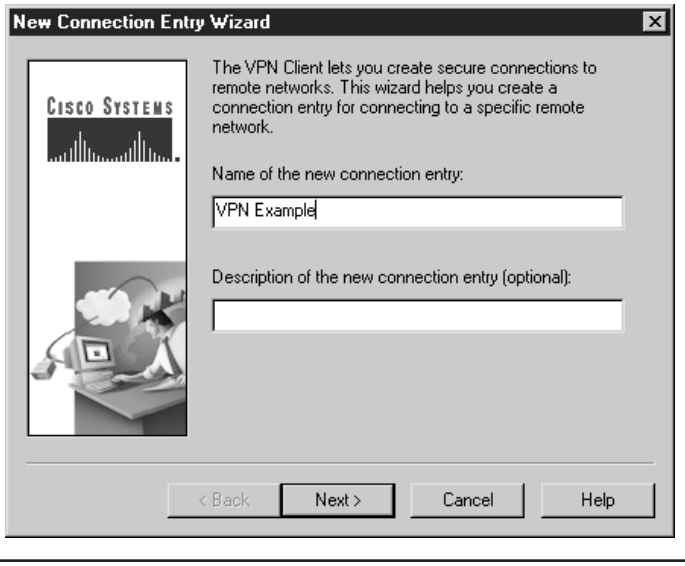

Figure 14-44 New Connection Entry Wizard box

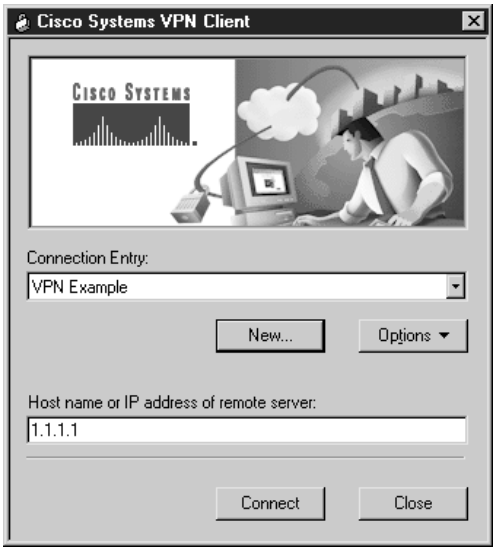

Figure 14-45 Cisco Systems VPN Client box

VPN Client Autoinitiation Feature

VPN Client Autoinitiation (Automatic VPN initiation) is a new feature that provides secure connections to hosts using a wireless LAN (WLAN) environment by connecting through a VPN 3000 Series Concentrator. With autoinitiation configured on the VPN Client, the Client becomes active immediately after the PC boots up, or after exiting Standby or Hibernation mode. The client establishes a VPN tunnel to the Concentrator defined for its network, prompts the user to authenticate, and allows that user network access.

In the WLAN network, the wireless client first associates itself to a wireless Access Point (AP). The installed VPN Client uses the IP address range it receives from the wireless connection to launch a VPN connection request automatically to the corresponding VPN Concentrator on site. The resulting IPSec VPN connection provides secure wireless 802.11*x* traffic for the wireless host. Without a successful VPN connection, the wireless host won't have access to the network resources.

The vpnclient.ini File

Currently, no wizard exists to perform this configuration, so it's necessary to edit the vpnclient.ini file manually for the VPN Client to activate autoinitiation. This file is

located in the VPN Client folder, under Program Files in the Windows environment. The file created previously in Chapter 12 looks like the following:

```
[main]
StatefulFirewall=1
EnableLog=1
[LOG.IKE]
LogLevel=1
[LOG.CM]
LogLevel=1
[LOG.PPP]
```

The changes can be made to the [Main] section by double clicking the file name. The file will then open in Notepad. Saving a copy before you begin might be a good idea.

Preparation

As with any configuration, gathering the needed information before you begin makes sense. The following information is needed to configure autoinitiation.

- The network IP addresses for the client network
- The subnet mask for the client network
- The names for all connection entries users are using for their connections

Configuration

To configure autoinitiation, you need to add the following three keywords and appropriate values in the [Main] section of the vpnclient.ini file:

- **AutoInitiationEnable**—enables or disables autoinitiation. 1 = enable, 0 = disable.
- **AutoInitiationRetryInterval**—defines the number of minutes to wait before retrying the autoinitiation connection. Range is one to ten minutes. The default is one minute.
- **AutoInitiationList**—defines a series of section names that follow this entry. Each one contains the network details needed to autoinitiate. Entries include network address, subnet mask, and a connection entry name, specifying a connection entry profile (.pcf file). You can have a maximum of 64 section (network) entries.

Next, you need to define the networks listed that are associated with the section names in the AutoInitiationList section. While the following is a simple example, it shows enabling autoinitiation for two networks. This feature could represent a significant

advantage in a corporate environment, where some personnel frequently work in multiple locations.

```
[main]
AutoInitiationEnable=1                (turns the feature on)
AutoInitiationRetryInterval=3         (sets a 3 minute wait)
AutoInitiationList=TacomaWLAN,
ViennaWLAN                            (identifies 2 WLANs)
[TacomaWLAN]                          (config info for Tacoma)
Network=10.95.254.0
Mask=255.255.255.0
ConnectionEntry=TacProf               (connection profile named TacProf.pcf)
[ViennaWLAN]                          (config info for Vienna)
Network=192.168.1.0
Mask=255.255.255.0
ConnectionEntry=VieProf               (connection profile named VieProf.pcf)
StatefulFirewall=1
EnableLog=1
[LOG.IKE]
LogLevel=1
[LOG.CM]
LogLevel=1
[LOG.PPP]
```


VPN 3000 Configuration

The configuration steps for this feature required on the 3000 Concentrator are the same as any other VPN Client group. Defining a new group and confirming that all Security Policy issues are addressed and enabled might be wise. This can also make administering and monitoring these users easier, as well as providing autoupdate configurations, if necessary.

Administer and Monitor Remote Access Networks

This section looks briefly at the Administration and Monitoring features of the Cisco VPN 3000 Concentrator Series.

Administration

The VPN 3000 Concentrator Series provides a rich set of administration tools and features that keep the system operational and secure. Configuring the system sets the parameters that govern its use and functionality as a VPN device, but administration involves higher-level activities, such as who is allowed to configure the system and what software runs on it. Only those logged in as administrators can access and use the Administration tools. Figure 14-48 shows a breakdown of the Administration menu.

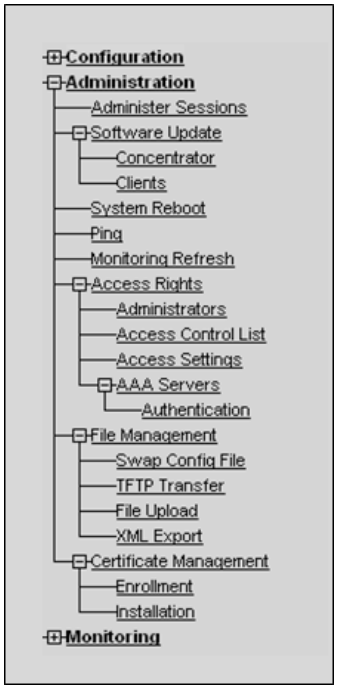

Figure 14-46 VPN 3000 Administration menu

Earlier sections in this chapter looked at some of the Administration features, such as managing Concentrator access hours and digital certificates. This section briefly discusses some other features that might be useful in a remote access environment.

Summarizing VPN Activity

The VPN 3000 Concentrator Series provides a recap of the current sessions currently underway by using two similar screens, the Administration | Administer Sessions screen option and the Monitoring | Sessions screen. Figure 14-49 shows the Administration | Administer Sessions screen.

Both screens have the Group drop-down box that allows looking at only a single group. Both screens have two additional outputs rows, which are not shown in the previous figure. One is a detail view of the current Remote Access sessions and the other is a detail view of the current administrative sessions.

The primary difference between the views is this: the Administration | Administer Sessions screen has links, just below the Group drop-down box, which can be used to log out of all active sessions of a given tunnel type at once. This could be handy if new

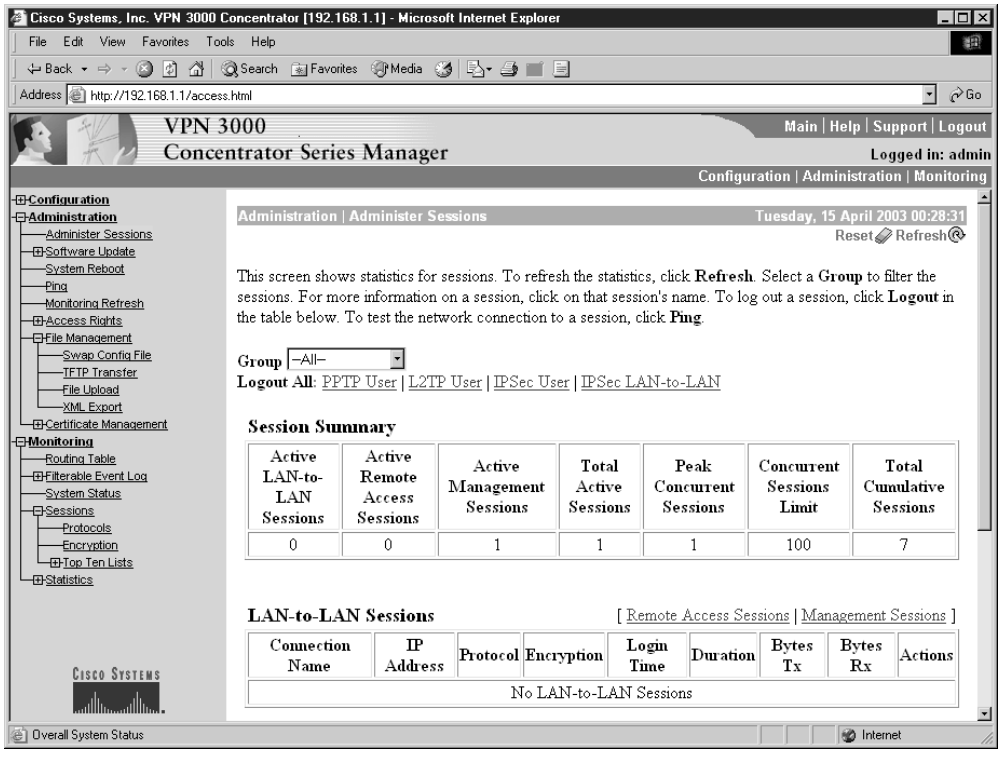

Figure 14-47 The Administration | Administer Sessions screen summarizing VPN activity

security parameters were configured and you decide to force all current sessions to comply, instead of waiting for the next session.

To log out of the sessions, click the appropriate label. The Manager displays a prompt to confirm the action. This action immediately terminates all sessions of the given tunnel type. No user warning or Undo option occurs.

Ping

The Administration | Ping screen, shown in Figure 14-50, lets you use the ICMP ping utility to test network connectivity. This is most useful when working with and troubleshooting remote user connections. The VPN Concentrator sends an ICMP Echo Request message to the defined host. If the host is reachable, the screen displays a simple IP address is alive message, such as *192.168.1.20 is alive.* If the host is unreachable, the Manager displays an Error message. You can also ping hosts from the Administration | Administer Sessions screen.

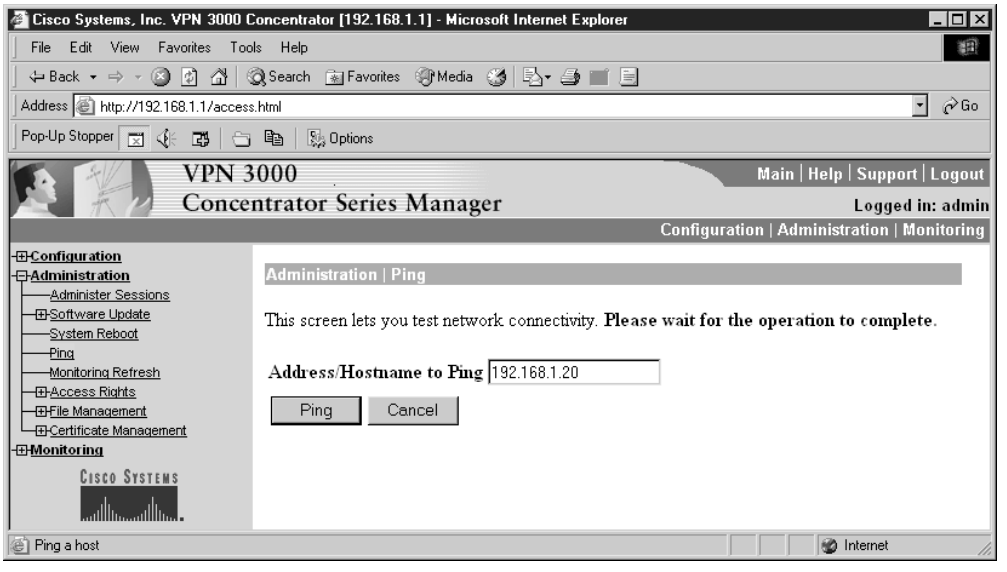

Figure 14-48 Ping screen

System Rebooting

The Administration | System Reboot screen, as shown in Figure 14-51, allows the administrator to reboot or shut down the VPN Concentrator with various options. Rebooting the system automatically logs you out and displays the main login screen. If the browser appears to hang during a reboot, preventing you from logging in, wait a minute for the reboot to finish.

The choices should be self-explanatory.

Software Update

The Administration | Software Update screen has only two links that allow the administrator to update either the VPN Concentrator executable system software or the VPN Client software. The two links are as follows:

- **Concentrator**—Uploads the software image to the VPN Concentrator
- **Client**—Updates the VPN 3002 Hardware Client software

Clicking the Concentrator link brings up the Administration | Software Update | Concentrator screen, as shown in Figure 14-52. The process uploads the executable system software to the VPN Concentrator, which then verifies the integrity of the software image.

To specify the new software file, enter the complete path name of the new image file or click the Browse... button to find and select the file from the workstation or network.

This process can take a few minutes to upload and verify the software. The system will display a simple progress bar.

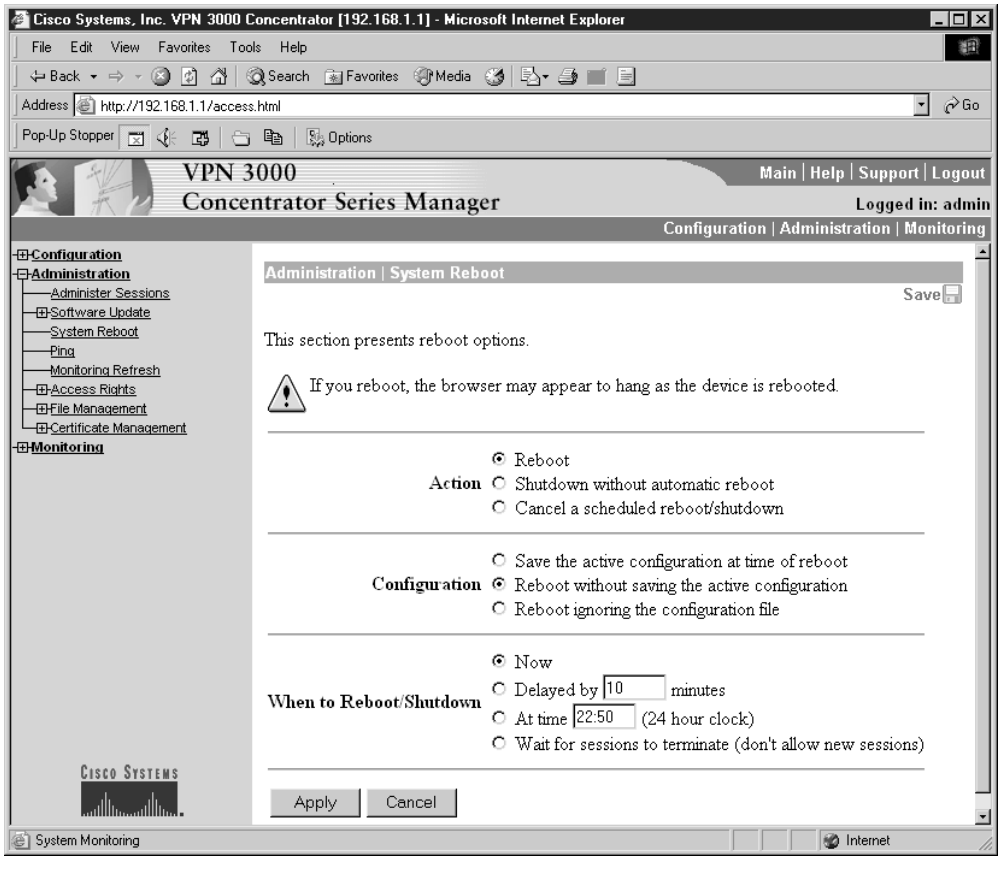

Figure 14-49 System reboot and shutdown options

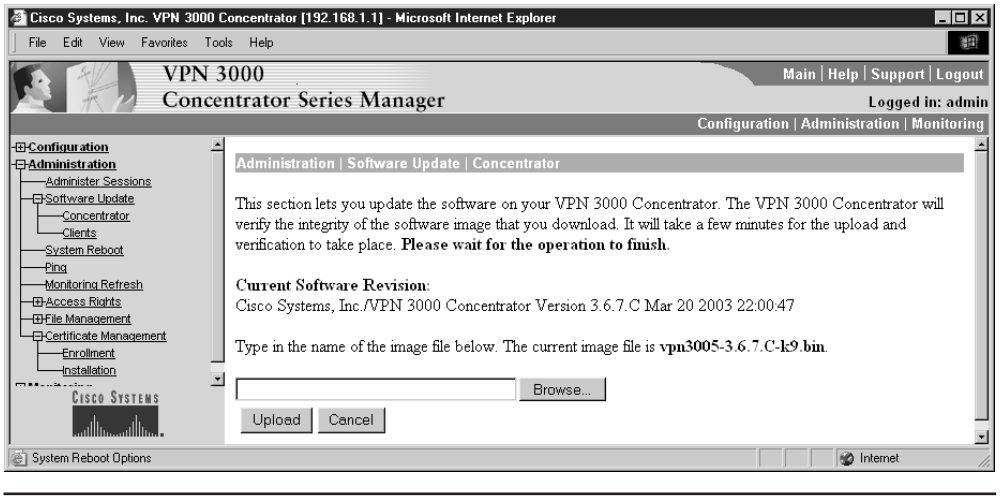

Figure 14-50 VPN Concentrator software upgrade screen

PART III

You must reboot the VPN Concentrator to run the new software image. The system prompts you to reboot when the update is finished.

Updating the Client software is covered in Chapter 15.

Monitoring

The VPN 3000 Concentrator compiles many statistics, and it tracks the status of many processes and activities essential to system administration and management. The Monitoring windows can be used to view the status items and statistics. You can see the state of LEDs that show the status of hardware subsystems in the device. Figure 14-53 shows a breakdown of the Monitoring menu.

The Monitor menu option provides opportunities to see snapshot summaries of activity broken down by protocol or encryption type. An option, Monitoring | Sessions | Top Ten Lists, shows statistics for the top ten currently active sessions, sorted by the following:

- Data—total bytes transmitted and received
- Duration—total time connected
- Throughput—average throughput (bytes/sec)

Routing Table

The Monitoring | Routing Table screen displays the current VPN Concentrator routing table. As with the routers, the routing table shows the best valid forwarding paths the system knows about. These routes can be static routes, or learned via routing protocols or interface configurations. Figure 14-54 shows a simple route table with related information.

