



FIRST-STEP SERIES

Network Security **first-step**

Second Edition

Your first step into the world
of **network security**

- No security experience required
- Includes clear and easily understood explanations
- Makes learning easy



Network Security First-Step

Tom Thomas
Donald Stoddard

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

Network Security First-Step

Tom Thomas
Donald Stoddard

Copyright© 2012 Cisco Systems, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

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

First Printing December 2011

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

ISBN-13: 978-1-58720-410-4

ISBN-10: 1-58720-410-X

Warning and Disclaimer

This book is designed to provide information about network security. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

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

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419 corpsales@pearsoned.com

For sales outside of the U.S. please contact: **International Sales** international@pearsoned.com

Feedback Information

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

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

We greatly appreciate your assistance.

Publisher: Paul Boger	Business Operation Manager, Cisco Press: Anand Sundaram
Associate Publisher: Dave Dusthimer	Manager Global Certification: Erik Ullanderson
Executive Editor: Brett Bartow	Senior Development Editor: Christopher Cleveland
Managing Editor: Sandra Schroeder	Copy Editor: Apostrophe Editing Services
Senior Project Editor: Tonya Simpson	Technical Editors: Phil Lerner, James Risler
Editorial Assistant: Vanessa Evans	Proofreader: Mike Henry
Cover Designer: Sandra Schroeder	Indexer: Cheryl Lenser
Composition: Mark Shirar	



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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

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

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Authors

Tom Thomas, CCIE No. 9360, claims he never works because he loves what he does. When you meet him, you will agree!

Throughout his many years in the networking industry, Tom has taught thousands of people how networking works and the secrets of the life of a packet. Tom is the author or coauthor of 18 books on networking, including the acclaimed *OSPF Network Design Solutions*, published by Cisco Press and now in its second edition. Beyond his many books, Tom also has taught computer and networking skills through his roles as an instructor and training-course developer.

In addition to holding the Cisco Certified Internetwork Expert (CCIE) certification—the pinnacle of networking certifications—Tom holds Cisco CCNP Security, CCDA, and CCNA certifications and is a certified Cisco Systems instructor (CCSI). These certifications support his industry-proven, problem-solving skills through technical leadership with demonstrated persistence and the ability to positively assist businesses in leveraging IT resources in support of their core business. He has also completed his Master of Science degree in network architecture and is looking at a doctorate next.

Tom currently is the CIO of Qconcert, a Cisco Gold Partner in Southern Florida that has an affiliated arm known as CCPrep.com, a Cisco Learning Partner, where he provides strategic direction and a little hands-on for customers of all types.

Donald Stoddard began his career in information technology in 1998, designing networks and implementing security for schools in North Dakota and South Dakota. He then went on to design and implement Geographical Information Systems (GIS) for a firm in Denver, Colorado. While there, he earned his Bachelor of Science degree in computer information systems management from Colorado Christian University. From Colorado, he then moved south, learned the ins-and-outs of Cisco VoIP, and began working through designing and securing VoIP solutions throughout the southeast. Don holds Microsoft MCSA and Linux+ and Security+ certifications and is presently wading through the CISSP material.

Currently, Don works for the Department of the Navy as the Information Assurance Officer for one of the premier Navy research and development labs, where he provides certification and accreditation guidance for the various projects being developed for implementation and deployment.

About the Technical Reviewers

Phil Lerner, CISSP, GFSP, GAWN, CHS-IV, CGEIT, ECSA, C-EH is an industry veteran with 20 years of experience covering information security. Most recently, Phil was one of the few senior technical solutions architects at Cisco Systems focused on Data Center and Security. Phil's areas of expertise include sanctioned attack and penetration, digital and network forensics, wireless security, network security architecture, and policy work. Phil is also an adjunct professor at St. John's University in Queens, New York, teaching wireless security to all levels of undergraduate students. Phil earned his MS-CIS (Cyber Security) from Boston University in 2009 and is a frequent information security show speaker and trusted advisor to many large firms.

James Risler, CCIE No. 15412, is a systems engineer education specialist for Cisco. His focus is on security technology and training development. James has more than 18 years of experience in IP internetworking, including the design and implementation of enterprise networks. Prior to joining Cisco, James provided Cisco training and consulting for Fortune 500 companies and government agencies. He holds two bachelor's degrees from University of South Florida and is currently working on his MBA at the University of Tampa.

Dedications

Tom Thomas: How do you put into words the importance someone has in your life? Love and time strengthens the emotions until they are so powerful they make you want to express them in a meaningful way. I dedicate this book and this poem to my partner and soul mate, Kristi. During the course of this writing we found out together that we are having a child, twins in fact, and I welcome them into our life with open arms.