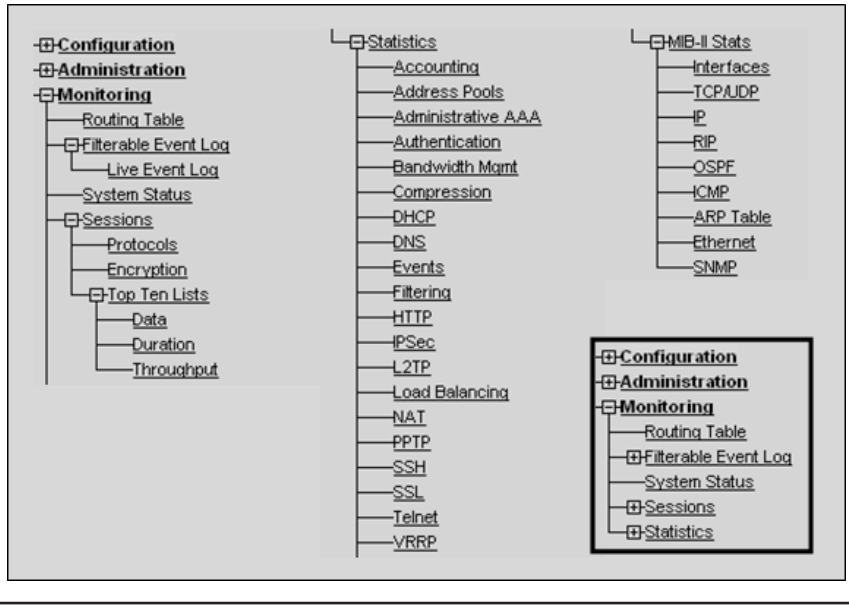

Figure 14-51 Monitoring menu options

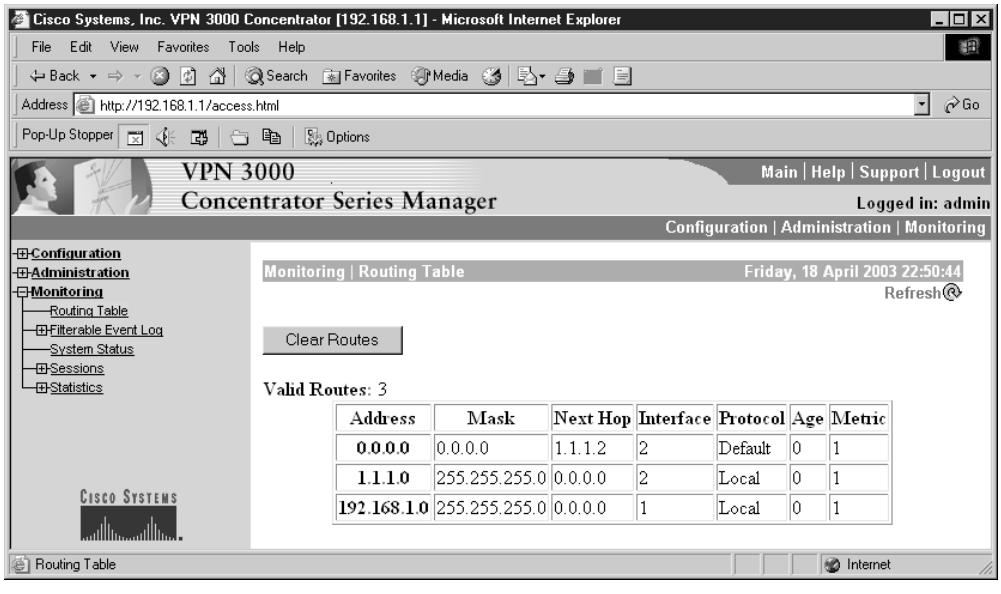

Figure 14-52 Routing table entries

Additional monitoring options are included in Chapter 15.

Chapter Review

The Cisco VPN 3000 Concentrator devices are a series of specialty appliances that perform VPN gateway services for organizations of all sizes. The VPN Concentrator can be used in remote access implementations, providing secure connections for mobile users and Small Office/Home Office (SOHO) locations, as well as LAN-to-LAN connections providing VPN access between branch locations. Many networks use the Concentrators for both types of networks.

Concentrator configuration and operations are via a menu-driven architecture, accessed by either a text-based CLI or a web-based interface. Generally, all tasks and functions can be performed using either interface, but the CLI must be used initially to configure an IP address on the private interface to allow web access.

The Concentrator functions and menu options break into three areas: Configuration, Administration, and Monitoring. Configuring the system sets the parameters that govern its use and functionality as a VPN device, but administration involves higher level activities, such as who is allowed to configure the system and what software runs on it. The Monitoring screens can be used to view the status of the many processes and activities essential to system administration and management, as well as the statistics the Concentrator compiles.

Questions

1. To make the VPN Concentrator accessible to a web browser, which interface must be configured by the CLI?
   - **A.** Ethernet 0
   - **B.** Ethernet 1
   - **C.** Ethernet 2
   - **D.** Ethernet 3
2. The VPN 3000 Concentrator CLI console connection uses which of the following?
   - **A.** A standard Cisco console kit with cable
   - **B.** A rollover cable with no adapters required
   - **C.** A straight-through serial connection
   - **D.** A straight-through RJ-45 cable with no adapters required
3. What is the following screen?

   Welcome to
   Cisco Systems
   VPN 3000 Concentrator Series
   Command Line Interface
   Copyright (C) 1998-2003 Cisco Systems, Inc.
   -- : Set the time on your device. ...
   > Time
   Quick -> [10:13:37]

   - **A.** CLI startup menu
   - **B.** Manager startup menu
   - **C.** CLI Quick Configuration
   - **D.** Manager Quick Configuration
4. Which of the following is *not* one of the Quick Configuration steps?
   - **A.** Define which tunneling protocols and encryption options will be used
   - **B.** Change the admin password to improve system security
   - **C.** Define the method(s) for assigning IP addresses to protected clients
   - **D.** Define the IP routing method to be used
5. Which of the following wouldn't a VPN Concentrator normally connect to in a remote access implementation?
   - **A.** VPN 3002 Client
   - **B.** VPN Client software

C. VPN 3000 Concentrator

D. PIX Firewall

6. Which is *not* one of the three types of preshared keys?

A. Unique

B. User

C. Group

D. Wildcard

7. When using the VPN Concentrator internal authentication server, which is *not* an accurate maximum number of entries (groups and users combined)?

A. Model 3005/3015—100

B. Model 3030—500

C. Model 3060—800

D. Model 3080—1,000

8. Which of the following is *not* one of the four possible methods for the VPN Concentrator to assign IP addresses to the remote users?

A. Use Address from Authentication Server

B. Use a DHCP Server

C. Use NAT inside

D. Use Address Pools

E. Use Client Address

9. What does the term split tunneling refer to?

A. The capability to establish multiple simultaneous VPN connections

B. The capability to allow multiple users to share a VPN connection

C. The capability to allow the IPSec client to go directly to the Internet in Cleartext form for those destinations that don't require encryption

D. The capability to create different VPN connections based on security requirements

10. Which statement is true about VPN Concentrator client firewall requirements?

A. Used to require a personal firewall for all VPN clients

B. Used to configure firewall parameters for VPN Clients running on PIX firewalls

C. Used to configure firewall parameters for VPN Clients running on non-Windows PCs

D. Used to configure firewall parameters for VPN Clients running Microsoft Windows

11. Which is *not* a supported type of user authentication server?
    - **A.** RADIUS server
    - **B.** TACACS+ server
    - **C.** Internal server
    - **D.** RSA Security Inc. SecurID (SDI) server
    - **E.** NT Domain server
12. When using the Identity Parameters tab to define a group, which of the following is *not* included?
    - **A.** User name
    - **B.** Group name
    - **C.** Password
    - **D.** Authentication server type
13. Which routing method is *not* supported using Configuration | System | IP Routing?
    - **A.** Static routes
    - **B.** OSPF
    - **C.** Default gateways
    - **D.** EIGRP
14. Which of the following is *not* a digital certificate type?
    - **A.** Subordinate certificate
    - **B.** CA certificate
    - **C.** Secondary certificate
    - **D.** Root certificate
    - **E.** Identity certificate
15. Which is *not* one of the three keywords added to the [Main] section of the vpnclient.ini file for VPN Client Autoinitiation?
    - **A.** AutoInitiationList
    - **B.** AutoInitiationRetryInterval
    - **C.** AutoInitiationNetworks
    - **D.** AutoInitiationEnable

Answers

1. **B.** Ethernet 1
2. **C.** A straight-through serial connection
3. **C.** CLI Quick Configuration
4. **D.** Define the IP routing method to be used
5. **C.** VPN 3000 Concentrator (used for LAN-to-LAN implementations)
6. **B.** User
7. **C.** Model 3060—800 (should be 1,000)
8. **C.** Use NAT inside
9. **C.** Enables the IPSec client to go directly to the Internet in Cleartext form for those destinations that don't require encryption
10. **D.** Used to configure firewall parameters for VPN Clients running Microsoft Windows
11. **B.** TACACS+ server
12. **A.** User name
13. **D.** EIGRP
14. **A.** Secondary certificate
15. **C.** AutoInitiationNetworks

CHAPTER 15

Configuring Cisco VPN 3002 Remote Clients

In this chapter, you will learn to:

- Configure the VPN 3002 using the CLI
- Configure the VPN 3002 device Remote Access
- Configure a variety of VPN 3000 and 3002 features
- Configure for a backup server and for load balancing
- Configure the client auto-update feature

This chapter looks at configuring the Cisco VPN 3002 Remote Hardware Client for basic operations, as well as advanced features and options. The 3002 device was introduced and described in Chapter 13. That information won't be repeated here, but the features introduced will be explored and the configuration steps defined.

The Cisco VPN 3002 Hardware Client communicates with a VPN 3000 Series Concentrator to create a VPN across a TCP/IP network (such as the Internet). The VPN 3002 requires minimal configuration, and allows the VPN administrator to monitor, configure, and upgrade multiple VPN 3002 hardware clients from a central location.

The VPN 3002 Remote Hardware Client is a small, footprint-dedicated device designed for organizations with many remote users and sites that require secure VPN connections. Figure 15-1 shows the book-sized (2 × 8.6 × 6.5 inches) 3002 device.

The VPN 3002 Hardware Client device can be used to connect one or more devices, including workstations, servers, hubs, cash registers, printers, and IP telephones to a company's central network. Each 3002 unit eliminates the need to install and configure VPN client software on the local workstation(s), plus it supports workstations running any TCP/IP supporting operating system (OS), including Windows, Sun Solaris UNIX, Mac, and Linux.

Currently, the 3002 case supports two models: the VPN 3002 has one public and one private 10/100BASE-T Ethernet interface, and the VPN 3002-8E has one public interface and a built-in 8-port 10/100-Mbps auto-sensing Ethernet switch as its private interface. Each 3002 can connect to LAN switches providing connections for up to 253 host devices. Figure 15-2 shows the back of the both models. The upper unit is the 3002-8E that incorporates the integrated switch, which can provide simplified installation for the

Figure 15-1 Cisco VPN 3002 Client device (front view)

small office or home needing to connect several computers to the network. The lower unit is the VPN 3002, which has a single private interface.

Both models share the following features found in many Cisco devices:

- Motorola PowerPC CPU
- Software-based encryption
- SDRAM memory for normal operation
- NVRAM for storing system configuration parameters
- Flash memory for file management and device OS
- Single power supply

The VPN 3002 in the Network

The VPN 3002 Hardware Client fits into the network anytime a relatively small group of users need secure VPN connections to the corporate network. Figure 15-3 shows both an overall view of the small branch connecting to the corporate network via a VPN 3002 and a more detailed view of the possible local connection using cable or DSL services.

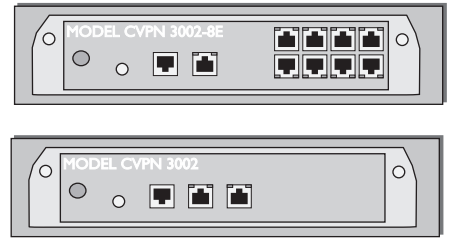

Figure 15-2 Cisco VPN 3002 Client models

VPN Modes

The Cisco VPN 3002 supports two modes of operation to offer implementation choices based on flexibility, security, and easy configuration. Those modes are as follows:

- Client mode
- Network Extension mode

A large VPN implementation might frequently have both types of operation.

Client Mode

In *Client* mode (also called *PAT* mode), the VPN 3002 emulates the VPN client software appearing to the main network like a single remote user, isolating all devices on the VPN 3002 private network from the corporate network. The private hosts protected behind the VPN 3002 are a separate network that remains invisible and can't be routed to by the central site hosts. The local hosts are assigned their IP addresses from the VPN 3002 private interface configured as a DHCP server, while the public network port can be configured to use DHCP client feature to acquire its IP address from an Internet service provider (ISP).

The 3002 device uses Port Address Translation (PAT) on the public interface to help secure (hide) the local network and to allow local hosts to travel out of the network in Client mode. Because all traffic to the central network will have the Public interface IP address, PAT supplies and manages unique port number mappings to be used in combination with the IP address.

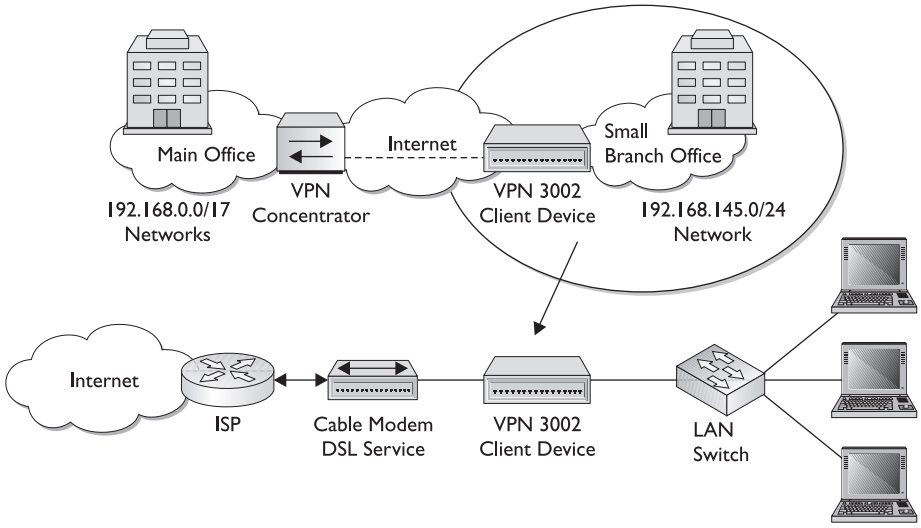

Figure 15-3 VPN 3002 connection overview and detail view

Because the VPN 3002 configured for Client mode can only create outbound connections, there's no way for an outside source—even from the corporate network—to initiate a connection with the 3002 unit or through it to the workstations behind.

Client Mode and Split Tunneling *Split tunneling* provides the capability to have a secure tunnel to the central site, while simultaneously maintaining an unsecured clear-text tunnel to the Internet through the ISP. PAT is used to protect the local workstations during split tunneling to the Internet. The network and addresses on the private side of the VPN 3002 remain hidden and can't be accessed directly from the Internet. If the organization security policy prohibits split tunneling, it can be blocked by creating a policy on the central site device, which is then pushed down to the 3002 Client.

Network Extension Mode

In *Network Extension* mode, the VPN 3002 establishes a secure, site-to-site connection with the central site device. The local stations behind the VPN 3002 are fully routable and the local network is visible to the central site. As the name implies, the local network becomes part of the organization's intranet. VPN and device configuration and security policies are pushed from the central site. The VPN 3002 must initiate the tunnel to the central site but, after the tunnel is up, either side can initiate data exchange.

In Network Extension mode, the private address can be assigned using the DHCP server. Any shared resources in the protected network that must be accessed by the central area hosts should be assigned manually to allow central site hosts and applications to reliably reach any local server, printer, POS terminal, IP phone, or other device critical to the business.

Network Extension Mode and Split Tunneling PAT provides security for local host traffic heading to the Internet through split tunneling. The network and addresses on the private side of the VPN 3002 are accessible over the tunnel, but are protected from the Internet because they can't be accessed directly. This outbound PAT on the VPN 3002 provides centralized security control because no configuration parameters exist for local users to adjust, which might cause the central site to be compromised.

Network Extension Mode per Group VPN software versions 3.6 and later let a network administrator restrict the use of Network Extension mode. The administrator can now enable/disable Network Extension mode on the VPN Concentrator for VPN 3002 hardware clients on a per-group basis.

Network Extension mode is the default setting on the VPN Concentrator. If the concentrator is configured to disallow Network Extension mode for a group, all VPN 3002s in the group must be configured for Client (PAT) mode.

IPSec VPNs

The VPN 3002 Hardware Client supports IPSec for secure connections to a central-site VPN Concentrator over a VPN tunnel. The VPN 3002 Hardware Client, which supports one tunnel at a time, running software release 3.6 or higher, supports the following IPSec implementations, but only one for each tunnel.

- IPSec over TCP
- IPSec over NAT-T
- IPSec over UDP

IPSec over TCP

IPSec over TCP encapsulates encrypted data traffic within TCP packets. This allows the VPN 3002 to operate in networks where standard ESP (Protocol 50) or IKE (UDP 500) can't function, or they can only function by modifying existing firewall rules. IPSec over TCP enables secure tunneling through NAT and PAT devices, and through firewalls by encapsulating both the IKE and IPSec protocols within TCP packets.

To use IPSec over TCP, both the VPN 3002 and the VPN Concentrator must meet the following requirements:

- Run version 3.5 or later software.
- Enable IPSec over TCP.
- Configure both the VPN 3002 and the VPN Concentrator to use the same port for IPSec over TCP.

NOTE IPSec over TCP doesn't work with proxy-based firewalls.

IPSec over NAT-T

NAT Traversal (NAT-T) allows IPSec peers to establish a connection through a device using NAT. NAT-T accomplishes this by encapsulating IPSec traffic in UDP datagrams (port 4500), thereby providing NAT devices with needed port information. NAT-T technology auto-detects any NAT devices and only encapsulates IPSec traffic when necessary.

The VPN 3002 hardware client uses NAT-T by default and requires no special config uration. The VPN 3002 first attempts NAT-T, and then uses IPSec over UDP if a NAT device isn't autodetected. The UDP packets allow IPSec traffic to pass through firewalls, which would normally reject and discard it.

To use NAT-T, the VPN 3002 must meet the following requirements:

- Run version 3.6 or later software.
- Port 4500 on any firewall between the VPN 3002 and the VPN peer must be open.
- Reconfigure any existing IPSec over UDP using port 4500 to a different port.
- Use the Configuration | Interfaces | Public screen to select the second or third options for the Fragmentation Policy parameter. These options let traffic travel across NAT devices that don't support IP fragmentation, while not impeding NAT devices that do support IP fragmentation.

IPSec over UDP

The VPN 3002 supports UDP NAT/Firewall Transparent IPSec. This technology encapsulates encrypted data traffic within UDP packets to provide secure connections between a VPN 3002 and a VPN Concentrator through a device, such as a firewall performing NAT.

The VPN 3002 uses frequent keepalives to ensure the mappings on the NAT device remain active. The VPN 3002 doesn't require special configuration for this feature, but the following minimum requirements must be met.

- Both the VPN 3002 and the VPN Concentrator must be running Release 3.0.3 or higher.
- IPSec over UDP must be enabled on the VPN Concentrator for the group to which the VPN 3002 belongs.

NOTE Cisco technology doesn't currently support a topology with multiple VPN 3002 Hardware Clients behind a single NAT device.

Configuring the 3002 Device

The VPN 3002 has been designed for simplicity and reliability of installation. It has few local setup parameters that must be configured. Basic configuration parameters, security policy, and even device upgrades are "pushed" to the device from the central site (head-end) device with the next connection. The user simply plugs the minimally configured VPN 3002 device into a DSL/cable connection, router, or other wide area networks (WANs) access device at the remote site.

This central control and management approach minimizes the need for skilled users or dedicated IT staff to deploy or maintain the connection. Troubleshooting aids and centralized monitoring features are built into the 3002 software to ensure proper operation after the unit has been set up.

The 3002 supports both a specialized command-line interface (CLI) and a Hardware Client Manager (web-based interface). In reality, they are similar to each other, much like the built-in web interface for routers like the 2600 series. From a practical standpoint, you can do exactly the same tasks with either interface. The primary difference is shown in Figure 15-4, where you can see the Explorer-like program structure on the left side and links to the three program modules in the upper-right corner. These two features make navigating the web-based interface significantly easier and quicker. The Help feature on the Hardware Client Manager is much better than the CLI.

Command-Line Interface (CLI)

The VPN 3002 Hardware Client CLI is a built-in, menu-based configuration, administration, and monitoring system that can be accessed via the system console port or a Telnet (or Telnet over SSL) session. Both Telnet options are enabled by default on the private network interface. The CLI supports the same configuration options as the HTML-

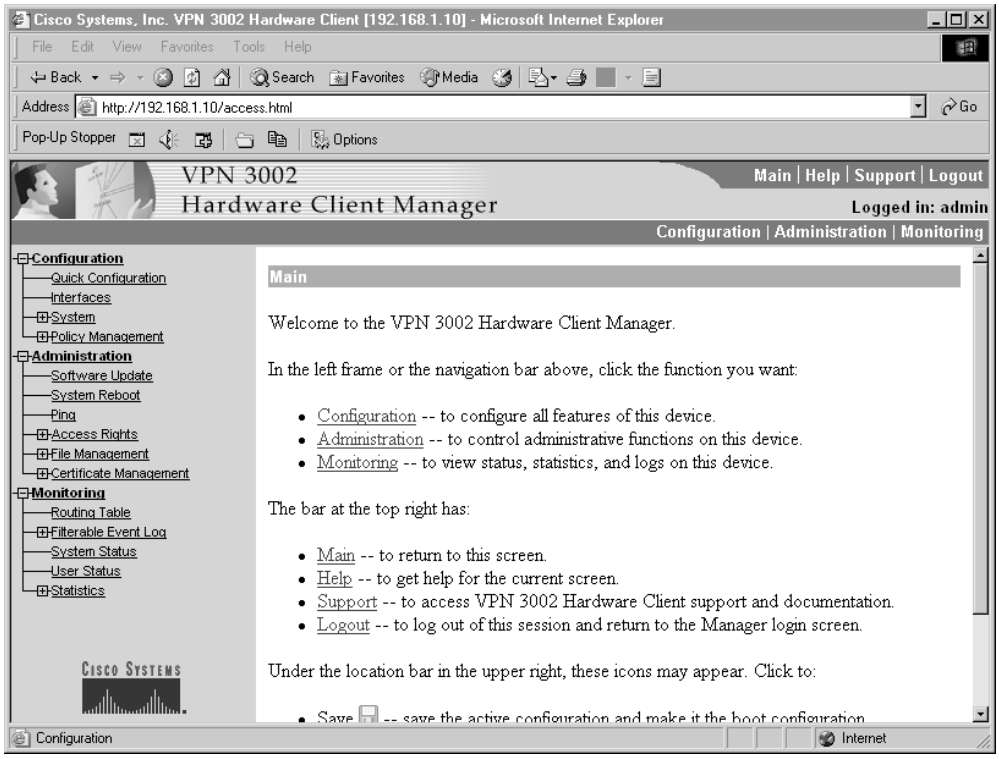

Figure 15-4 VPN 3002 Hardware Client Manager

based VPN 3002 Hardware Client Manager covered in the section "The Hardware Client Manager."

NOTE The VPN 3002 uses a standard Cisco console kit and plugs into a RJ-45 interface on the device. The VPN 3000 concentrators use a straight-through jumper cable and needs one of the RJ-45 to DB-9 converters.

Console port access is similar to the IOS routers using a terminal emulator program, such as HyperTerminal. You might need to press ENTER until the login prompt appears. Login user names and passwords for both console and Telnet access are the same. The factory-supplied default is configured and enabled for administrators using admin for both the login and the password. Entries are case-sensitive. Access and user names/ passwords are set using the Administration | Access Rights | Administrators menus. The following output shows the initial login and main menu:

```
Login: admin
Password:                          (doesn't display)
```

```
                Welcome to
               Cisco Systems
         VPN 3002 Hardware Client
          Command Line Interface
Copyright (C) 1998-2003 Cisco Systems, Inc.

1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main ->
```

Help

The *Help* menu system is somewhat limited, displaying only the following information when 5 is entered at the Main menu. Context-sensitive Help isn't available and the familiar question mark (?) doesn't activate Help features. The Help feature in the Hardware Client Manager is much better, offering context-sensitive assistance like most Windows applications.

```
Main -> 5
Cisco Systems.  Help information for the Command Line Interface

From any menu except the Main menu.
-- 'B' or 'b' for Back to previous menu.
-- 'H' or 'h' for Home back to the main menu.

For Data entry
-- Current values are in '[ ]'s. Just hit 'Enter' to accept value.

1) View Help Again
2) Back

Help ->
```

As it turns out, the *B* and *H* options—and particularly the *H* option—will come in handy when you navigate the Device menus. Pressing *H* returns you to the Main menu.

Saving Configuration File Changes

Configuration and administration changes made using menu options 1 and 2 on the Main menu take effect immediately and become a part of the active, or running, configuration. Like the Cisco routers, if the VPN 3002 is rebooted without saving the active configuration, any changes will be lost.

Saving changes to the system configuration (CONFIG) file is a one-step process from the Main menu. At the Main -> prompt, typing 4 will save changes without additional steps or confirmation.

```
1) Configuration
2) Administration
```

```
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 4
```

The system writes the current (active) configuration to the CONFIG file and redisplays the main menu.

Second Level Menus

Familiarity with the menu system will come with experimentation and experience, but the following examples expand the menus one level. The following output reflects choosing the Configuration (1) option. Notice the prompt changes to reflect the new menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 1

1) Quick Configuration
2) Interface Configuration
3) System Management
4) Policy Management
5) Back

Config ->
```

The following output reflects choosing the Administration (2) option from the Main menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 2

1) Software Update
2) System Reboot
3) Ping
4) Access Rights
5) File Management
6) Certificate Management
7) Back

Admin ->
```

The following output reflects choosing the Monitoring (3) option from the Main menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 3

1) Routing Table
2) Event Log
3) System Status
4) User Status
5) General Statistics
6) Back

Monitor ->
```

Shortcut Numbers

Once you become familiar with the structure of the CLI, you can quickly access any level by entering a series of numbers, corresponding to menu choices, separated by periods. For example, entering 2.2.2.1.2 at the Main-> prompt saves the configuration and reboots the device immediately. The result looks like the following, beginning at the Main menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 2.2.2.1.2

Done

Login:
```

The following are the steps that were fast-forwarded through, beginning at the Main menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
6) Exit

Main -> 2
```

```
1) Software Update
2) System Reboot
3) Ping
4) Access Rights
5) File Management
6) Certificate Management
7) Back

Admin -> 2

1) Cancel Scheduled Reboot/Shutdown
2) Schedule Reboot
3) Schedule Shutdown
4) Back

Admin -> 2

1) Save active Configuration and use it at Reboot
2) Reboot without saving active Configuration file
3) Reboot ignoring the Configuration file
4) Back

Admin -> 1

1) Cancel Scheduled Reboot/Shutdown
2) Reboot Now
3) Reboot in X minutes
4) Reboot at time X
5) Reboot wait for sessions to terminate
6) Back

Admin -> 2

123 03/31/2003 15:41:12.460 SEV=1 REBOOT/1 RPT=1
Reboot scheduled immediately.
Done
```


The Hardware Client Manager (Web Interface)

The VPN 3002 Hardware Client Manager is an HTML-based interface that makes it possible to configure, administer, monitor, and manage the VPN 3002 device with a web browser. The easiest way to use the web interface is to connect to the VPN 3002, using any PC with a web browser on the private network behind the VPN 3002.

By default, the Client Manager uses HTTP, which is convenient, but messages are in clear text. If security requires it, the Client Manager supports a secure, encrypted HTTP connection over Secure Sockets Layer (SSL) protocol, known as HTTPS.

Browser Requirements

The VPN 3002 Hardware Client Manager supports either Microsoft Internet Explorer (IE) version 4.0 or higher or Netscape Navigator version 4.5–4.7. For the best results, Cisco recommends Internet Explorer, and JavaScript and cookies must be enabled in the browser. The other recommendation is that any updates and patches be installed.

Recommended Display Settings

Cisco recommends the following monitor display settings for best viewing:

- Screen area 1,024 × 768 pixels or greater (Minimum 800 × 600 pixels)
- Colors 256 colors or higher

Browser Navigation Toolbar

Earlier implementations of the Client Manager were basically the CLI converted simply to a web interface. Each new version includes much better Windows function integration. *Help*, a Java-based applet, in particular, is getting friendlier and more useful.

Cisco still doesn't recommend using the browser navigation toolbar buttons Back, Forward, or Refresh/Reload with the Client Manager unless specifically instructed to do so. To maintain access security, clicking the Refresh/Reload button automatically logs out the Manager session and returns to the login screen. Using the Back or Forward buttons could possibly display old Manager displays with incorrect data or settings. If you're concerned about this, the IE View | Full screen (F11) feature will eliminate the temptation.

Connecting to the Client Manager

To access the VPN 3002 Client Manager application using HTTP over a web browser, type the VPN 3002 private interface IP address (such as 192.168.1.10) in the browser Address or Location field. The browser will automatically supply the http:// prefix.

The browser displays the VPN 3002 Hardware Client Manager login screen, as shown in Figure 15-5.

Logging in to the Manager application is the same for clear-text HTTP or secure HTTPS. The 3002 supports three types of accounts that can access the device: Administrator, Config, and ISP. Only the Administrator account is enabled by default using admin/admin for the user name/password. Internet Explorer users can use the TAB key to move from field to field. The Clear button can be used to start over.

Figure 15-6 shows the opening screen that appears, offering access to the three main application modules. This screen provides a good overview of the various screen components and options to maneuver through the application. The application tree on the left-hand side offers Explorer-like navigation capabilities to move quickly from feature to feature. This feature alone makes the web interface significantly easier to use than the CLI.

The VPN 3002 Hardware Client Reference, available online or in the CD-ROM documentation that came with the device, covers how to set up the device for installing an SSL Certificate in the browser for HTTPS connectivity.

Client Manager Organization

The Client Manager, exactly like the CLI, is made up of three major sections and many second and third level subsections:

- **Configuration**—Sets all VPN 3002 parameters that govern the unit's use and functionality as a VPN device.

- **Quick Configuration**—A series of steps that supply the minimal parameters needed to make the VPN 3002 operational.
- **Interfaces**—Ethernet parameters for public (outside) and private (inside) interfaces.
- **System**—Sets system-wide function parameters, such as server access, IPSec tunneling protocol, built-in management servers, event handling, IP-routing, and system identification.
- **Policy Management**—Enables PAT and certificate validation.
- **Administration**—Manages the higher-level functions that keep the 3002 unit operational and secure, such as who is allowed to configure the system and what software runs on it, as well as managing its configuration files and digital certificates.
- **Monitoring**—Views routing tables, event logs, system LEDs and status, and statistics and user session data.

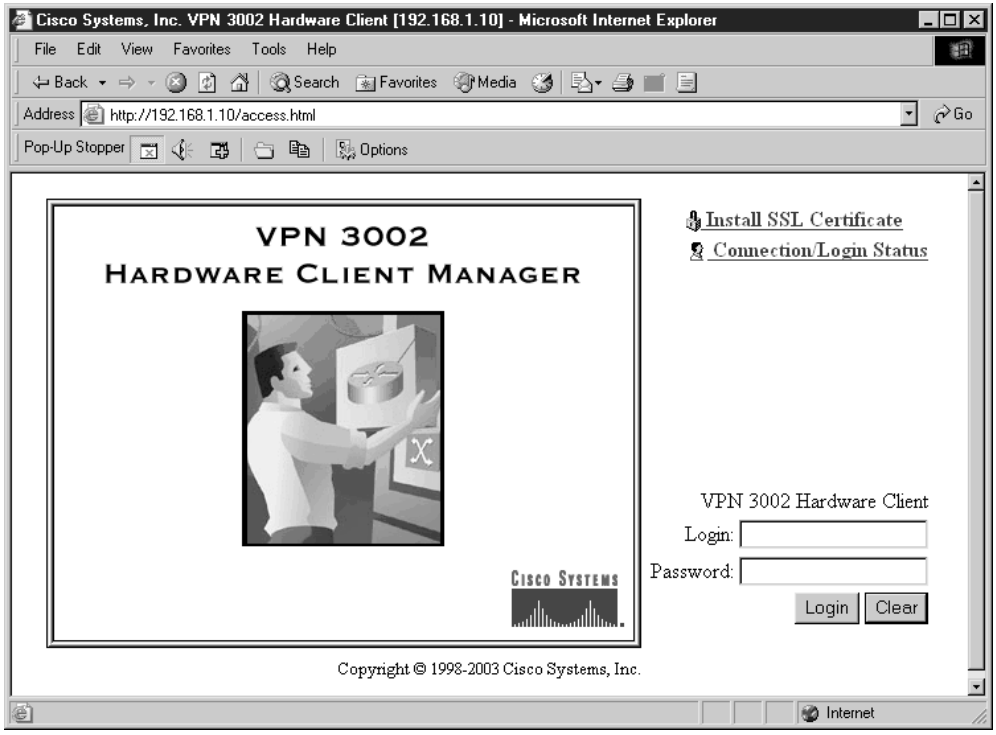

Figure 15-5 VPN 3002 Hardware Client Manager login screen

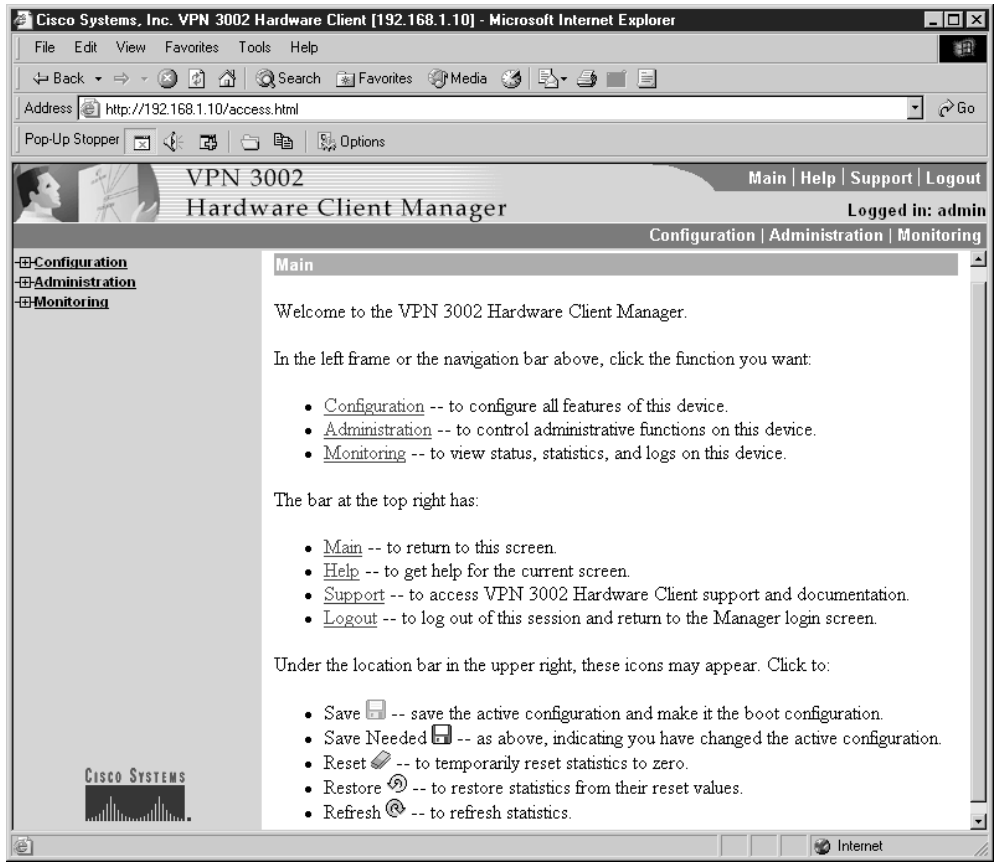

Figure 15-6 Client Manager opening screen

Help

Figure 15-7 shows the result of selecting Administration | Ping in the left panel, and then clicking on the Help button in the upper-right corner. The Help window works much like any Windows help document.

Second-Level Menus

The Client Manager structure tree in the leftmost panel can be expanded using standard Windows techniques. Figure 15-8 shows the three menus expanded and the Configuration menu expanded to three levels.

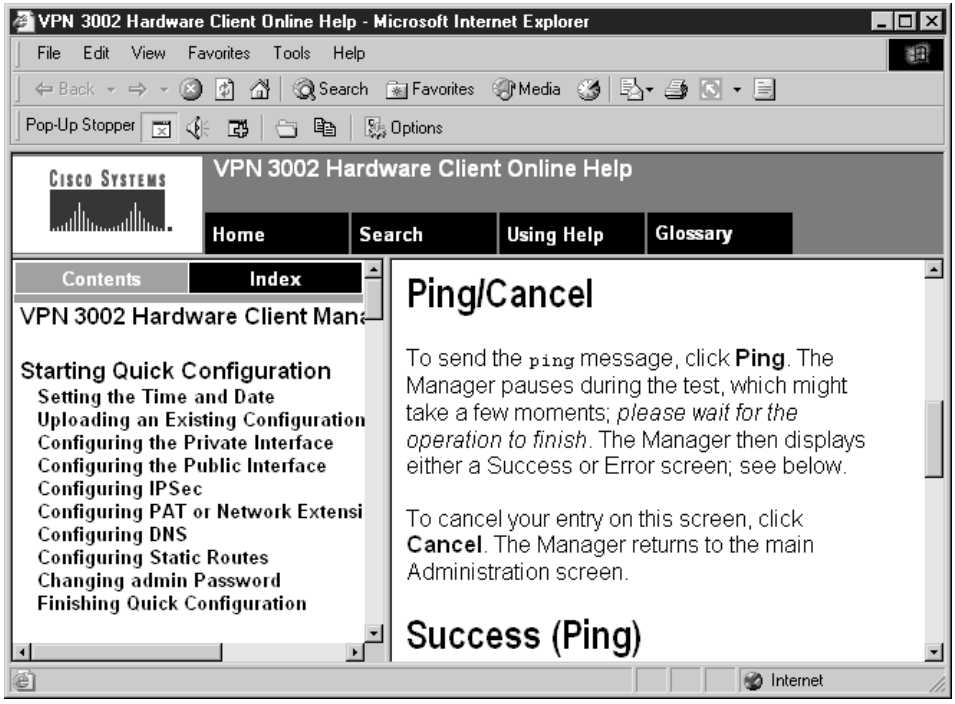

Figure 15-7 VPN 3002 Client Manager help system

Common Configuration Tasks

This section looks briefly at how to use both interfaces to accomplish routine configuration tasks. While conceptually these will be familiar, please remember that neither interface has been aligned with the familiar IOS and this can be a little confusing.

Upgrading the Software

When upgrading with the CLI, you must download the latest OS for the 3002 and store it in the source directory for the TFTP server. Make sure the TFTP server is running. The web interface uses an upload process that only requires the new software is installed on the computer.

The following output shows the upgrade steps to use when working with the CLI, starting at the Main menu:

```
1) Configuration
2) Administration
3) Monitoring
4) Save changes to Config file
5) Help Information
```

```
6) Exit

Main -> 2                                   (select Administration)

1) Software Update
2) System Reboot
3) Ping
4) Access Rights
5) File Management
6) Certificate Management
7) Back

Admin -> 3                  (select Ping to test connectivity to TFTP)

> Ping host                 (prompt is confusing. Enter IP address only)
Admin -> 192.168.1.20
Host 192.168.1.20 (192.168.1.20) is alive. The round trip time is 3.96 ms

> Ping host
Admin ->                    (press ENTER to return to Admin menu)

1) Software Update
2) System Reboot
3) Ping
4) Access Rights
5) File Management
6) Certificate Management
7) Back

Admin -> 1                                        (select Software Update)
```

The first prompt asks for the upgrade file name. The file name for the current OS, or for the OS you tried to get on a previous attempt that failed, is displayed in square brackets. Type the new name and press ENTER. The second prompt is asking for the TFTP server address. Type the address and press ENTER. Finally, a prompt will allow modifying the two initial entries, continuing to file transfer, or exiting from the process.