How do I begin to tell you how
lucky I am to have you in my life?
I'll start by saying what a gift you
gave me the day you became my wife.

In you I have truly found
An Angel who walks upon the ground.
You go beyond all limits for me
Just to show your love endlessly.
I could search my whole life through
And never find another "you."
You are so special that I wanted you to know
I truly, completely love you so.

You must be an angel without wings
To put up with all of my bothersome things
My anger, my love, my sometimes weary heart
What others hated about me you love
How could I not love you with all that I am
You are the steady I need for my trembling hand
You simply must be an angel without wings!

You're my best friend in the good times
and my rock in times of sorrow.
You're the reason for sweet yesterdays
and my promise for tomorrow.

I never thought I could feel this loved
until you became my wife.
You made this year and every year
the best one of my life.

Donald Stoddard: To AJ, my friend, my lover, my wife and queen. You have done the impossible...you've made me believe in myself again. From the moment I saw you across the room I knew you were the other half my soul longed for. Thank you for your love, support, and strength: ost min kis mik.

‡ † † : † | † : † | † : † | †

Acknowledgments

Tom Thomas: Special acknowledgments go to my good friend and the best editor, Chris Cleveland. His insight, abilities, and editorial comments take a rough manuscript and gave it life beyond what a simple nerd was able to envision. I have had the pleasure of working with Chris for many years, and I do not think I would ever want to write a book without his involvement.

As always, I would like to thank my technical editors for their friendship, insight, and awesome comments. Your knowledge helped to fine-tune my thoughts. I know that this book will help many people, and that was the goal. Thank you.

Don, we have been friends for years and you have always been a part of my life through the good and the bad; I am lucky to call you brother.

Donald Stoddard: I would like to extend a great thank-you for a great staff: Brett Bartow, Vanessa Evans, Chris Zahn, Chris Cleveland, and the technical reviewers (James Risler and Phil Lerner); without your patience and attention to detail this book would not be in the hands of readers today. Honestly, without you to guide, push, and correct, none of this is possible. Thank you all for your hard work and contributions throughout the long months from start to finish...truly this has been a marathon, not a sprint, and it has been a pleasure from the beginning.

And finally, I want to acknowledge a man who has guided my career and life for a long time. Tom, we've known each other for many years, and you have always been there to guide me when my career was derailed. You have been an inspiration. I will always remember you telling me to get focused. In fact, I think your words to me were, "...Don, you know what your problem is? You lack focus...." We've never been people who mince words, have we? I have focus now, I have a plan, and I have a career set before me all because of you. Thank you for your professional guidance and your friendship.

Contents at a Glance

	Introduction	xxii
Chapter 1	There Be Hackers Here!	1
Chapter 2	Security Policies	45
Chapter 3	Processes and Procedures	85
Chapter 4	Network Security Standards and Guidelines	105
Chapter 5	Overview of Security Technologies	127
Chapter 6	Security Protocols	169
Chapter 7	Firewalls	193
Chapter 8	Router Security	217
Chapter 9	IPsec Virtual Private Networks (VPNs)	257
Chapter 10	Wireless Security	299
Chapter 11	Intrusion Detection and Honeypots	331
Chapter 12	Tools of the Trade	359
Appendix A	Answers to Review Questions	389
Index		403

Contents

Introduction xxii

Chapter 1 **There Be Hackers Here!** 1

Essentials First: Looking for a Target	2
Hacking Motivations	3
Targets of Opportunity	4
Are You a Target of Opportunity?	6
Targets of Choice	7
Are You a Target of Choice?	7
The Process of an Attack	9
Reconnaissance	9
Footprinting (aka Casing the Joint)	11
Scanning	18
Enumeration	23
<i>Enumerating Windows</i>	24
Gaining Access	26
<i>Operating System Attacks</i>	27
<i>Application Attacks</i>	27
<i>Misconfiguration Attacks</i>	28
<i>Scripted Attacks</i>	29
Escalating Privilege	30
Covering Tracks	31
Where Are Attacks Coming From?	32
Common Vulnerabilities, Threats, and Risks	33
Overview of Common Attacks and Exploits	36
Network Security Organizations	39
CERT Coordination Center	40
SANS	40
Center for Internet Security (CIS)	40
SCORE	41
Internet Storm Center	41
National Vulnerability Database	41
Security Focus	42
Learning from the Network Security Organizations	42
Chapter Summary	43
Chapter Review	43