```
Name of the file for main code upgrade? [old.bin] vpn3002-3.6.7.C-k9.bin
IP address of the host where the file resides? [10.1.0.1] 192.168.1.20

(M)odify any of the above (C)ontinue or (E)xit? [M] c
Erasing flash...This can take several seconds to complete!

Starting the TFTP download...
Loading.....................................
Verifying.....................................complete
SUCCESS: New code image will become active on next reboot!

Reboot now? (Y/N) [Y]                             (press ENTER to confirm)
Reboot scheduled immediately...
64 03/29/2003 21:44:38.110 SEV=1 REBOOT/1 RPT=1
Reboot scheduled immediately.

Done

Login: admin
Password:
```

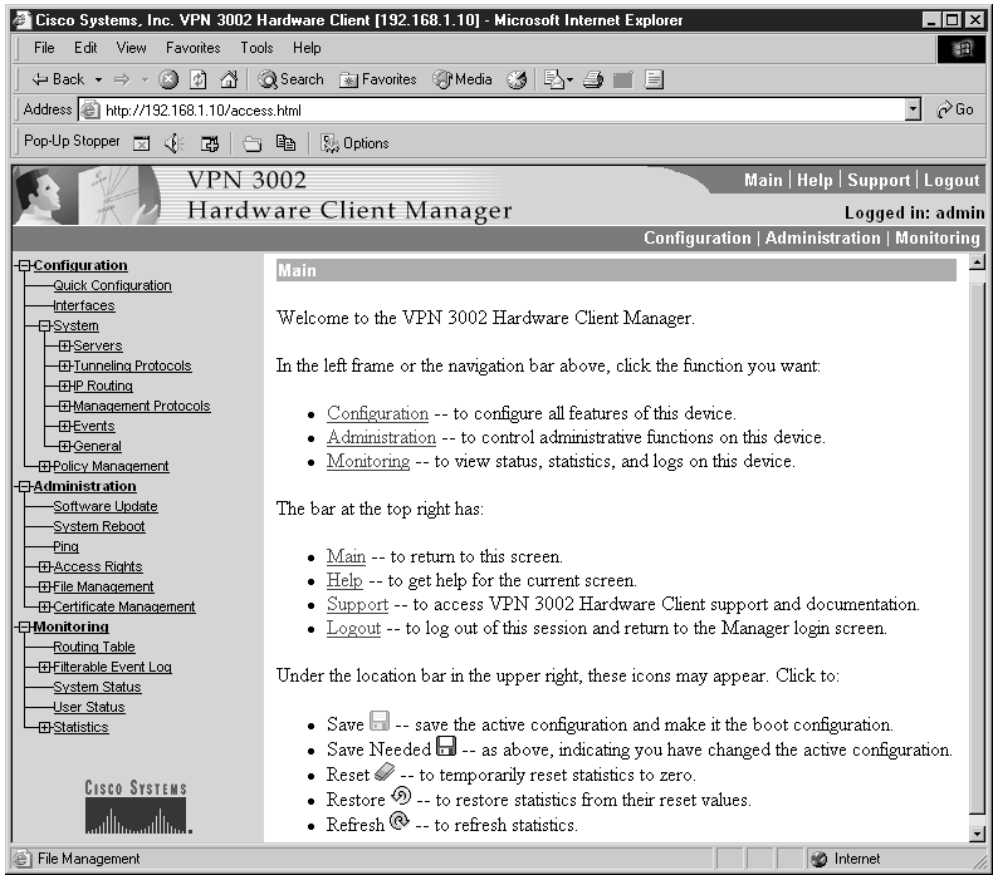

Figure 15-8 Client Manager menu tree expanded on the left side

Figure 15-9 shows the Update process using the Web interface. Expanding Administration in the left-panel tree and choosing Software Update brings up this screen. Note, a Ping option should also be used first to verify connectivity to the TFTP server.

NOTE A Ping option exists that should be used first to verify connectivity to the TFTP server. The screen is straightforward: browse for the OS file on your local computer or a server you have access to, and then click the Upload button.

Quick Configuration

Quick Configuration (Configuration | Quick Configuration) starts a series of screens to provide basic connectivity for the VPN 3002. See the section "VPN Hardware Configuration"

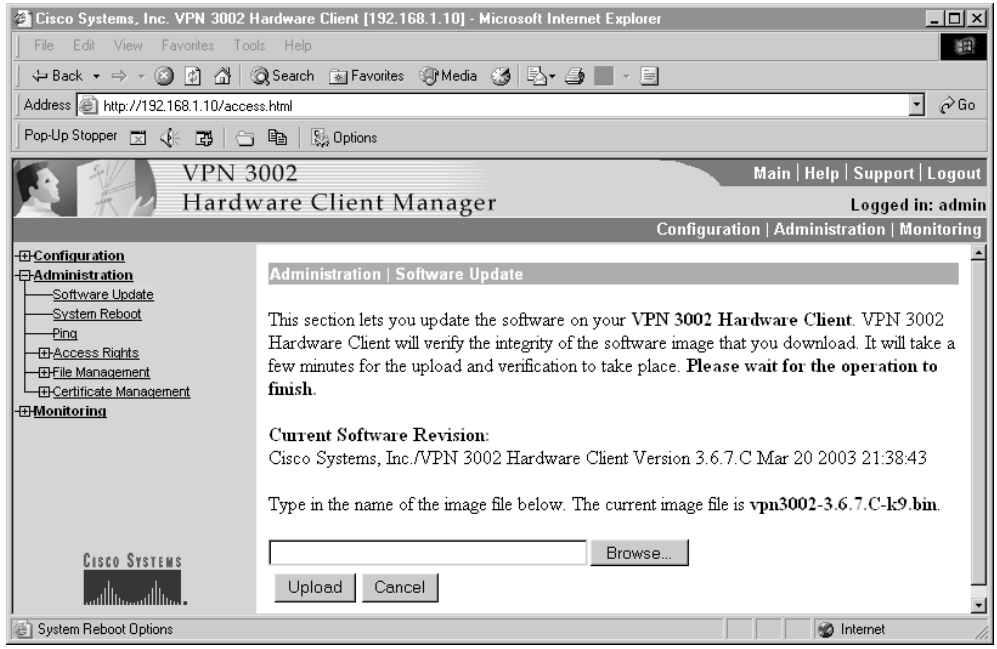

Figure 15-9 VPN 3002 Update process using the web interface

for an example of these features. The Quick Configuration on the web interface consists of the following ten steps (the CLI steps are slightly different). When appropriate, the current settings will appear in square brackets as the default value.

1. Set the system time, date, and time zone.
2. Configure the Private Interface Ethernet interface. To use Network Extension mode, you must configure an IP address other than the default.
3. Optionally upload an already existing configuration file.
4. Configure the Public Interface Ethernet interface to a public network.
5. Specify a method for assigning IP addresses.
6. Configure the IPSec tunneling protocol with group and user names, and passwords and encryption options.
7. Set the VPN 3002 to use either PAT or Network Extension mode.
8. Configure DNS.

9. Configure static routes.
10. Change the admin password for security.

System Status

Use the Monitoring | System Status menu to check the status of several software and hardware variables. The resulting data is a snapshot of the device feature at the time the screen is displayed. From this screen, you can display the status of the IPSec tunnel security associations (SAs) and tunnel duration. On the web interface, the device front and rear panels are displayed with embedded links that display port statistics. Figure 15-10 shows the web-based output for displaying system status.

The following is the CLI output from choosing Monitoring | System Status (3.3) from the menus:

```
System Status
-------------
VPN Concentrator Type: 3002
Serial Number: CAM02223438
Bootcode Rev:
  Cisco Systems, Inc./VPN 3002 Hardware Client Version 3.0.Rel Feb 26
  2001 10:39:17
Software Rev:
  Cisco Systems, Inc./VPN 3002 Hardware Client Version 3.6.7.C Mar 20
  2003 21:38:43
Up For 3:29:12
Up Since 03/30/2003 19:41:58
RAM Size: 16 MB
No Tunnel Established - Public Interface not configured.

1) Refresh System Status
2) Reset System Status
3) Restore System Status
4) Connect Now
5) Disconnect Now
6) View Memory Status
7) Back