Chapter 2 Security Policies 45

- Responsibilities and Expectations 50
 - A Real-World Example 50
 - Who Is Responsible? You Are! 50
 - Legal Precedence* 50
 - Internet Lawyers* 51
 - Evolution of the Legal System* 51
 - Criminal Prosecution 52
 - Real-World Example* 52
 - Individuals Being Prosecuted* 53
 - International Prosecution* 53
- Corporate Policies and Trust 53
 - Relevant Policies 54
 - User Awareness Education 54
 - Coming to a Balance 55
 - Corporate Policies 55
- Acceptable Use Policy 57
 - Policy Overview 57
 - Purpose 58
 - Scope 58
 - General Use and Ownership 58
 - Security and Proprietary Information 59
 - Unacceptable Use 60
 - System and Network Activities* 61
 - Email and Communications Activities* 62
 - Enforcement 63
 - Conclusion 63
- Password Policy 64
 - Overview 64
 - Purpose 64
 - Scope 64
 - General Policy 65
 - General Password Construction Guidelines 66
 - Password Protection Standards 67
 - Enforcement 68
 - Conclusion 68

Virtual Private Network (VPN) Security Policy	69
Purpose	69
Scope	69
Policy	70
Conclusion	71
Wireless Communication Policy	71
Scope	72
Policy Statement	72
<i>General Network Access Requirements</i>	72
<i>Lab and Isolated Wireless Device Requirements</i>	72
<i>Home Wireless Device Requirements</i>	73
Enforcement	73
Definitions	73
Revision History	73
Extranet Connection Policy	74
Purpose	74
Scope	74
Security Review	75
Third-Party Connection Agreement	75
Business Case	75
Point of Contact	75
Establishing Connectivity	75
Modifying or Changing Connectivity and Access	76
Terminating Access	76
Conclusion	76
ISO Certification and Security	77
Delivery	77
ISO/IEC 27002	78
Sample Security Policies on the Internet	79
Industry Standards	79
Payment Card Industry Data Security Standard (PCI DSS)	80
Sarbanes-Oxley Act of 2002 (SOX)	80
Health Insurance Portability and Accounting Act (HIPAA) of 1996	81
Massachusetts 201: Standards for the Protection of Personal Information of Residents of the Commonwealth	81
SAS 70 Series	82
Chapter Summary	82
Chapter Review	83

Chapter 3 Processes and Procedures 85

- Security Advisories and Alerts: Getting the Intel You Need to Stay Safe 86
 - Responding to Security Advisories 87
 - Step 1: Awareness* 88
 - Step 2: Incident Response* 90
 - Step 3: Imposing Your Will* 95
 - Steps 4 and 5: Handling Network Software Updates (Best Practices)* 96
- Industry Best Practices 98
 - Use a Change Control Process 98
 - Read All Related Materials 98
 - Apply Updates as Needed 99
 - Testing 99
 - Uninstall 99
 - Consistency 99
 - Backup and Scheduled Downtime 100
 - Have a Back-Out Plan 100
 - Forewarn Helpdesk and Key User Groups 100
 - Don't Get More Than Two Service Packs Behind 100
 - Target Noncritical Servers/Users First 100
 - Service Pack Best Practices 101
 - Hotfix Best Practices 101
 - Service Pack Level Consistency* 101
 - Latest Service Pack Versus Multiple Hotfixes* 101
 - Security Update Best Practices 101
 - Apply Admin Patches to Install Build Areas* 102
 - Apply Only on Exact Match* 102
 - Subscribe to Email Notification* 102
- Summary 102
- Chapter Review and Questions 104

Chapter 4 Network Security Standards and Guidelines 105

- Cisco SAFE 2.0 106
 - Overview 106
 - Purpose 106
- Cisco Validated Design Program 107
 - Branch/WAN Design Zone Guides 107
 - Campus Design Zone Guides 107

Data Center Design Zone Guides	108
Security Design Zone Guides	109
Cisco Best Practice Overview and Guidelines	110
Basic Cisco IOS Best Practices	110
<i>Secure Your Passwords</i>	110
<i>Limit Administrative Access</i>	111
<i>Limit Line Access Controls</i>	111
<i>Limit Access to Inbound and Outbound Telnet (aka vty Port)</i>	112
<i>Establish Session Timeouts</i>	113
<i>Make Room Redundancy</i>	113
<i>Protect Yourself from Common Attacks</i>	114
Firewall/ASAs	115
<i>Encrypt Your Privileged User Account</i>	115
<i>Limit Access Control</i>	116
<i>Make Room for Redundant Systems</i>	116
<i>General Best Practices</i>	117
<i>Configuration Guides</i>	117
<i>Intrusion Prevention System (IPS) for IOS</i>	117
NSA Security Configuration Guides	118
Cisco Systems	119
<i>Switches Configuration Guide</i>	119
<i>VoIP/IP Telephony Security Configuration Guides</i>	119
Microsoft Windows	119
<i>Microsoft Windows Applications</i>	120
<i>Microsoft Windows 7/Vista/Server 2008</i>	120
<i>Microsoft Windows XP/Server 2003</i>	121
Apple	121
Microsoft Security	121
Security Policies	121
<i>Microsoft Windows XP Professional</i>	122
<i>Microsoft Windows Server 2003</i>	122
<i>Microsoft Windows 7</i>	122
<i>Windows Server 2008</i>	123
Microsoft Security Compliance Manager	124
Chapter Summary	125
Chapter Link Toolbox Summary	125

Chapter 5 Overview of Security Technologies 127

- Security First Design Concepts 128
- Packet Filtering via ACLs 131
 - Grocery List Analogy 132
 - Limitations of Packet Filtering 136
- Stateful Packet Inspection 136
 - Detailed Packet Flow Using SPI 138
 - Limitations of Stateful Packet Inspection 139
- Network Address Translation (NAT) 140
 - Increasing Network Security 142
 - NAT's Limitations 143
- Proxies and Application-Level Protection 144
 - Limitations of Proxies 146
- Content Filters 147
 - Limitations of Content Filtering 150
- Public Key Infrastructure 150
 - PKI's Limitations 151
- Reputation-Based Security 152
 - Reactive Filtering Can't Keep Up 154
 - Cisco Web Reputation Solution 155
- AAA Technologies 156
 - Authentication 156
 - Authorization 157
 - Accounting 157
 - Remote Authentication Dial-In User Service (RADIUS) 158
 - Terminal Access Controller Access Control System (TACACS) 159
 - TACACS+ Versus RADIUS 160
- Two-Factor Authentication/Multifactor Authentication 161
 - IEEE 802.1x: Network Access Control (NAC) 162
 - Network Admission Control* 163
 - Cisco TrustSec 164
 - Solution Overview* 164
 - Cisco Identity Services Engine* 166
- Chapter Summary 168
- Chapter Review Questions 168

Chapter 6	Security Protocols	169
	Triple DES Encryption	171
	Encryption Strength	171
	Limitations of 3DES	172
	Advanced Encryption Standard (AES)	172
	Different Encryption Strengths	173
	Limitations of AES	173
	Message Digest 5 Algorithm	173
	MD5 Hash in Action	175
	Secure Hash Algorithm (SHA Hash)	175
	Types of SHA	176
	SHA-1	176
	SHA-2	176
	Point-to-Point Tunneling Protocol (PPTP)	177
	PPTP Functionality	177
	Limitations of PPTP	178
	Layer 2 Tunneling Protocol (L2TP)	179
	L2TP Versus PPTP	180
	Benefits of L2TP	180
	L2TP Operation	181
	Secure Shell (SSH)	182
	SSH Versus Telnet	184
	SSH Operation	186
	Tunneling and Port Forwarding	187
	Limitations of SSH	188
	SNMP v3	188
	Security Built In	189
	Chapter Summary	192
	Chapter Review Questions	192
Chapter 7	Firewalls	193
	Firewall Frequently Asked Questions	194
	Who Needs a Firewall?	195
	Why Do I Need a Firewall?	195
	Do I Have Anything Worth Protecting?	195
	What Does a Firewall Do?	196
	Firewalls Are “The Security Policy”	197
	We Do Not Have a Security Policy	200

- Firewall Operational Overview 200
 - Firewalls in Action 202
 - Implementing a Firewall 203
 - Determine the Inbound Access Policy 205
 - Determine Outbound Access Policy 206
- Essentials First: Life in the DMZ 206
- Case Studies 208
 - Case Study: To DMZ or Not to DMZ? 208
- Firewall Limitations 214
- Chapter Summary 215
- Chapter Review Questions 216

Chapter 8 Router Security 217

- Edge Router as a Choke Point 221
 - Limitations of Choke Routers 223
- Routers Running Zone Based Firewall 224
 - Zone-Based Policy Overview 225
 - Zone-Based Policy Configuration Model 226
 - Rules for Applying Zone-Based Policy Firewall 226
 - Designing Zone-Based Policy Network Security 227
 - Using IPsec VPN with Zone-Based Policy Firewall 228
- Intrusion Detection with Cisco IOS 229
 - When to Use the FFS IDS 230
 - FFS IDS Operational Overview 231
 - FFS Limitations 233
- Secure IOS Template 234
- Routing Protocol Security 251
 - OSPF Authentication 251
 - Benefits of OSPF Neighbor Authentication* 252
 - When to Deploy OSPF Neighbor Authentication* 252
 - How OSPF Authentication Works* 253
- Chapter Summary 254
- Chapter Review Questions 255

Chapter 9 IPsec Virtual Private Networks (VPNs) 257

- Analogy: VPNs Securely Connect IsLANds 259
- VPN Overview 261
 - VPN Benefits and Goals 263