Status ->
```

PPPoE Support

Point-to-Point Protocol over Ethernet (PPPoE) incorporates two widely used and understood standards: PPP and Ethernet. The *PPPoE* specification connects hosts on an Ethernet to the Internet through a common broadband medium, such as DSL line, cable modem, or a wireless device. With PPPoE, the principles of Ethernet supporting multiple users in a LAN are combined with the principles of PPP, which uses serial

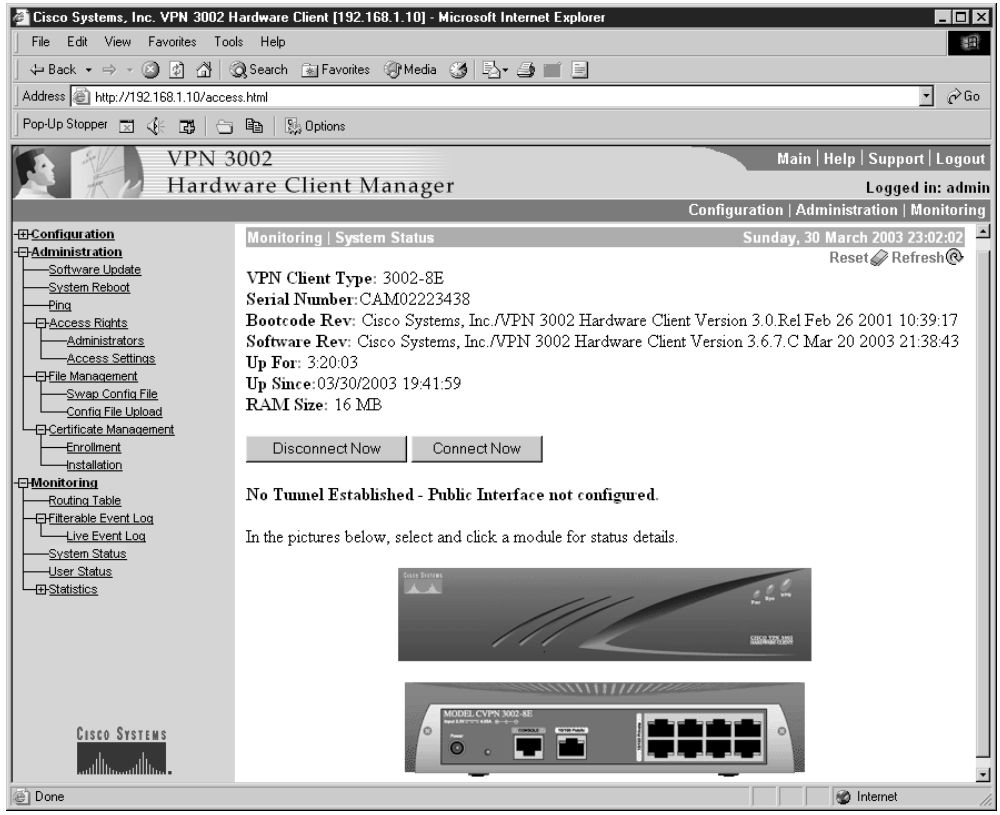

Figure 15-10 Web interface displaying system status

connections. The VPN 3002 supports PPPoE Client mode on the public interface to access these networks. Users need only to authenticate to the PPPoE server the first time and, for all subsequent attempts, VPN 3002 will authenticate for the user.

Figure 15-11 shows the public interface screen (Configuration | Interface | Public). The feature is supported on the CLI by following the same menu options. To configure PPPoE, you must furnish the following information, which is generally provided by the ISP.

- A valid PPPoE user name
- The PPPoE password for the user name entered previously
- The PPPoE password again to verify it

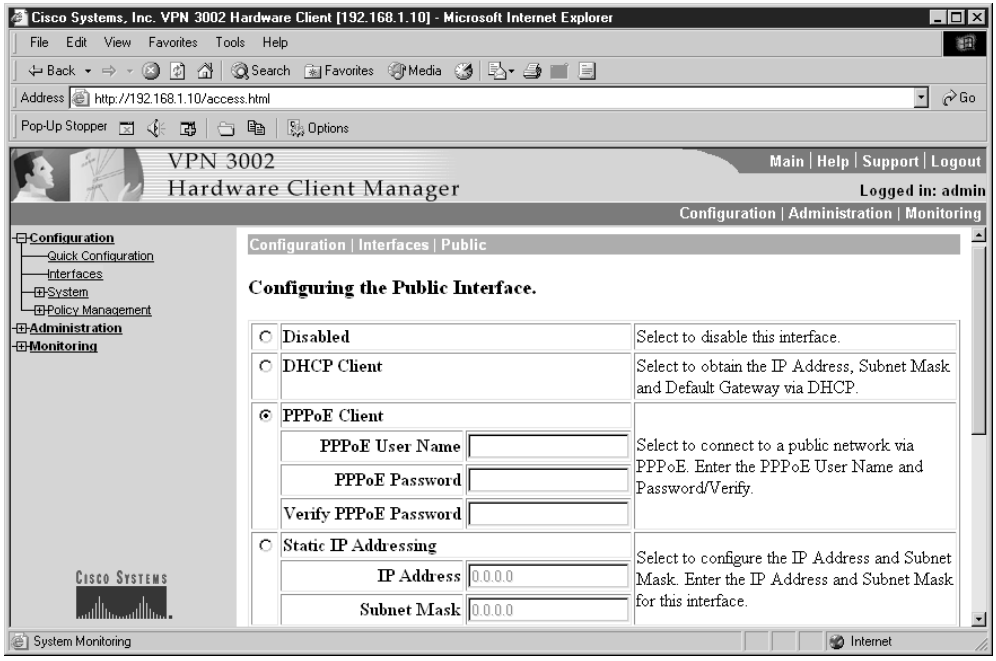

Figure 15-11 PPPoE configuration on the public interface screen

Basic Configuration for the VPN 3002

The Quick configuration ten-step process in either CLI or web-based Client Manager can be used to supply the minimal parameters needed to make the VPN 3002 operational. The Client Manager is used in this section, but it shouldn't be any trick to follow along in the CLI.

The following discussion assumes a successful login and choosing Configuration | Quick Configuration in the left-side panel. The actual configuration is based on Figure 15-12, showing a small branch location connecting through a VPN 3002 client to a VPN Concentrator at the main office.

The scenario assumes the main office has reserved the networks 192.168.0.0 to 192.168.127.0 for its internal use. The other private class C addresses have been assigned as needed to the company's branch locations. The figure shows a small branch location assigned the 192.168.145.0 network.

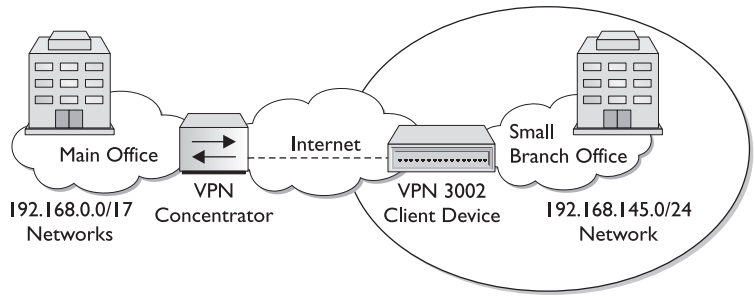

Figure 15-12 VPN 3002 configuration scenario

Set the System Time, Date, and Time Zone

The Client Manager window displays the Configuration | Quick | Time and Date screen.

Figure 15-13 shows the entry screen used to set the time and date on this device. The choices are self-explanatory. Notice all ten steps are listed at the top of the screen, allowing a person to jump to any feature. Click the Continue button to advance to the next screen.

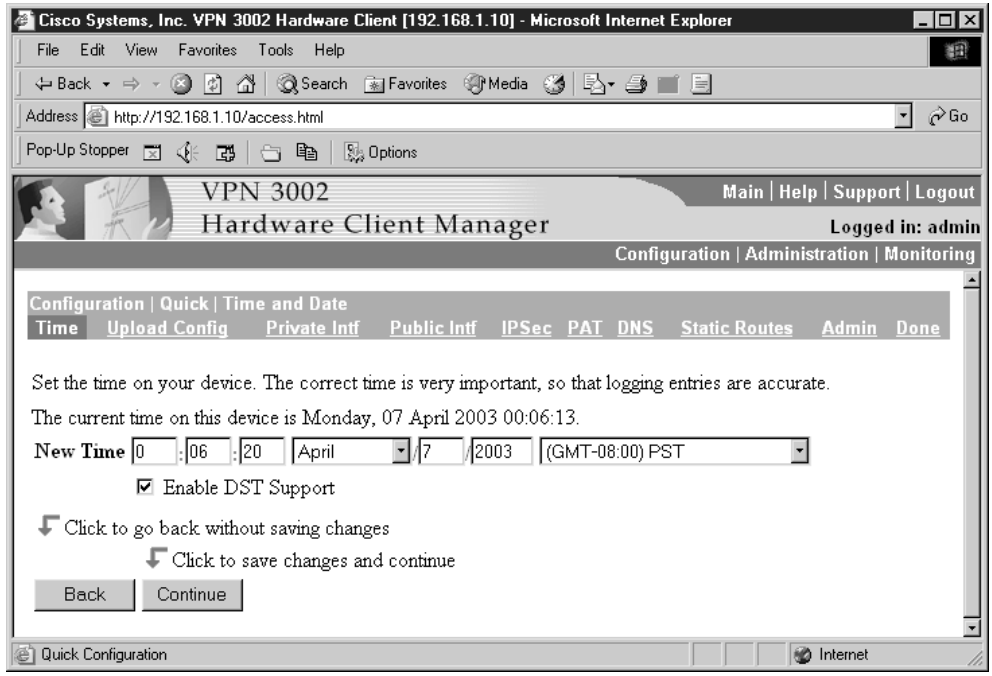

Figure 15-13 Screen to set date, time, and time zone for the device

Optional—Upload an Existing Configuration File

The Client Manager window displays the Configuration | Quick | Upload Config screen.

Figure 15-14 shows the option screen used to use HTTP or HTTPS to transfer (upload) configuration files from a host to the VPN 3002 flash memory. This could be a time-saver if you need to restore a lost or damaged configuration.

Click No to continue to the next section or click Yes to upload an existing configuration file.

Configure the Private Interface

The Client Manager window displays the Configuration | Quick | Private Interface screen.

Figure 15-15 shows the screen used to configure the VPN 3002 private interface. This is the protected LAN interface of the network. The top portion of the screen displays the current configuration settings. The following is a possible example:

IP Address 192.168.1.10/ 255.255.255.0
DHCP Server Enabled (192.168.1.21–192.168.1.254)

The first question determines whether to reconfigure the IP address for the private interface. The Yes/No choices are self-explanatory.

The second question deals with using DHCP to define the address for the private interface. The first choice ultimately brings up a screen to configure the DHCP server

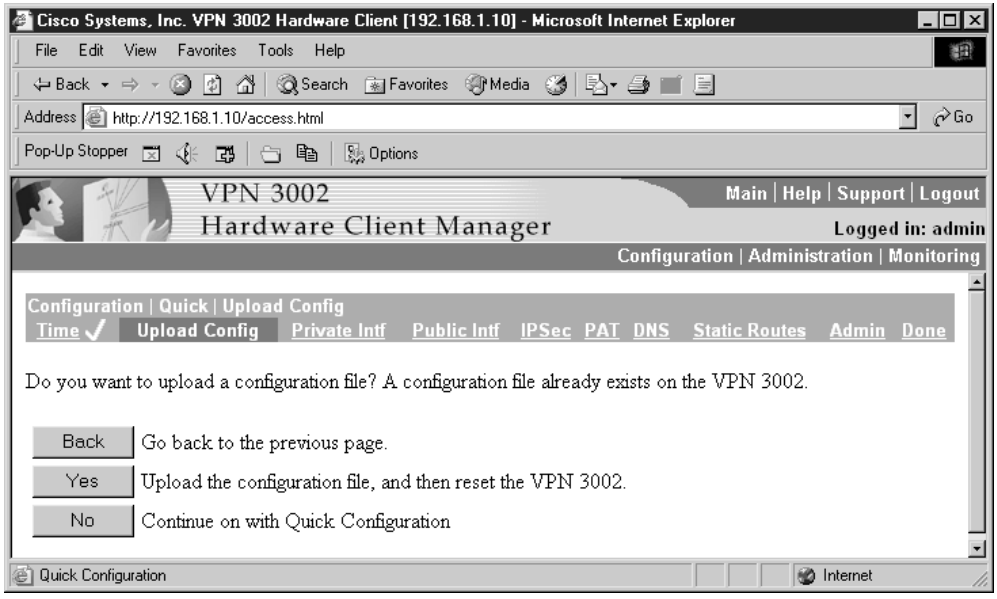

Figure 15-14 Choice to upload an existing configuration file

PART III

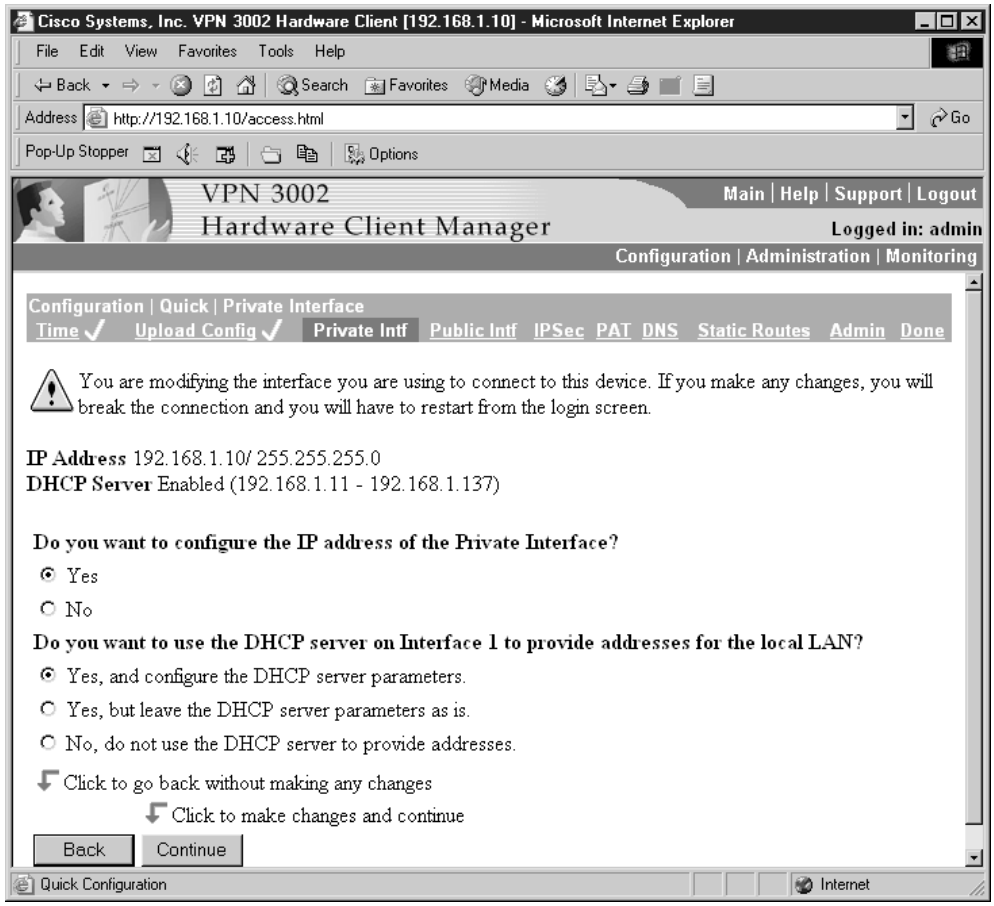

Figure 15-15 Configure the private interface

parameters, as shown in Figure 15-16. The second choice would be used while reviewing the settings. The third choice would be used to configure the interface address manually. Click the Continue button to implement the advance of the choice to the appropriate next screen.

For the VPN 3002 to operate in Network Extension mode, you must change the private interface IP address from the default setting 192.168.10.1.

Configuring the DHCP Server

The Client Manager Configuration | Quick | Private Interface | DHCP Server screen is used to enable and configure the VPN 3002 private interface to serve as a DHCP server for the private network hosts. This allows IP hosts on the LAN to obtain IP addresses automatically from a limited pool of addresses for a fixed length of time or for a lease

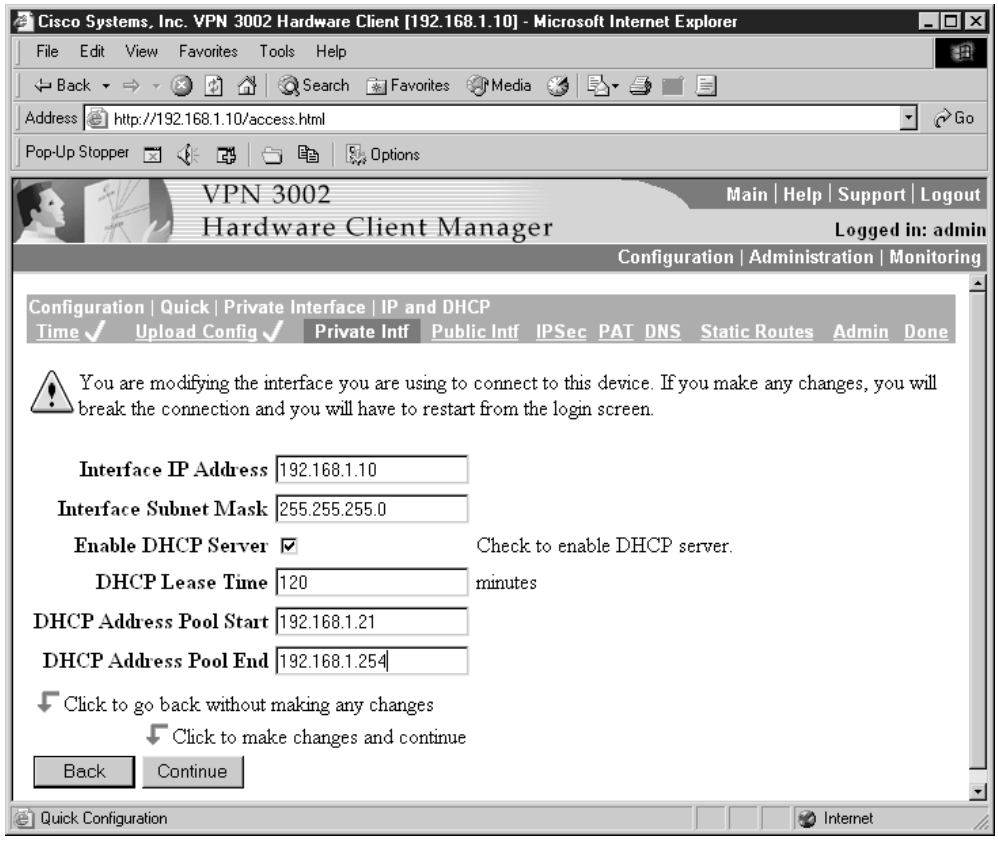

Figure 15-16 Configuring the DHCP server

period. DHCP simplifies host configuration by allowing the network settings to be learned from the DHCP server rather than statically configured.

Check the Enabled box to enable DHCP services for the private interface.

The DHCP Lease Timeout field can be set between 5 and 500,000 minutes. The default is 120 minutes. After the lease time expires, the lease can be renewed or returned to the address pool. This Lease Timeout period applies only when the tunnel to the VPN Concentrator is established. When the tunnel isn't established, the Lease Timeout period is five minutes.

The next two fields are used to define the starting and ending address of the DHCP. The example reflects the first 20 addresses were excluded from the pool, so they might be permanently assigned to shared resources, such as printers and servers. The default pool is 127 IP addresses and the start of the range is the next IP address after that of the current private interface.

If changes were made, the Manager displays the Configuration | Quick |

Private Interface | DHCP server address pool screen used confirms that the DHCP server address pool range was entered.

Configure the Public Interface

The Client Manager, displaying the Configuration | Quick | Public Interface screen, is shown in Figure 15-17.

This is the interface used to connect to an ISP and to the central network. The public interface can obtain an IP address in one of three ways:

- DHCP client
- PPPoE
- Static addressing

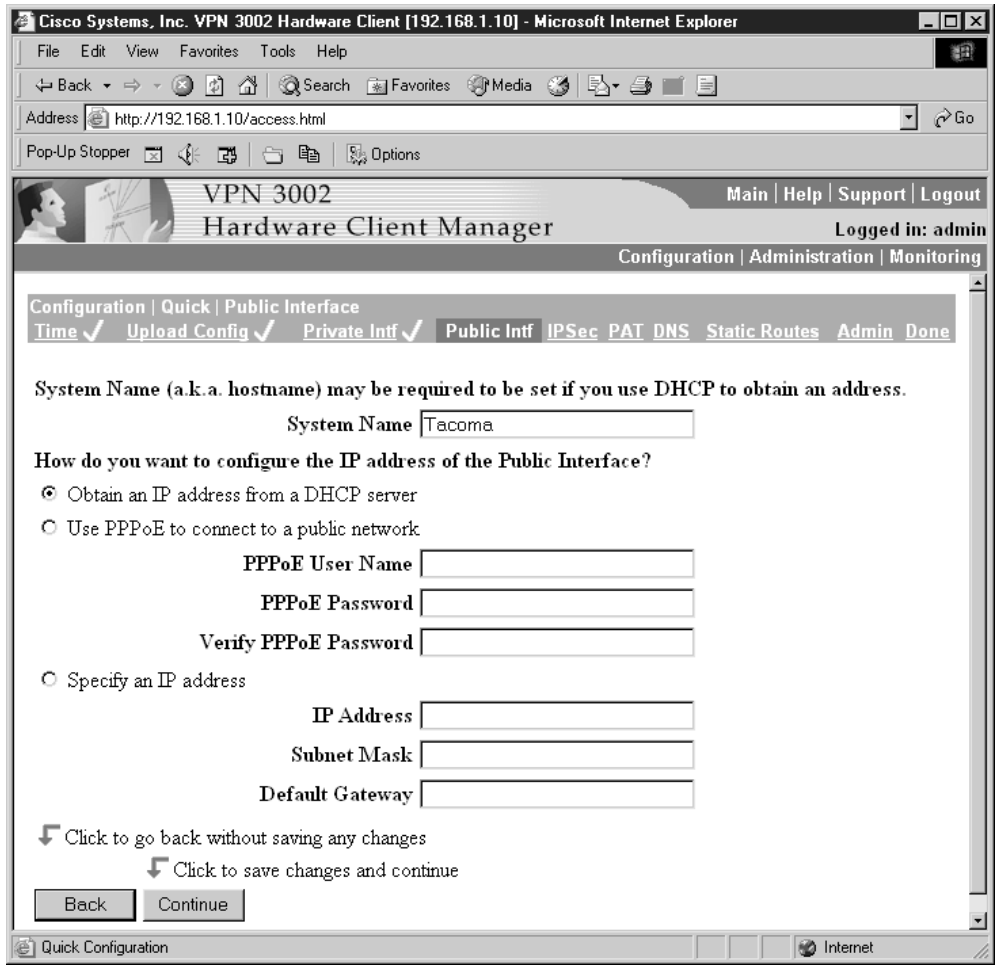

Figure 15-17 Screen to configure the public interface

Once the choice is made, PPPoE and static addressing need some address information. The PPPoE information would generally be provided by the ISP.

The system name, also known as a host name, is optional unless DHCP is chosen to obtain an IP address and the ISP requires a host name. As a DHCP client, the upstream DHCP server assigns the public interface IP address, subnet mask, and default gateway.

Configure the IPSec

The Client Manager displays the Configuration | Quick | IPSec screen.

This screen lets you configure the IPSec parameters, enabling the VPN 3002 to connect to the VPN Concentrator or to other IPSec security gateways, such as the Cisco PIX firewall or Cisco IOS routers. Figure 15-18 shows the IPSec configuration screen.

The Remote Server field is for the IP address or host name of the VPN Concentrator to which this VPN 3002 hardware client connects. If a host name is used, a DNS server must be available to resolve the name.

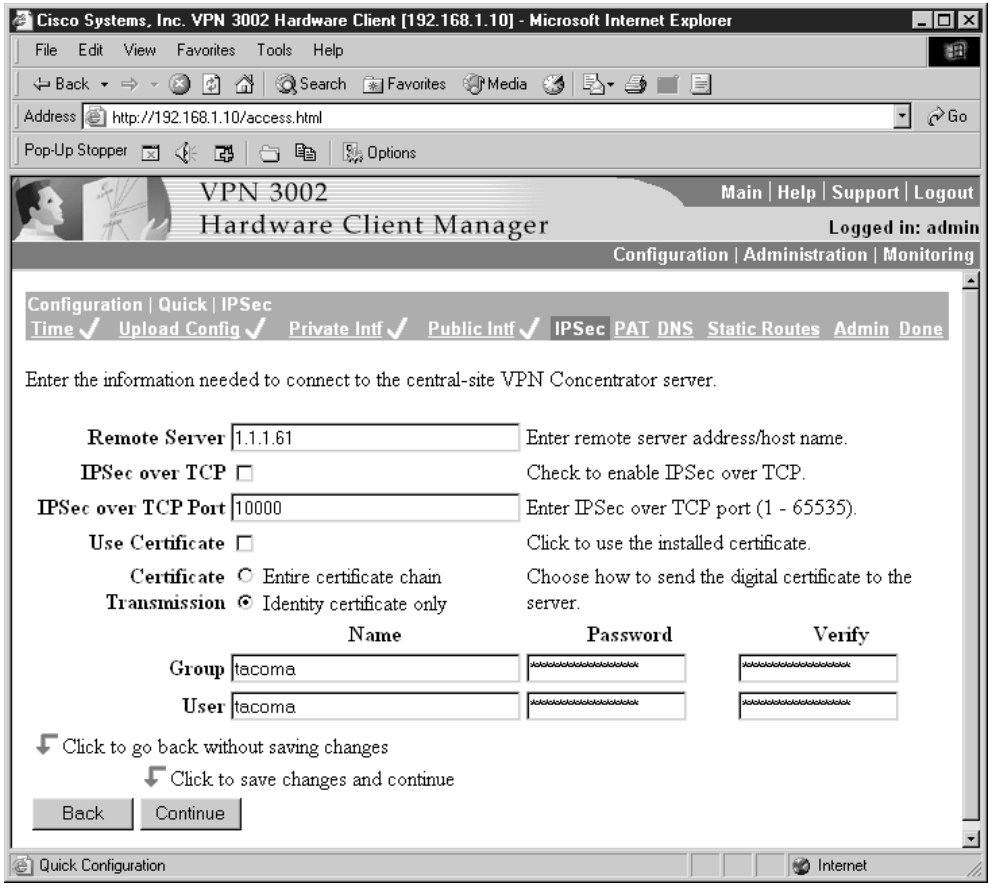

Figure 15-18 IPSec configuration screen

As mentioned previously, NAT-T is the default, but you can check the IPSec over TCP box to use TCP. The TCP feature must also be enabled on the VPN Concentrator to which this VPN 3002 connects.

Specify the IPSec over TCP port number; only one port number can be specified. The VPN 3002 port must also be configured on the VPN Concentrator to which this VPN 3002 connects.

The *Use Certificate* box specifies digital certificates for authentication. With digital certificates, you needn't enter a group name and group password.

Digital Certificates

You have two Select a Certificate Transmission options:

- **Entire certificate chain**—to send the peer the identity certificate and all issuing certificates, including the root certificate and any subordinate CA certificates.
- **Identity certificate only**—to send the peer only the identity certificate.

Preshared Keys

The following information has to be consistent with that configured for this VPN 3002 on the central-site VPN Concentrator.

- Group Name field—unique name for this group (up to 32 characters, case-sensitive).
- Group Password field—unique password for this group (4 to 32 characters, case-sensitive). The field displays only asterisks.
- Group Verify field—reenter the group password.
- User Name field—unique name for this user in the group (up to 32 characters, case-sensitive).
- User Password field—unique password for this user (4 to 32 characters, case-sensitive). The field displays only asterisks.
- User Verify field—reenter the user password.

Choose Client (PAT) Mode or Network Extension Mode

The Client Manager displays the Configuration | Quick | PAT screen.

The next screen is used to specify either Client (PAT) mode or Network Extension mode. The default Yes selects Client mode; No selects Network Extension mode. Figure 15-19 shows the selection screen.

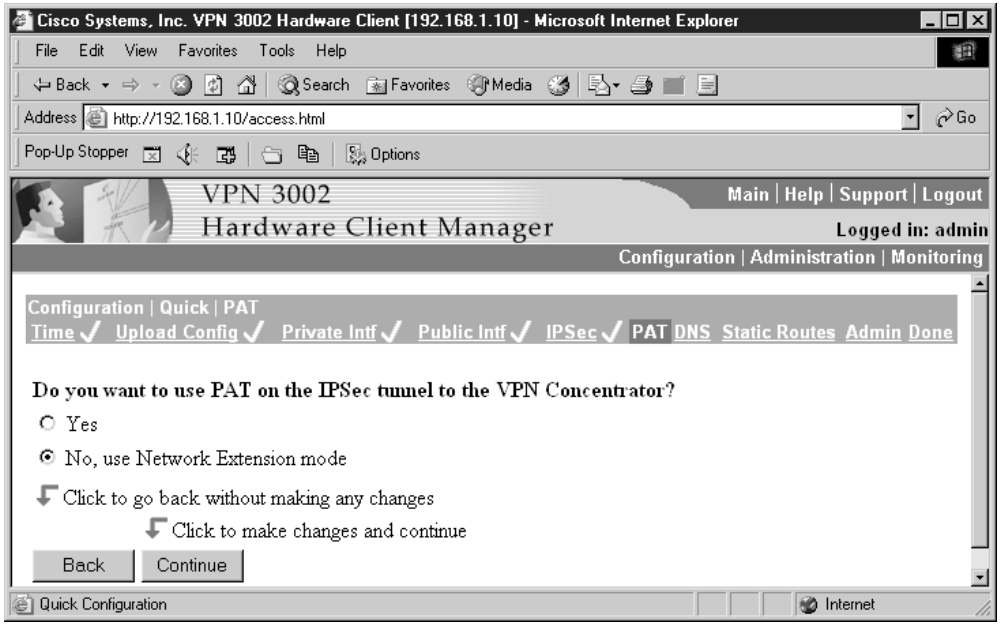

Figure 15-19 Client mode or Network Extension mode choice

Configure DNS

The Client Manager displays the Configuration | Quick | DNS screen.

As shown in Figure 15-20, this screen is used to specify a DNS server for the local ISP, so Internet host names can be used instead of IP addresses for servers when configuring and managing the VPN 3002. While host names are easier to remember, using IP addresses avoids problems that might occur with the DNS server offline or congested.

If a host name was used to identify the central-site VPN Concentrator on the IPSec configuration screen, a DNS server must be configured on the VPN 3002.

If used, the IP address of the local DNS server is entered in the DNS Server field. The local ISP domain name is entered in the Domain field.

Configure Static Routes

The Client Manager displays the Configuration | Quick | Static Routes screen.

The Static Routes list shown in Figure 15-21 displays any existing static IP routes that were configured. The format is destination network address/subnet mask -> outbound destination. Use this screen to add or delete static routes for IP routing.

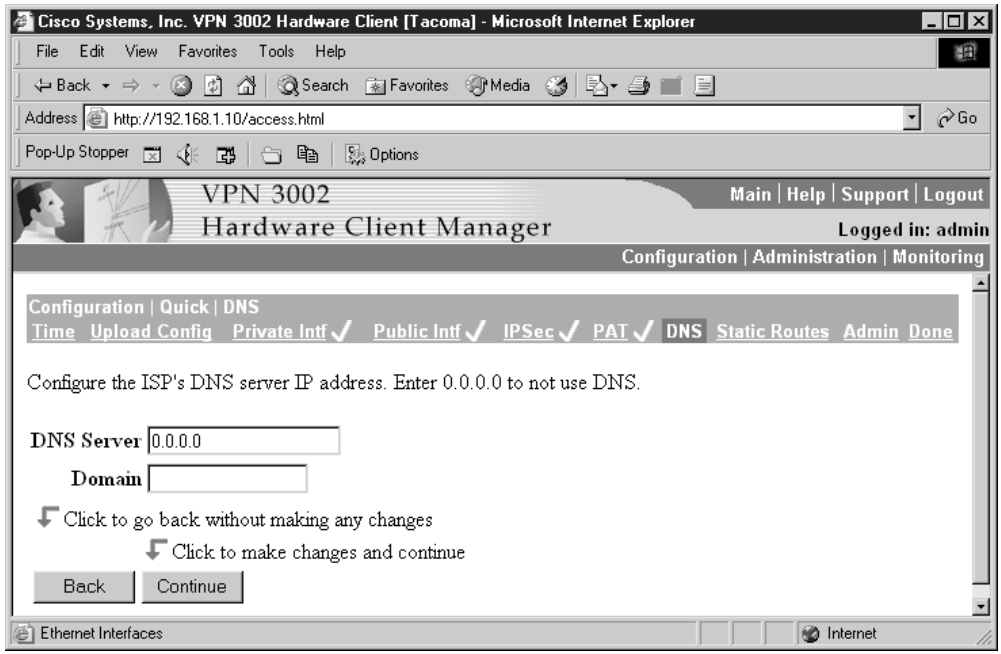

Figure 15-20 Define ISP DNS server and domain name (optional)