VPN Implementation Strategies	264
Split Tunneling	265
Overview of IPsec VPNs	265
Authentication and Data Integrity	268
Tunneling Data	269
VPN Deployment with Layered Security	270
IPsec Encryption Modes	271
<i>IPsec Tunnel Mode</i>	271
<i>Transport Mode</i>	272
IPsec Family of Protocols	272
Security Associations	273
ISAKMP Overview	273
Internet Key Exchange (IKE) Overview	274
<i>IKE Main Mode</i>	274
<i>IKE Aggressive Mode</i>	275
IPsec Security Association (IPsec SA)	275
IPsec Operational Overview	276
<i>IKE Phase 1</i>	277
<i>IKE Phase 2</i>	278
<i>Perfect Forward Secrecy</i>	278
<i>Diffie-Hellman Algorithm</i>	279
Router Configuration as VPN Peer	281
Configuring ISAKMP	281
<i>Preshared Keys</i>	282
Configuring the ISAKMP Protection Suite	282
Configuring the ISAKMP Key	283
Configuring IPsec	284
<i>Step 1: Create the Extended ACL</i>	284
<i>Step 2: Create the IPsec Transforms</i>	284
<i>Step 3: Create the Crypto Map</i>	285
<i>Step 4: Apply the Crypto Map to an Interface</i>	286
Firewall VPN Configuration for Client Access	286
Step 1: Define Interesting Traffic	288
Step 2: IKE Phase 1 ^[udp port 500]	288
Step 3: IKE Phase 2	288
Step 4: Data Transfer	289
Step 5: Tunnel Termination	289

- SSL VPN Overview 289
- Comparing SSL and IPsec VPNs 290
- Which to Deploy: Choosing Between IPsec and SSL VPNs 292
- Remote-Access VPN Security Considerations 293
 - Steps to Securing the Remote-Access VPN 294
 - Cisco AnyConnect VPN Secure Mobility Solution* 295
- Chapter Summary 296
- Chapter Review Questions 297

Chapter 10 Wireless Security 299

- Essentials First: Wireless LANs 301
 - What Is Wi-Fi? 302
 - Benefits of Wireless LANs 303
 - Wireless Equals Radio Frequency 303
- Wireless Networking 304
 - Modes of Operation 305
 - Coverage 306
 - Bandwidth Availability 307
- WarGames* Wirelessly 307
 - Warchalking 308
 - Wardriving 309
 - Warspamming 311
 - Warspying 312
- Wireless Threats 312
 - Sniffing to Eavesdrop and Intercept Data 313
 - Denial-of-Service Attacks 315
 - Rogue/Unauthorized Access Points 316
 - Misconfiguration and Bad Behavior 317
 - AP Deployment Guidelines* 317
 - Wireless Security 318
 - Service Set Identifier (SSID) 318
 - Device and Access Point Association 319
 - Wired Equivalent Privacy (WEP) 319
 - WEP Limitations and Weaknesses* 320
 - MAC Address Filtering 320
- Extensible Authentication Protocol (EAP) 321
 - LEAP 322
 - EAP-TLS 322

EAP-PSK	323
EAP-TTLS	323
Essential Wireless Security	323
Essentials First: Wireless Hacking Tools	325
NetStumbler	325
Wireless Packet Sniffers	326
Aircrack-ng	327
OmniPeek	327
Wireshark	329
Chapter Summary	329
Chapter Review Questions	330

Chapter 11 Intrusion Detection and Honeypots 331

Essentials First: Intrusion Detection	333
IDS Functional Overview	335
<i>Host Intrusion Detection System</i>	340
<i>Network Intrusion Detection System</i>	341
<i>Wireless IDS</i>	343
<i>Network Behavior Analysis</i>	344
How Are Intrusions Detected?	345
Signature or Pattern Detection	346
Anomaly-Based Detection	346
Stateful Protocol Analysis	347
Combining Methods	347
Intrusion Prevention	347
IDS Products	348
<i>Snort!</i>	348
Limitations of IDS	350
Essentials First: Honeypots	354
Honeypot Overview	354
Honeypot Design Strategies	356
Honeypot Limitations	357
Chapter Summary	357
Chapter Review Questions	357

Chapter 12 Tools of the Trade 359

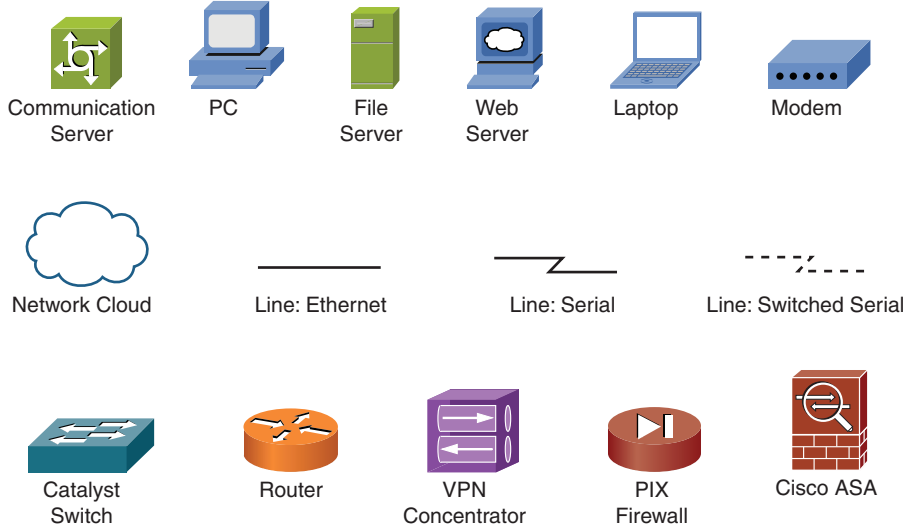
Essentials First: Vulnerability Analysis	361
Fundamental Attacks	361
<i>IP Spoofing/Session Hijacking</i>	362

- Packet Analyzers* 363
- Denial of Service (DoS) Attacks* 363
- Other Types of Attacks* 366
- Back Doors* 368
- Security Assessments and Penetration Testing 370
 - Internal Vulnerability and Penetration Assessment 370
 - Assessment Methodology* 371
 - External Penetration and Vulnerability Assessment 371
 - Assessment Methodology* 372
 - Physical Security Assessment 373
 - Assessment Methodology* 373
 - Miscellaneous Assessments 374
 - Assessment Providers* 375
- Security Scanners 375
 - Features and Benefits of Vulnerability Scanners 376
 - Freeware Security Scanners 376
 - Metasploit* 376
 - NMAP* 376
 - SAINT* 377
 - Nessus* 377
 - Retina Version 5.11.10* 380
- CORE IMPACT Pro (a Professional Penetration Testing Product) 382
 - In Their Own Words 383
 - Scan and Detection Accuracy 384
 - Documentation 384
 - Documentation and Support 386
 - Vulnerability Updates 386
- Chapter Summary 386
- Chapter Review Questions 387

Appendix A Answers to Review Questions 389

Index 403

Icons



Command Syntax Conventions

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

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

Introduction

This book was written to address the need for increased understanding of network security. Many texts are available on the subject, and they have value. However, many people and companies are now considering increasing their network security. Where do you start? Perhaps you want to deploy wireless and you need to ensure that it is secure. What single resource can provide you with a good overview of wireless security or firewalls, and so on? This book provides you with enough security information that you can leverage your newfound knowledge for your own benefit and for the benefit of your organization.

This book was written from the standpoint that every reader needs security but does not actually understand the risks and available techniques and possibilities. Each chapter addresses a specific aspect of an overall layered security model and enables you to see and understand why security for each area is needed, what you should consider, and how you should proceed.

Goals and Methods

The goal of this book is to provide a resource for every person concerned with security. Readers do not have to be networking professionals or CIOs to benefit from this book, although they can as well. It is our hope that all readers, from students to professionals, will benefit from this book.

You can explore each component of the network and verify how it can be securely deployed. When complex security technologies or concepts are encountered, they are explained with real-world examples and practical analogies. This book covers serious topics, but it should also be fun and easy to read. We have endeavored to meet this goal.

Who Should Read This Book?

This book was written with a broad audience in mind. Consider students who are hearing all about the importance of network security and want to focus on this area. This book helps them by providing an understanding of all the major components of securing a network. Perhaps you are a networking professional with in-depth expertise in routing and switching, and now you have been asked to deploy wireless (securely). This book provides a solid foundation upon which to explore the subject matter in more depth, while understanding the different components necessary for accomplishing your goals. You might even be a CIO who has been tasked with determining whether you should invest in an intrusion detection system (IDS). Perhaps you need to understand why this is needed, how it works, and when/where to use it.

Regardless of your expertise or role in the IT industry, this book has a place for you; it takes concepts and simplifies them to give you a solid foundation of understanding. What you do with that knowledge is up to you. This book might give you what you need, or it might be the first step in your journey.

How This Book Is Organized

Although you could read this book cover-to-cover, it is designed to be flexible and enable you to easily move between chapters and sections of chapters to cover only the material you need. If you do intend to read them all, the order in which they are presented is an excellent sequence.