Clicking the Add button displays the Configuration | Quick | Static Routes | Add screen, as shown in Figure 15-22. This screen lets you add a new static route to the IP routing table. The options are pretty straightforward. The Subnet Mask automatically defaults to a standard classful subnet mask, but it can be changed as needed.

The Metric field allows assigning a cost for the route. The range is 1 to 16, where 1 is the lowest cost. The device always tries to use the least costly route. This makes creating floating static routes possible, where two routes to the same network can be given different metrics to reflect a preference.

The last choice is between using a next-hop address or the local VPN 3002 interface. For the Interface option, the drop-down menu button can be used to select a configured VPN 3002 interface as the outbound destination.

Change the Admin Password

The Client Manager displays the Configuration | Quick | Admin Password screen.

The screen is used to change the password for the administrator account (admin). The default password is also admin. Obviously, this isn't secure for the most powerful account on the device. Changing this password makes sense to improve device security.

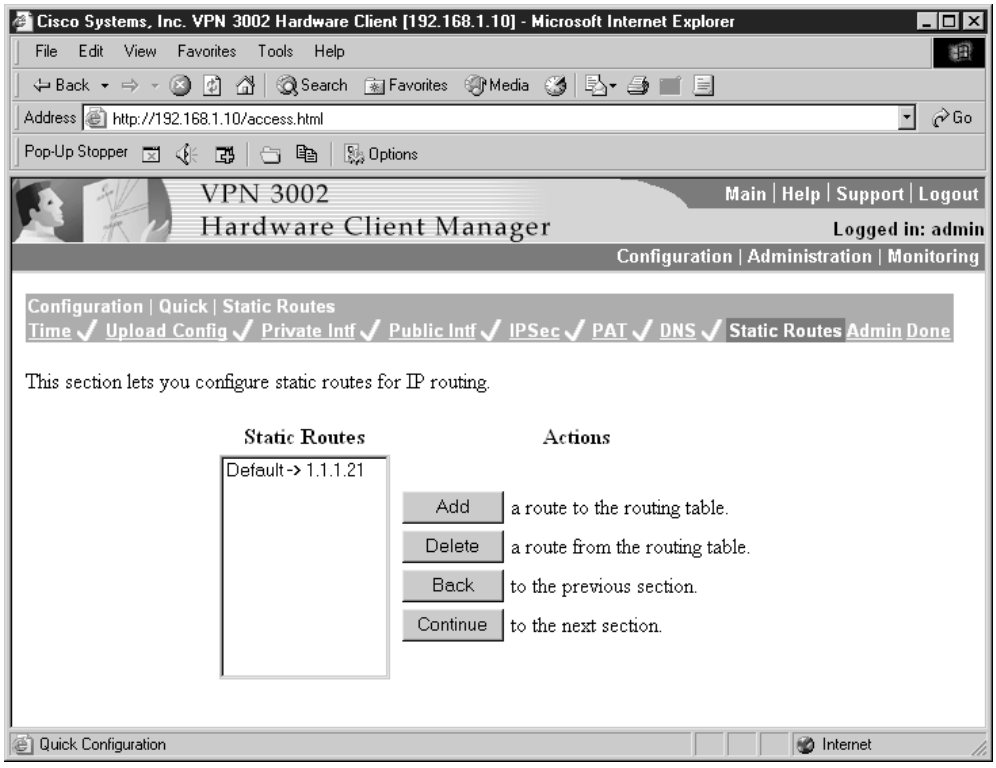

Figure 15-21 Existing static routes

When the password is set, the Quick configuration is done and a message screen much like the opening screen confirms that.

The Quick configuration can be used again to make changes or the Configuration menu can be used to change specific features or add options.

Modifying Options

The Quick configuration is used to configure the minimum requirements for connecting to a VPN Concentrator. Modifying or adding options later to a VPN Concentrator is easy. For example, when the DHCP server was configured on the private interface, only minimal features were defined. Once the initial configuration is in place, use the Configuration menu to set additional features. Figure 15-23 shows the DHCP Options screen and the left-side panel shows how to get there. This is where additional servers could be defined.

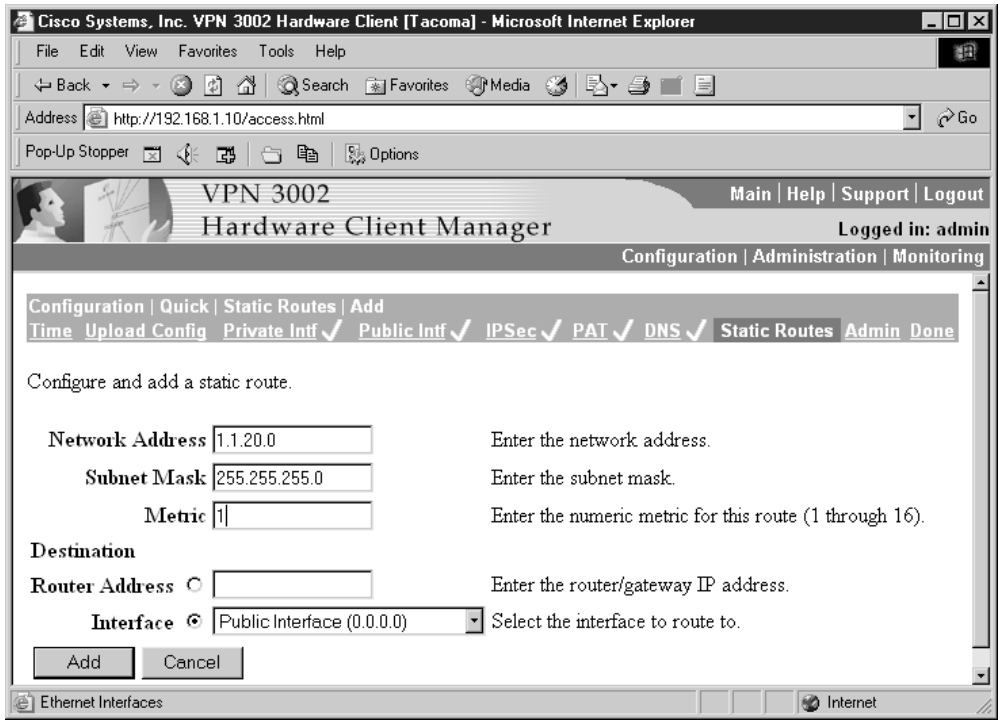

Figure 15-22 Adding a static route to the VPN 3002

Other VPN 3002 Software Features

The VPN 3002 software supports the following features:

Interactive Hardware Client Authentication

Interactive hardware client authentication, sometimes called *interactive unit authentication*, prevents VPN 3002 private LAN users from accessing the central site until the VPN 3002 unit authenticates. In this scenario, the VPN 3002 doesn't use a saved user name and password for authentication. Instead, a valid user name and a password for the 3002 must be manually entered each time.

The VPN 3002 sends the user name and the password to the VPN Concentrator when it initiates a tunnel session. The VPN Concentrator can authenticate the connection using either an internal or an external server. The tunnel is only established if the user name/password combination is valid.

Interactive hardware client authentication is configured on the VPN Concentrator, which then pushes the policy down to the VPN 3002 at the next connection.

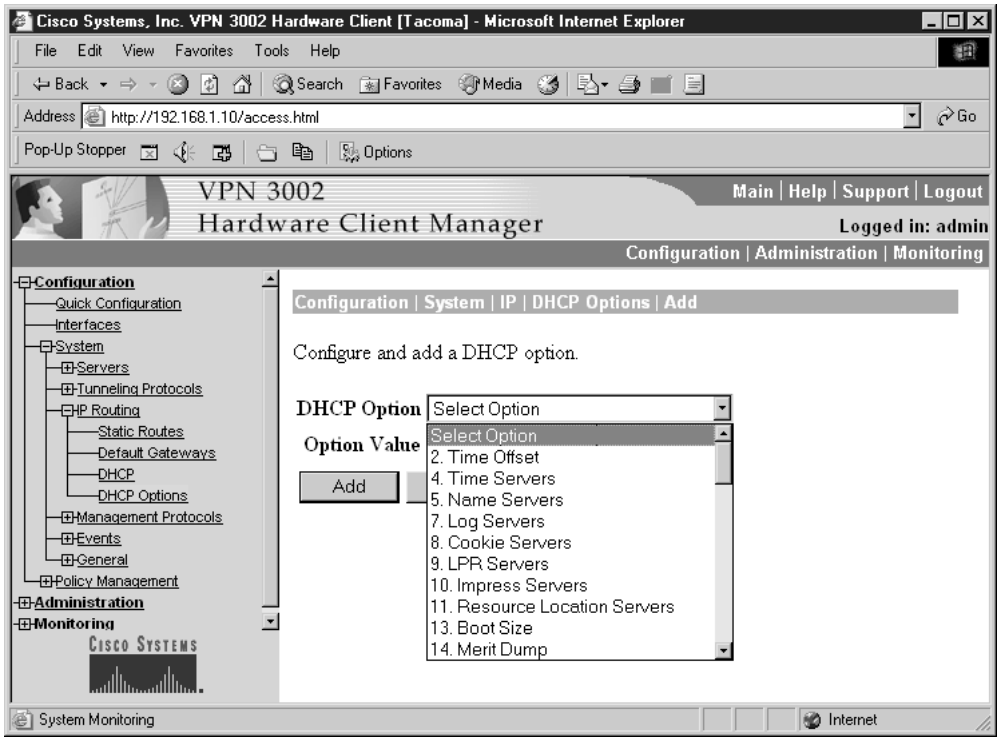

Figure 15-23 DHCP Options configuration screen

Configuring Interactive Unit Authentication

The Hardware Client parameters tab on the VPN 3000 Concentrator Series Manager is used to configure several features for the VPN 3002 and its users in the base group. The feature will be "pushed" down to the client devices the next time the VPN 3002 establishes a session. The menu selection is Configuration | User Management | Base Group, HW Client parameters tab.

Check the Require Interactive Hardware Client Authentication check box, as shown in Figure 15-24.

Individual User Authentication

Individual user authentication protects the central site from access by unauthorized individuals on the VPN 3002 private network. It accomplishes this by requiring each user to open a web-browser session and manually enter a valid user name and password combination to gain access to the network behind the VPN Concentrator, regardless of

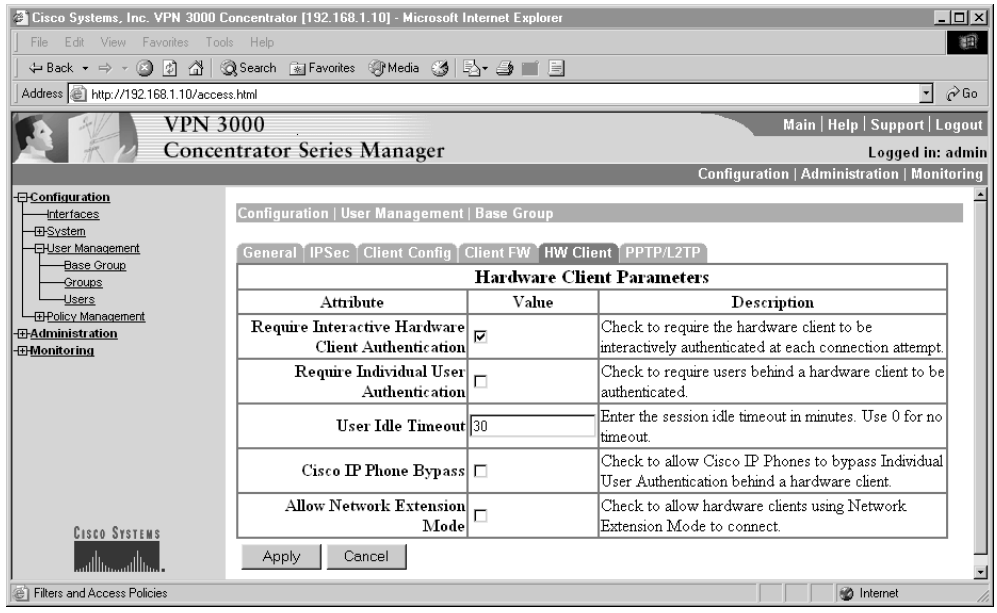

Figure 15-24 Configure Interactive Hardware Client Authentication

whether a VPN tunnel already exists. A successful login results in the browser displaying the appropriate default home page.

This feature can only be used with a browser, not the command-line interface. Attempts to access non-web-based resources, such as e-mail, on the network behind the VPN Concentrator will fail until a successful browser authentication occurs.

To simplify the process and make it as transparent as possible to the end users, this technology automatically intercepts any users attempting to traverse the VPN tunnel and redirects them to a browser page to authenticate. Users needn't initiate or remember to initiate the security authentication because it's done automatically. If users are only attempting to access the Internet via Split Tunneling, they aren't prompted to authenticate. Each user can maintain a maximum of four simultaneous login sessions.

Individual user authentication is configured on the VPN Concentrator, which then pushes the policy down to the VPN 3002 at the next connection.

Configuring Individual User Authentication

This feature is enabled on the same screen as the last feature. The menu selection is Configuration | User Management | Base Group, HW Client parameters tab, as shown in Figure 15-24 in the previous section.

Check the Require Individual User Authentication check box. This feature can be used separately or in conjunction with the Interactive Hardware Client Authentication.

LEAP Bypass

IEEE 802.1X is a standard for authentication on wired and wireless networks providing wireless LANs with strong mutual authentication between clients and authentication servers. 802.1*X* provides dynamic per-user, per-session wireless encryption privacy (WEP) keys, thereby removing administrative overhead and security concerns related to static WEP keys.

Lightweight Extensible Authentication Protocol (LEAP) is Cisco Systems 802.1*X* wireless authentication technology that implements mutual authentication between a wireless client and a RADIUS server. The authentication credentials, including a password, are always encrypted before they're transmitted over the wireless medium.

LEAP Bypass allows LEAP packets from devices behind a VPN 3002 to travel across a VPN tunnel before individual user authentication. This allows wireless workstations using access point devices to establish LEAP authentication, and then authenticate again using individual user authentication, if enabled.

Without this technology, LEAP users behind a VPN 3002 are caught in a Catch-22. They can't authenticate on the wireless network because they can't access the VPN tunnel to get to the RADIUS. They can't access the VPN tunnel because they haven't authenticated on the wireless network.

The VPN Concentrator administrator enables LEAP Bypass on a per group basis at the central site, using a check box on the HW Client tab on the Group configuration page. The LEAP packets travel over the tunnel to a RADIUS server via ports 1645 or 1812.

LEAP Bypass functions properly if the following conditions are met.

- Interactive unit authentication must be disabled, otherwise, a non-LEAP (wired) device needs to authenticate the VPN 3002 before LEAP devices can connect using the tunnel.
- Individual user authentication is enabled, otherwise, LEAP Bypass isn't needed.
- The VPN 3002 device can be in either Client mode or Network Extension mode.
- Wireless Access points must be Cisco Aironet Access Points.
- The Cisco Aironet Access Point must be running Cisco Discovery Protocol (CDP).
- The wireless NICs for the PCs can be from other manufacturers.

While the LEAP and LEAP Bypass technologies are sound, some security risk always exists in allowing any unauthenticated traffic to traverse the secure tunnel.

IPSec Backup Servers

The IPSec backup servers feature provides alternatives for the VPN 3002 hardware client to connect to the central site when its primary VPN Concentrator is unavailable. Backup servers can either be configured individually on the VPN 3002 device or on a per-group basis on the central-site VPN Concentrator. When configured on the central-site VPN Concentrator, the Concentrator pushes the backup server policy to all VPN 3002 hardware clients in the group.

The following characteristics apply to the IPSec backup server feature:

- Each VPN 3002 must connect to the primary VPN Concentrator at least once to download a backup server list. A backup server list can't be downloaded from a backup server.
- If the primary VPN Concentrator is unavailable to download the backup server list and the VPN 3002 has a previously configured backup server list, it can continue to connect to the servers on that list.
- If the VPN 3002 has tried all designated backup servers on the list and can't connect, it doesn't automatically retry. The following trigger a new round of attempts:
  - In Network Extension mode, the VPN 3002 attempts a new connection after four seconds.
  - In Client mode, the VPN 3002 attempts a new connection when the user clicks the Connect Now button on the Monitoring | System Status screen or when data passes from the VPN 3002 to the VPN Concentrator.
- Any changes to the configuration of the backup server's list during an active VPN 3002 session won't take effect until the next time the VPN 3002 connects to its primary VPN Concentrator.
- The VPN Concentrator backup servers needn't be aware of each other.

The group name, user name, and any passwords configured for the VPN 3002 must be identical for the primary VPN Concentrator and all backup servers. Also, if interactive hardware client authentication and/or individual user authentication are configured for the VPN 3002 on the primary VPN Concentrator, they must be configured on backup servers as well.

Configure IPSec Backup Servers—VPN 3002 Client

You can configure the backup server feature from the primary VPN Concentrator or the VPN 3002. Use the Configuration | System | Tunneling Protocols | IPSec screen to configure backup servers directly on the VPN 3002. From this screen, shown in Figure 15-25, you can configure up to ten servers, ranging from the highest priority on the top to the lowest priority on the bottom. The Backup Easy VPN Servers window is only a small text box allowing direct entry and insertions.

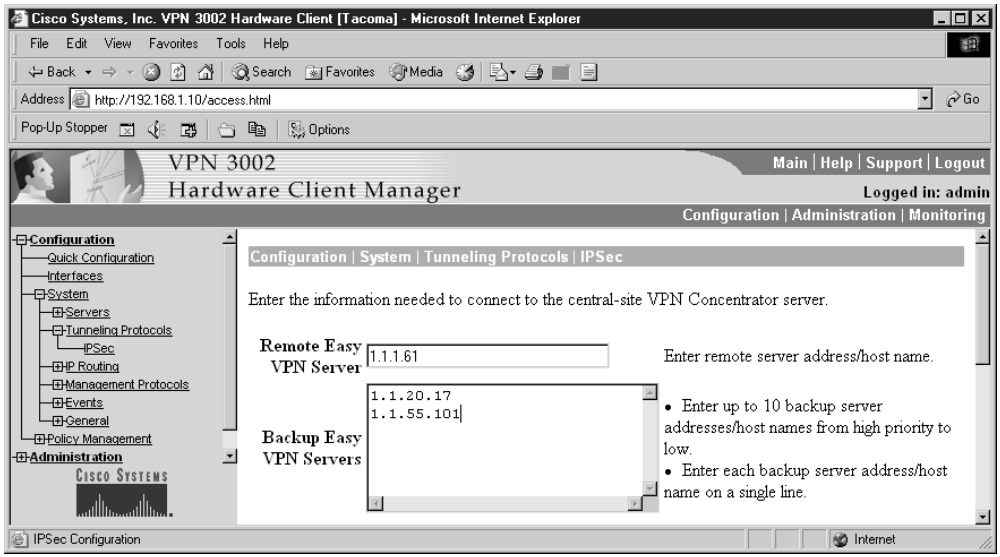

Figure 15-25 Configure backup servers directly on the VPN 3002 Client

Configure IPSec Backup Servers—VPN 3000 Concentrator

To configure backup servers for the VPN 3002 from the VPN Concentrator, use the Configuration | User Management | Base Group, Client Config tab, as shown in Figure 15-26. The backup server's list will apply the next time the 3002 connects to its primary concentrator.