Chapters 1 through 12 cover the following topics:

- **Chapter 1, “There Be Hackers Here”:** Provides a glimpse into the mind and motivation of the individuals who attack your systems. This chapter covers tools, techniques, and attacks.
- **Chapter 2, “Security Policies”:** Starts the defense-in-depth concept with the foundation of securing your network, which is the security policy. This chapter goes over roles and responsibilities within your organization, defines various corporate policies, and then goes over industry standards in use that you should be aware of. When you finish with the chapter, you will understand the role that policies play and one of the ways to prepare/respond to incidents.
- **Chapter 3, “Processes and Procedures”:** Discusses common security operating processes and provides an overview of how to implement those processes and procedures from the ground up. This chapter also includes some industry best practices that are sure to help you and your organization.
- **Chapter 4, “Network Security Standards and Guidelines”:** Goes into depth on the industry standards and guidelines for security implementation within your organization for Cisco, Microsoft, and Macintosh products. It then gives some best practices for implementing and configuring various security devices, such as your Cisco IOS, firewall/ASA, and intrusion prevention system (IPS).
- **Chapter 5, “Overview of Security Technologies”:** Discusses the nuts and bolts of how to use security technologies from the most basic access control lists available in every router to global solutions such as PKI. Many of these technologies are used today without your needing to fully understand when or where they operate. After reading this chapter, you will understand the benefits of these technologies, where they operate, and some of the risks associated with them.
- **Chapter 6, “Security Protocols”:** Looks at security from an encryption protocol implementation point of view. In addition, it considers the limitations of each covered security protocol because nothing is perfect.
- **Chapter 7, “Firewalls”:** Covers firewalls and how they operate. It examines who needs a firewall and why they are an essential part of your network’s defense.

- **Chapter 8, “Router Security”:** If you have a network, you have a router; they have evolved over the years and are now effective security devices. This chapter discusses the expanded security capabilities of routers.
- **Chapter 9, “IPsec Virtual Private Networks (VPN)”:** Discusses the role of VPNs and how they are reshaping the public Internet, encrypting all information that flows across the Internet. This includes the functional characteristics and operational parameters.
- **Chapter 10, “Wireless Security”:** Discusses the hottest technology, wireless, and explains that all is not well in this IT nirvana. Hackers have also come here, and they bring a full complement of tools. Many think that wireless is safe and easy; this chapter ensures that those people become security conscious.
- **Chapter 11, “Intrusion Detection and Honeypots”:** Discusses how you can detect a hacker’s attempt to gain access into your network by implementing an intrusion detection system (IDS) or intrusion prevention system (IPS). It compares and contrasts the two so that you understand the role of each device. In addition, it discusses one of the ways to confuse a hacker—through the use of a honeypot.
- **Chapter 12, “Tools of the Trade”:** Chapter 1 warns you that there be hackers . . . this chapter helps you understand what you are up against by discussing the various methods and tools used by hackers to infiltrate computer systems. This chapter then examines the available tools for identifying weaknesses in your network and the anatomy of a security audit, which is a crucial piece for ensuring that a network is secure and thus foiling the bad guy.

There Be Hackers Here!

When the ancient mapmakers reached the edge of the known world they wrote on their maps, “There Be Dragons Here!”

This chapter discusses in broad strokes the anatomy of a hacker attack from the beginning steps of finding the right target with recon and enumeration to executing the attack to cleanup. You learn some of the factors and footprints of hackers, enabling you to understand the emerging threats and potential exploits.

By the end of this chapter, you should know and be able to explain the following:

- What are hacker motivations and how are they evolving?
- What is the difference between a target of opportunity and a target of choice?
- What are the major components of an attack and the purpose of each?
- What are the breadth and scope of the possible attacks and exploits available to attackers?
- Where are the online security organizations and how can they assist you?

Answering these key questions will enable you to understand the overall characteristics and importance of network security. By the time you finish this book, you will have a solid appreciation for network security and understand its issues, how it works, and why it is important enough to include in every home and corporate network.

In today’s interconnected world, this ancient representation of the world beyond a person’s knowledge holds true. When you connect your home or corporate network to the Internet, everything beyond your network is literally the edge of the world to you and the beginning of the World Wide Web (the home of dragons), wherein hackers are looking to take advantage of the unwary.

There Be Hackers Here!

It is hard for people who are not involved in IT to understand why someone would want to hack or otherwise intentionally harm someone else. The motivations behind these behaviors might be easier to understand after you complete this book.

In a book about understanding network security, the obvious first step is to introduce and review what a hacker is and some of the methods a hacker employs to threaten your network.

From finding the right target to executing the attack, this chapter provides an overview of a hacker attack's anatomy. You learn some of the factors and footprints of hackers that will enable you to understand the threat that is present beyond the edge of your network.