In the IPSec Backup Servers section, use the drop-down list box to select a method to use or disable backup servers. The three choices are as follows:

- Use client configured list
- Disable and clear client configured list
- Use list below

Enter up to ten IPSec backup server addresses/names starting from the highest priority to the lowest. Enter each IPSec backup server address/name on a single line.

IPSec Server Load Balancing

The load balancing feature makes it possible to distribute remote sessions among two or more VPN Concentrators connected on the same network. Load balancing provides efficient use of system resources, while providing increased performance and high availability by directing remote sessions to the least-loaded device.

The load balancing is used only with remote sessions with VPN Concentrators initiated by either the Cisco VPN Client (3.0 or later) or the Cisco VPN 3002 Hardware

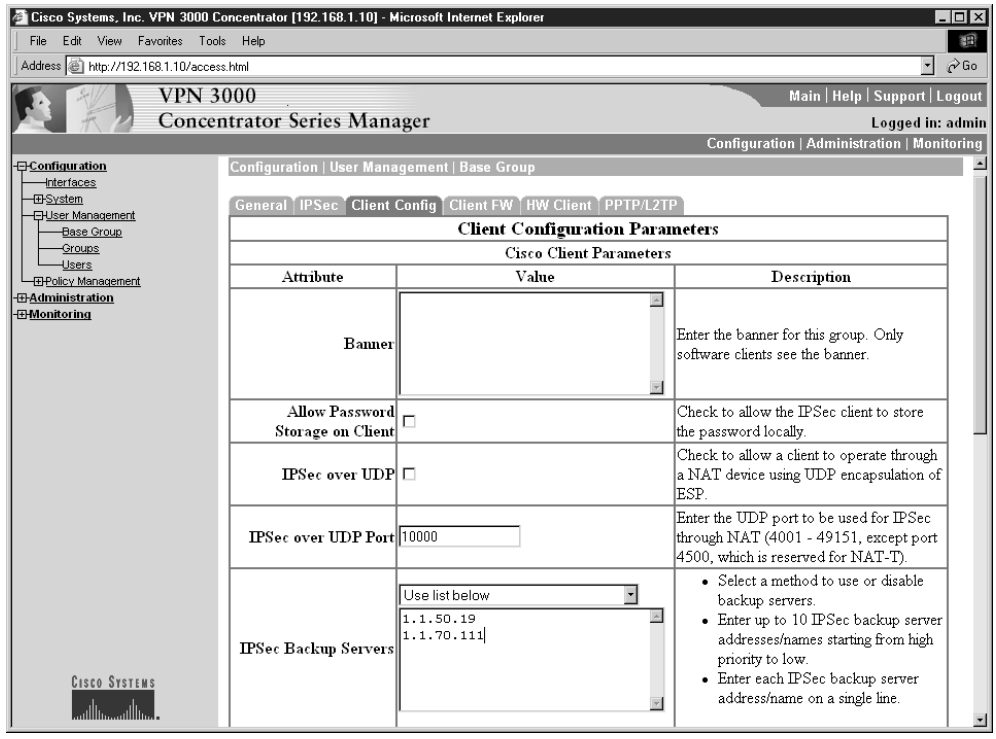

Figure 15-26 Configure backup servers on the VPN Concentrator

Client (3.5 or later). All other VPN clients, including LAN-to-LAN connections, can connect to a VPN Concentrator on which load balancing is enabled, but they can't participate in load balancing.

Load balancing requires no configuration on the VPN Client or VPN 3002.

Configure a Server for Load Balancing

Before configuring load balancing on a VPN Concentrator, you must complete the following two tasks:

- Configure the private and public interfaces.
- Configure the filters for the private and public interfaces to allow the Virtual Cluster Agent (VCA) load balancing protocol.

Configure Interfaces Use the Configuration | Interfaces window to check to see if the public and private interfaces were defined and each has status UP. If either interface is undefined, it must be defined before proceeding.

Virtual Cluster To implement load balancing, you must group together two or more VPN Concentrators logically on the same private LAN-to-LAN network, private subnet, and public subnet into a virtual cluster. The virtual cluster appears to outside clients as a single virtual cluster IP address.

All devices in the virtual cluster are used to distribute session loads. The *virtual cluster master* directs incoming calls to the other devices, referred to as *secondary devices*. By monitoring all devices, the virtual cluster master can distribute the session load based on the activity of each device. The virtual cluster master role isn't assigned to a specific physical device, but can shift among devices, as needed. This flexibility is particularly important if the current virtual cluster master fails. In this case, one of the secondary devices takes over and immediately becomes the new virtual cluster master.

A VPN Client wanting to initiate a session connects to the virtual cluster IP address. The virtual cluster master returns the public IP address of the cluster host with the least load to the client. The process is transparent to the user because the VPN client connects directly to that designated host without any user involvement or messages.

If a cluster machine fails, the terminated sessions reconnect immediately to the virtual cluster IP address where the virtual cluster master repeats the reassignment process. Even if the failed device is the virtual cluster master, one of the secondary cluster devices immediately and automatically takes over as the new virtual session master. Multiple device failures in the cluster should also be resolved, as long as one cluster device remains available.

Configure Filters Use the following VPN Concentrator steps to configure the filters for the private and public interfaces to allow the VCA load balancing protocol.

1. In the Configuration | Interfaces window, select Ethernet1 (Private). The Configuration | Interfaces | Ethernet1 window appears.
2. Choose the General tab.
3. Use the drop-down Filter menu button and select Private (Default).
4. Click on Apply.
5. In the Configuration | Interface window, select Ethernet2 (Public). The Configuration | Interfaces | Ethernet2 window appears.
6. Choose the General tab.
7. Use the drop-down Filter menu button and select Public (Default).
8. Click Apply.
9. Go to the Configuration | Policy Management | Traffic Management | Filters window.
10. Choose Private (Default) from the Filter list.
11. Select Assign Rules to Filter. The Configuration | Policy Management | Traffic Management | Assign Rules to Filter window appears.

12. Verify that VCA In (forward/in) and VCA Out (forward/out) are in the Current Rules in Filter list. If necessary, add them to the list.
13. Click on Done.
14. In the Configuration | Policy Management | Traffic Management | Filters window, choose Public.
15. Choose Assign Rules to Filter. The Configuration | Policy Management | Traffic Management | Assign Rules to Filter window appears.
16. Verify VCA In (forward/in) and VCA Out (forward/out) are in the Current Rules in Filter list. If necessary, add them to the list.
17. Click Done.
18. Click the Save Needed icon in the upper-right corner to save the changes.

Configure Load Balancing Use the Configuration | System | Load Balancing screen to enable load balancing on the VPN Concentrator, as shown in Figure 15-27. The process takes two steps:

1. Configure the cluster—Define the common virtual cluster IP address, UDP port (if necessary), and IPSec shared secret value for every device in the cluster.
2. Configure the device —Enable load balancing on the device, and then define the device-specific properties. These values can vary, based on device type and option features.

VPN Virtual Cluster IP Address	A single IP address identifying the virtual cluster. This address must be within the address range shared by all VPN Concentrators in the cluster.
VPN Virtual Cluster UDP Port	A UDP destination port number to use for load balancing if another application is already using the default port.
Encryption	Specifies all load-balancing communication between the VPN Concentrators is encrypted.
IPSec Shared Secret	Available only if Encryption is checked. The shared secret is a common password used to authenticate all virtual cluster members. IPSec uses this shared secret as a preshared key to establish secure tunnels between virtual cluster peers.
Priority	Priority (1 to 10) for this VPN Concentrator within the virtual cluster. The higher the value, the more likely this device could become the virtual cluster master either at startup or when an existing master fails.

H.323 Support in PAT Mode

VPN 3002 client software supports H.323, the packet-based multimedia communications standard developed by the International Telecommunication Union (ITU). A variety of multimedia applications use the H.323 standard to implement real-time audio,

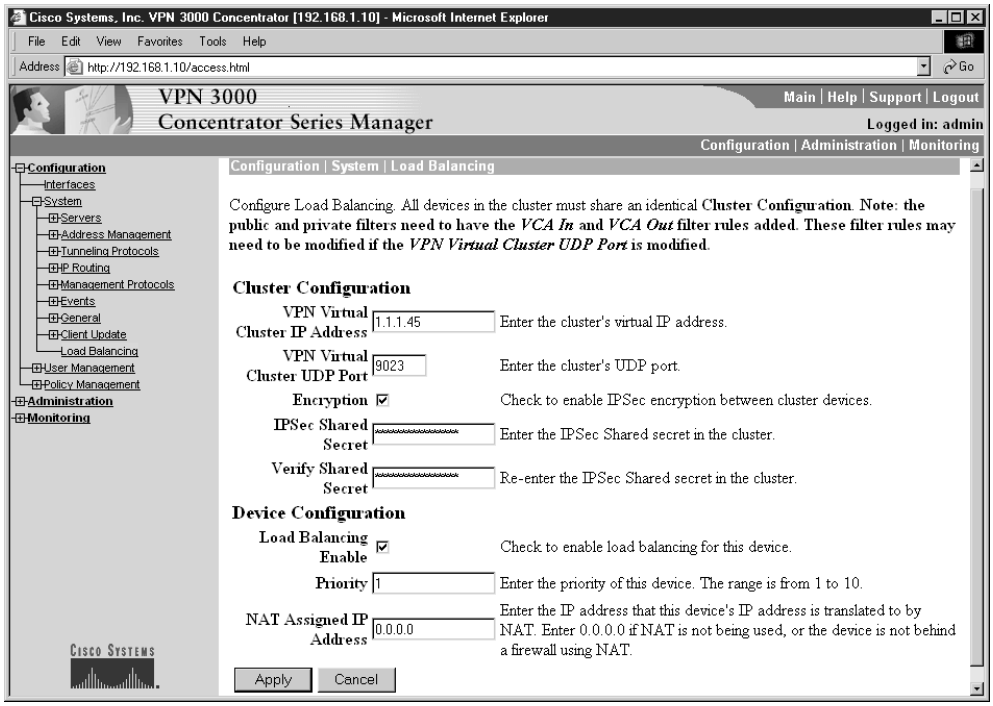

Figure 15-27 Configuring load balancing

video, and data communications. This H.323 support allows the VPN 3002 to support Microsoft NetMeeting. H.323 support requires no configuration on the VPN 3002.

Simple Certificate Enrollment Protocol (SCEP)

You can enroll and install digital certificates on the VPN 3002 manually or automatically. The automatic method is a new feature that uses the Simple Certificate Enrollment Protocol (SCEP) to streamline enrollment and installation. SCEP is a secure messaging protocol that requires minimal user intervention and was introduced in Chapter 11.

This automatic method is quicker than enrolling and installing digital certificates manually, but is available only if the following two conditions are met.

- The CA must support SCEP.
- Enrolling must be done via the Web.

If the CA doesn't support SCEP or if digital certificates are enrolled by other means, such as e-mail or floppy disk, then they must be processed using the manual method.

To allow retrieving the CA certificate via SCEP, use the Administration | Certificate Management | Install | CA Certificate | SCEP, as shown in Figure 15-28.

XML Management

The VPN 3002 now supports an XML-based interface that allows the administrator to use an external management application. These management applications can be Cisco products or third-party tools. XML data can be sent to the VPN Concentrator using HTTPS, SSH, or standard file transfer protocols, such as FTP or TFTP.

This feature is enabled by default and doesn't require configuration.

Reverse Route Injection (RRI)

A VPN Concentrator can be configured to add routes to its routing table for remote hardware or software clients. The VPN Concentrator then advertises these routes to its private network via RIP or OSPF, making the VPN 3002 protected networks known to the main network. This feature is called reverse route injection (RRI) and it was introduced in version 3.5 of the VPN 3000 Concentrator code.

Figure 15-29 shows the VPN scenario used earlier. The scenario assumes the main office has reserved the networks 192.168.0.0 to 192.168.127.0 for its internal use. The other private class C addresses were assigned as needed to branch locations. RRI could be implemented so all main office routers know about the branch office LAN. This assumes the branch office is configured for Network Extension mode.

RRI requires no configuration on the VPN 3002.

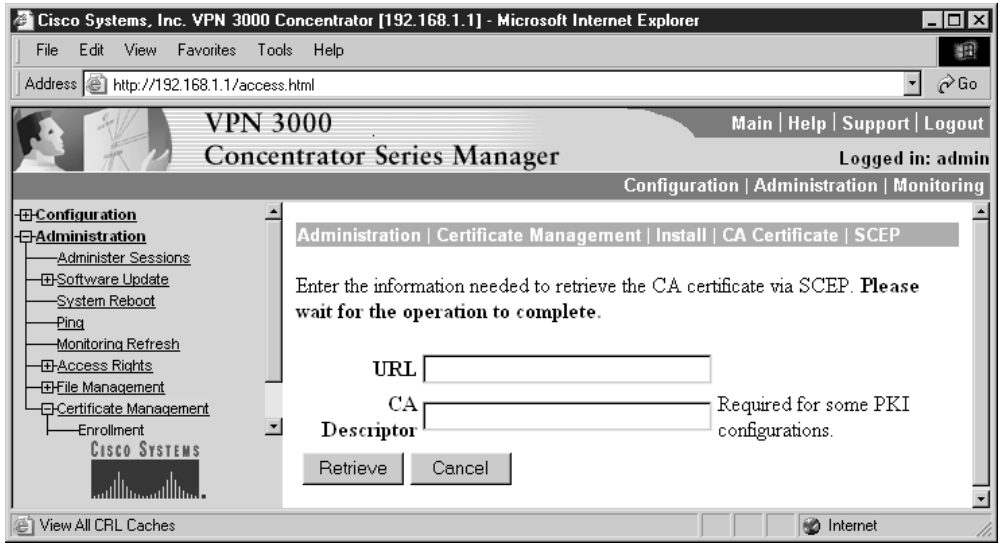

Figure 15-28 Using SCEP to enroll certificates

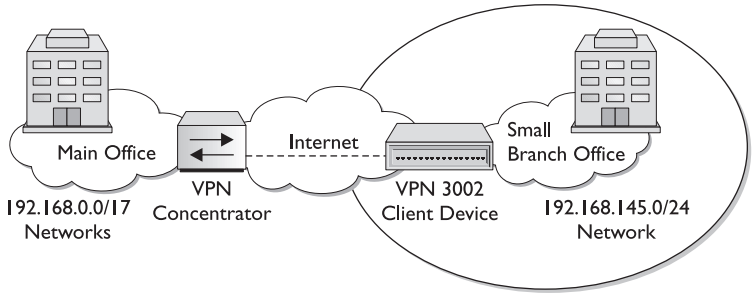

Figure 15-29 VPN scenario network with the main office using RRI

VPN Client Support

Client RRI can be used on all VPN Clients connecting to the VPN Concentrator. This option applies to all remote software clients and VPN 3002 Hardware Clients using Client (PAT) Mode. To configure Client RRI on the VPN Concentrator, go to Configuration | System | IP Routing | Reverse Route Injection, and then select the check box for Client Reverse Route Injection. This selection adds host routes for each remote client to the VPN Concentrator routing table. The VPN Concentrator adds a host route when the client connects and deletes it when the client disconnects. This box is unchecked by default.

Network Extension RRI

This option is for VPN 3002 Client in Network Extension mode (NEM) only. To configure Network Extension RRI on the VPN Concentrator, go to Configuration | System | IP Routing | Reverse Route Injection and select the check box for Network Extension Reverse Route Injection. This selection adds a network route for each network behind a VPN 3002 Hardware Client to the routing table on the VPN Concentrator. The VPN Concentrator adds the route when the VPN 3002 connects and deletes the route when it disconnects. This box is unchecked by default.

Figure 15-30 shows the Configuration | System | IP Routing | Reverse Route Injection screen where both of the RRI features can be configured. The example shows adding the protected LAN from the scenario example.

AES Support and Diffie-Hellman Group 5

VPN software version 3.6 introduced support for Advanced Encryption Standard (AES), which is more secure than DES and more efficient than 3DES. AES supports 128-, 192-, and 256-bit key strengths. 128-bit AES is significantly faster than 168-bit 3DES and little performance difference exists between 256-bit AES and 168-bit 3DES. AES support was also added to the PIX Firewall with v6.3(1).

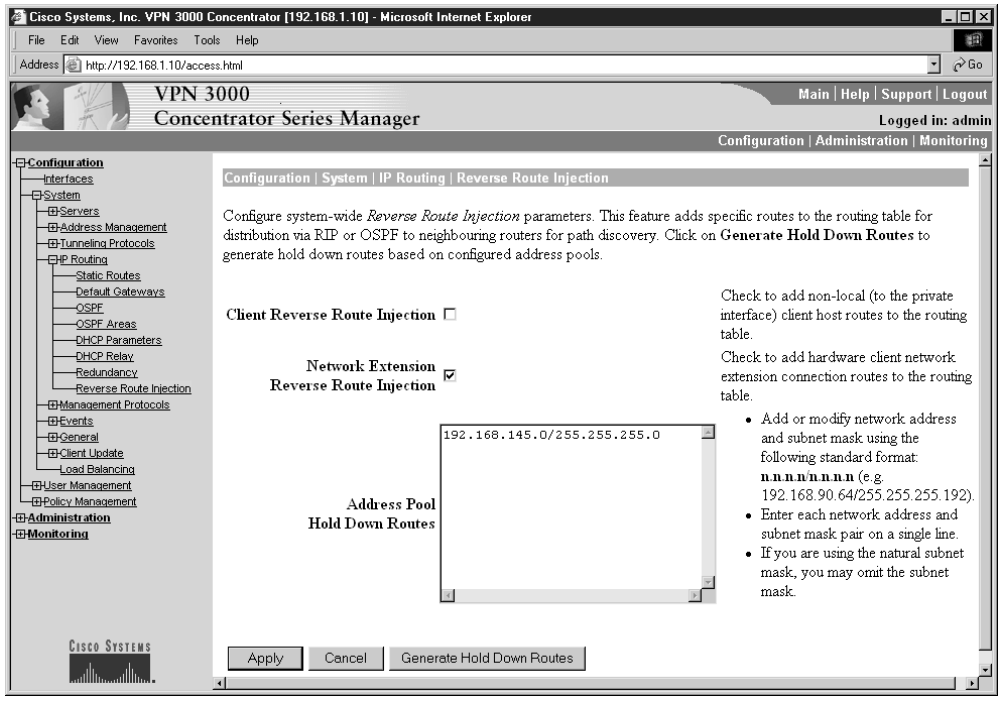

Figure 15-30 Configuring Reverse Route Injection

VPN software version 3.6 also introduced support for Diffie-Hellman Group 5 (1,536-bit key), which provides greater key exchange security. This feature is set as part of IPSec configuration on the VPN Concentrator, as shown in Figure 15-31.

Push Banner to VPN 3002

An administrator can create a banner on the VPN 3000 Concentrator and push it to the VPN 3002. This allows the organization to provide information to users about their network, terms for use, liability, and other issues. The maximum banner length is 510 characters. Any characters, including newline (ENTER key) can be used. Use the Configuration | User Management | Base Group, Client Configuration parameters tab on the VPN Concentrator, as shown in Figure 15-32.

The banner displays only when individual user authentication is enabled.

Delete with Reason

Delete with reason is a part of the system-messaging process that can be used for alerts and troubleshooting. Delete with reason causes the VPN Concentrator to send reasons for VPN Concentrator-initiated disconnects to any impacted software clients or VPN 3002 hardware clients. The client device then decodes the reason and displays it in the event log.

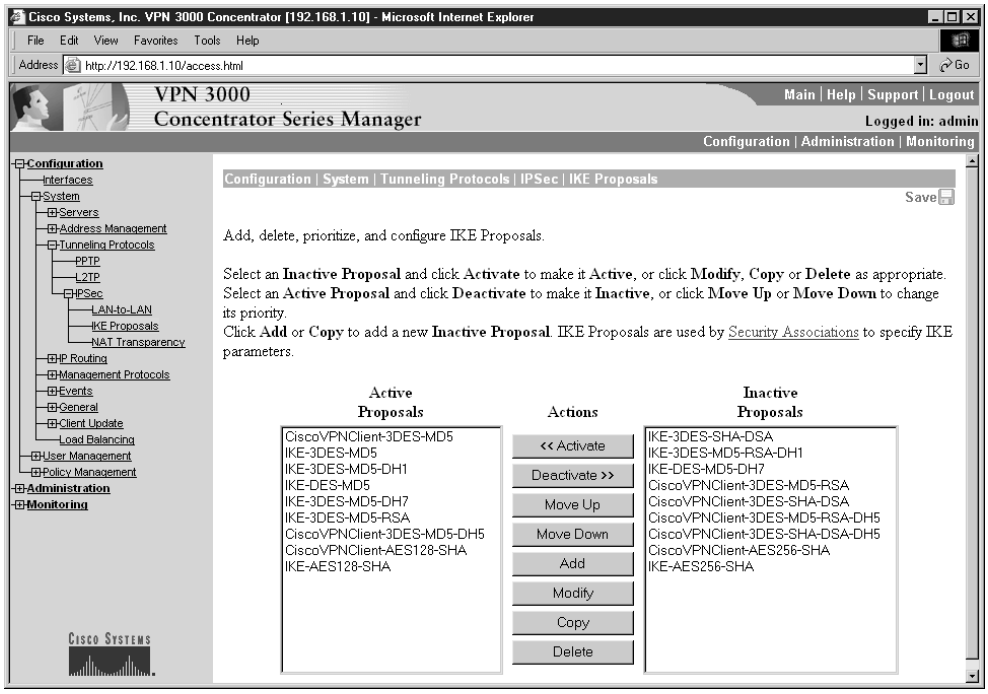

Figure 15-31 Configuring IKE proposals, including AES and DH 1, 5, and 7

Similarly, the VPN 3002 sends reasons for any VPN3002-initiated disconnects to the VPN Concentrator at the central site. The VPN Concentrator decodes the reason and displays it in the event log.

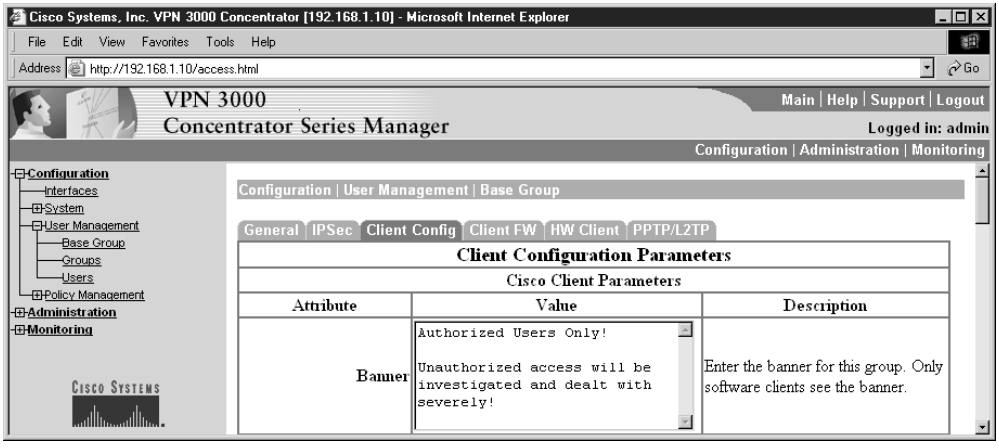

Figure 15-32 Configuring a banner to be pushed to VPN 3002s

PART III

The feature isn't currently supported by the Cisco PIX Firewalls.
The feature is active by default, but an administrator can disable it.

Auto-Update Feature

The *client update* feature was added in version 3.0 for the VPN 3002 Hardware Client allowing VPN central devices, such as VPN Concentrators, to upgrade the software and configuration on the client. In the case of the VPN 3002 Hardware Client, firmware upgrades can also be pushed down to the client.

VPN 3002 Hardware Clients

For VPN 3002 Hardware Clients, the client update allows an administrator to update software and firmware automatically for the 3002 device. If an upgrade is needed, the unit upgrades automatically from an internal TFTP server specified on the central site VPN Concentrator. The process of maintaining security, managing the system, and upgrading it is transparent to the end user.

To avoid update failures and reduce downtime, the VPN 3002 stores image files in two locations: the active and the backup location. The *active location* contains the image currently running on the system. Updating the image overwrites the image in the backup location. That new image is tested to validate it and, if successful, the image is identified as the active location for the next reboot. If the update isn't validated, the client doesn't reboot, and the invalid image doesn't become active. If the update process fails to download a valid image, it will retry up to 20 times, waiting three minutes between attempts. Any unsuccessful updates are logged with information indicating the type of failure.

Cisco VPN Software Clients

For Cisco VPN software clients, the process is a little less automatic: it's more of a notification mechanism with an assisted upgrade. The client update for the Cisco VPN software clients allows central location administrators to notify the client users automatically when it's time to update. Then action is required on the part of users to retrieve and install the newer software.

Configuring Auto-Update

Use the Configuration | System | Client Update menu options of the VPN 3000 Concentrator Manager to configure the client update feature. This screen, as shown in Figure 15-33, offers the following two options.

Enable	Enables or disables client update.
Entries	Configures updates by client type, acceptable firmware and software versions, and their locations.

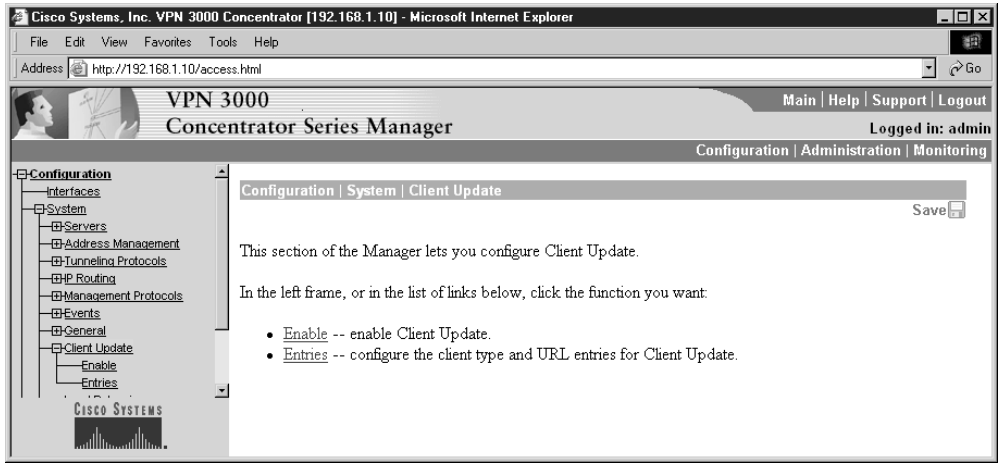

Figure 15-33 Client auto-update screen

1. The Configuration | System | Client Update | Enable screen has only an Enable check box. Select it and click the Apply button to enable or disable the client update feature. The Manager returns to the Configuration | System | Client Update screen.
2. The Configuration | System | Client Update | Entries screen, shown in Figure 15-34, is used to add, modify, or delete client update entries. The update entry list shows the available client updates. If no updates were configured, the list shows "--Empty--".

Use the Add button to configure a new client update entry. The Manager opens the Configuration | System | Client Update | Entries | Add screen, as shown in Figure 15-35.

To modify an existing entry, highlight the update in the Update Entry list and click the Modify button. The Manager opens the Configuration | System | Client Update | Modify screen, which looks like the Add screen.

The Delete button can be used to remove a client update entry.

Chapter Review

This chapter looked at configuring the VPN 3002 client and the VPN 3000 concentrators to support the VPN client features.

The VPN 3002 Hardware Client is a dedicated VPN device that can be used to connect one or more devices, including workstations, servers, hubs, cash registers, printers, and IP telephones to a company's central network. The 3002 unit eliminates the need to install and configure VPN client software on the local workstation(s), plus it supports workstations running any TCP/IP supporting OS.

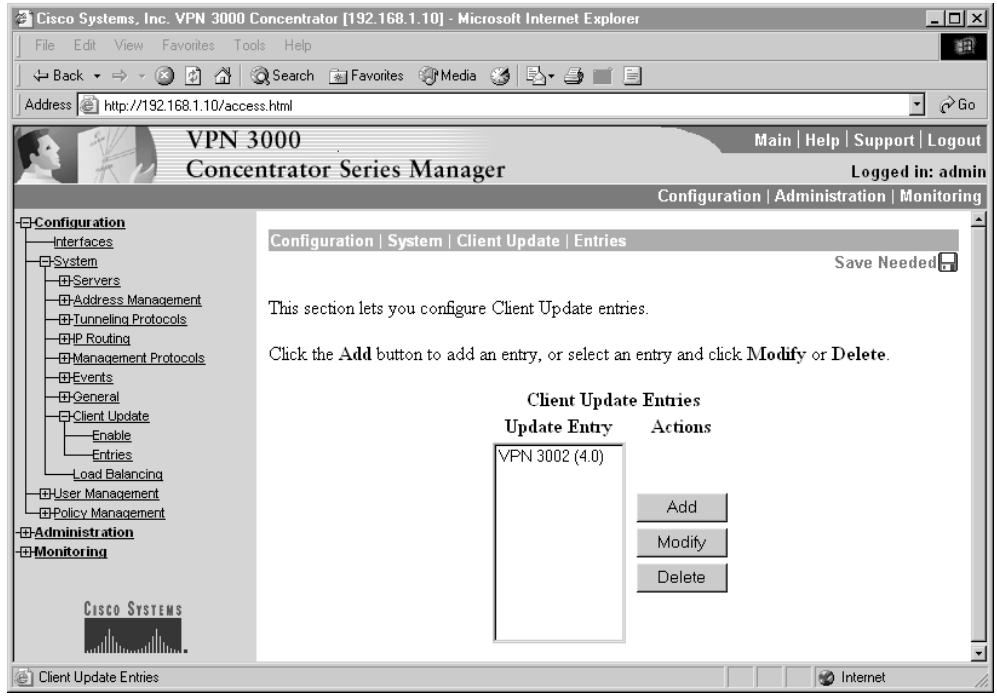

Figure 15-34 Configuration | System | Client Update | Entries screen

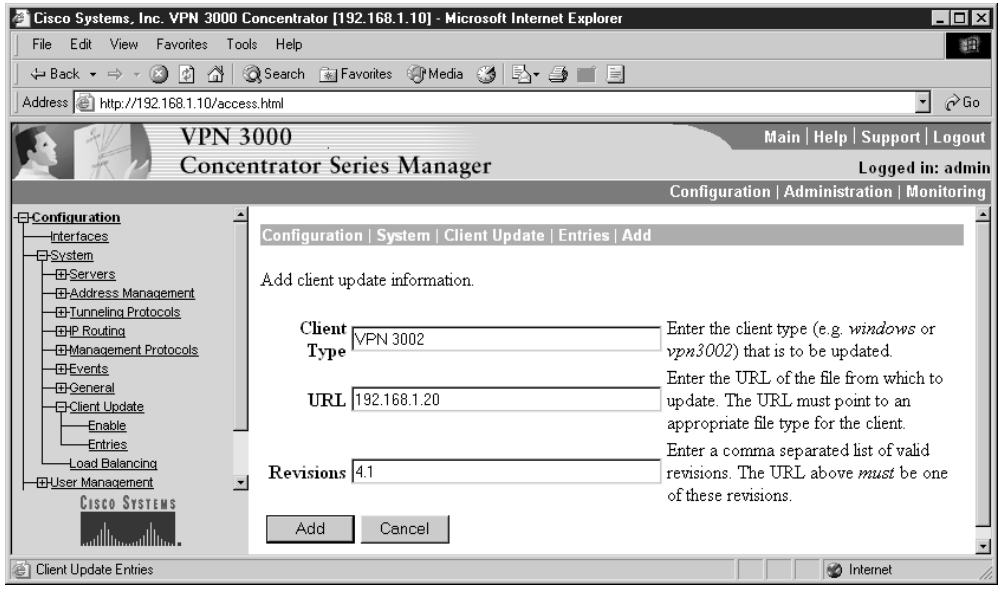

Figure 15-35 Adding or modifying a client update entry

Configuring the VPN 3002 can be accomplished using a menu-driven version of the CLI or the more user-friendly Client Manager web-based application. The two systems have virtually identical steps and menu options making it easier to use both as needed.

The Quick Configuration option is used to provide the basic parameters required for connecting to a VPN Concentrator, which will "push" the basic IPSec configuration and security policy features down to the client. A wide variety of features can be configured on the VPN Concentrator to be pushed out to either VPN 3002 or VPN software clients. An auto-update feature facilitates configuration, software, and firmware upgrades.

Questions

1 What is another name for VPN PAT mode?

A. Network Extension mode

B. Client mode

C. Network Extension Mode per Group

D. Split tunnel mode

2. Which of the following is *not* an IPSec implementation used by the VPN 3002?

A. IPSec over TCP

B. IPSec over UTP

C. IPSec over NAT-T

D. IPSec over PPTP

3. The VPN 3002 CLI is most like which of the following?

A. Cisco IOS routers

B. Cisco Catalyst 5000 and 6000 switches

C. The PIX Firewall

D. None of the above

4. Which one of the following is *not* one of the three main menu sections for the VPN 3002?

A. Monitoring

B. Reporting

C. Configuration

D. Administration

5. What is the command to upgrade the VPN 3002 software?

A. copy tftp flash

B. copy tftp nvram

C. copy tftp backup

D. None of the above

6. Which of the following is a growing protocol for connecting ISPs through broadband connections like DSL/cable systems?
   - **A.** PPP
   - **B.** PPPoE
   - **C.** Ethernet
   - **D.** BGP
7. What is the first feature(s) set with Quick Configuration?
   - **A.** Upload an Existing Configuration file
   - **B.** Configure the Public Interface
   - **C.** Set the System Time, Date, and Time Zone
   - **D.** Configure DNS
8. Where might the DHCP Server feature be configured on the VPN 3002?
   - **A.** Public interface
   - **B.** Private interface
   - **C.** Public or private interfaces
   - **D.** None of the above
9. In which two cases would PAT always be used on outbound traffic?
   - **A.** Network extension mode
   - **B.** Split tunnels
   - **C.** Client mode
   - **D.** PPPoE
10. Which feature can be configured on either the VPN 3002 or the Concentrator?
    - **A.** Interactive Unit Authentication
    - **B.** IPSec Server Load Balancing
    - **C.** Individual User Authentication
    - **D.** IPSec backup servers
11. Which feature requires Virtual Cluster protocol configuration?
    - **A.** Interactive Unit Authentication
    - **B.** IPSec Server load balancing
    - **C.** Individual User Authentication
    - **D.** IPSec backup servers

12. Reverse Route Injection refers to which of the following?
    A. Inserting main network routes in remote network route tables
    B. Making main network routes available to all remote VPN clients
    C. Inserting remote network routes in main network route tables
    D. Exchanging routes both ways between main and branch networks
13. Which Cisco protocol allows wireless users to authenticate over VPN links?
    A. NAT Traversal (NAT-T)
    B. LEAP Bypass
    C. RRI
    D. PPPoE
14. Which of the following is *not* supported by Client Update features?
    A. VPN Software
    B. Firmware on VPN 3002
    C. Push Banner to VPN 3002
    D. VPN configurations
15. Which of the following might create problems for wireless users behind the VPN 3002?
    A. IPSec over NAT-T
    B. Interactive Hardware Client Authentication
    C. Individual User Authentication
    D. AES Support and Diffie-Hellman Group

Answers

1. **B.** Client mode
2. **D.** IPSec over PPTP
3. **D.** None of the above. It is menu-driven, somewhat like the old 1900 switches
4. **B.** Reporting
5. **D.** None of the above. This is a menu-driven process
6. **B.** PPPoE
7. **C.** Set the System Time, Date, and Time Zone
8. **B.** Private interface. Where the LAN connects

9. **B.** Split tunnels, and **C.** Client mode
10. **D.** IPSec backup servers
11. **B.** IPSec Server load balancing
12. **C.** Inserting remote network routes in main network route tables
13. **B.** LEAP Bypass
14. **C.** Push Banner to VPN 3002
15. **B.** Interactive Hardware Client Authentication. The wireless clients would have to wait for a wired user to authenticate the VPN 3002 to the Concentrator.