Essentials First: Looking for a Target

The Internet has more than several billion possible public IP addresses, so how hard can it be to find a suitable target (also referred to as a *mark* or *subject*)? This is the first aspect of security on which people concentrate. Certainly your network's presence on the Internet is a way for hackers to find you; as a result, you should consider the security of your network from attackers and the value of anonymity. You might have purchased the best security technology to protect your PC, and you constantly ensure that it is up to date with the latest security patches. This includes your firewall, Internet router, VPNs, antivirus software, proxy server, biometrics, and all the best security technologies that money can buy. You have done this, right? Of course not, because these things are a pain to do and you believe that you have nothing anyone would want. We shall see....

It is natural to think that security technology can protect you from the malicious threats of hacker exploits. In this case, however, you might have been yearning for a sense of security but forgotten about the weakest security link: the human factor, which is what sits between the keyboard and the chair. It is this factor that thieves of any type count on; perhaps it's leaving your door unlocked, not patching your computer or antivirus/malware protection software, or believing you're safe behind your router or cable modem.

Consider for a moment whether your employees are trained in information and physical security. Would they know what to do if someone tried to fool them into giving away potentially sensitive information? How many sets of keys to the building exist? What are the cleaning people doing when you are not there? Are they disposing of your trash properly, or are they bagging and dropping it into the dumpster? Could an intruder break a window or pick a lock to enter your building undetected, or my favorite, how long have you had the same alarm PIN?

You might think that you have a great IT staff or even a team dedicated to network security, which is a good thing. Security professionals are expected to have a high level of technical competence and, for the most part, this is true. Now how does that awesome firewall completely protect you? What are the threats to the corporation from the inside behind those firewall controls, and what countermeasures do you have in place to protect your corporate assets today?

However, these same professionals often do not expect the same to be true of those attackers and intruders from whom they defend their sites. Many do not take heed of the

axiom that *“There’s always someone out there smarter, more knowledgeable, or better-equipped than you.”* Having engineers who think that they are the smartest people in the company is a recipe for disaster. Trust me, arrogance or a know-it-all attitude is a sure invitation to disaster and a magnet to those with something to prove. Segregation of duty is a very important concept ensuring that one employee does not have the complete keys to your kingdom.

Security is often simply an illusion facilitated and made more believable by the ignorance or naiveté of everyone in an organization. Do not place all your trust in security products; if you do, you settle for the illusion of security. Any security process must be implemented—that is, both technology *and* rules. (Specifically, all people in an organization must hold to these stated rules.) In addition, you must perform random and repeated audits to determine whether certain people in the company, such as the CEO who does not heed all the rules, bypass any rules or controls. The CEO or other senior executives usually have access to secrets and are the first target for a hacker. Letting the CEO bypass security policies, standards, and guidelines is a sure way to weaken a security policy.

In summary, true security is more than a product; it is a series of processes that encompass products and personnel across an organization—an end-to-end solution set that includes processes and controls with heavy policy governance. The following section covers the importance of having company personnel be aware of the security process.

Hacking Motivations

The introduction briefly touched on some of the confusion surrounding why hackers do the things they do. Although motivations are extremely diverse, there are some that are quite easy to identify. It is worth mentioning that several years ago these motivational categories did not exist, and as the Internet continues to evolve, so too will the hacker. The following list looks at some of the common motivators for hacking:

- **Human curiosity and fame:** In the early days of the Internet, hackers wrote viruses to see whether they could (and did) crash thousands of Windows PCs and gain global TV news coverage. It was also believed at that time hackers did so because they were curious or otherwise interested in technology. Certainly there are still many hacks occurring because people are curious or want recognition; however, this desire is shifting to the youth of the world, who get a charge out of hacking the cheerleader’s Facebook account. There are newer and more lucrative motivations driving the true threats, which have evolved past the script kiddie hacker today.
- **Anti-Establishment:** Hackers motivated by this category typically feel that the rules and regulations they are surrounded by do not or should not apply to them. You often hear of hackers striking out against a government or perhaps an employer. Oftentimes people on the inside of the target organization conduct threats motivated in this manner. One of the most recent examples was the Iranian presidential election of 2009; opposition parties whose freedoms had been restricted moved to the Internet; however, the authorities aggressively responded. This forced activists to “get creative” with getting the word out online and to media outlets outside the country’s borders.

- **Economic motivations:** There is an old saying that money makes the world go round, and there is an even older saying that says money is the root of all evil. In yet another example of the changing motivations behind hackers, many groups and individuals hack for cash. Certainly the most commonly known financial gains are through stealing credit card numbers or a person's identity. In the last several years, a new type of online hacker gang has emerged that is blackmailing businesses threatening to bring down their websites, thus impacting sales, if they do not pay money to ensure they are free from attacks for a year. Security organizations often report that financial reward is the largest reason why hackers keep coming back and upping their game.
- **Hactivism:** When you have a problem, and the police cannot help and the laws are silent, you might want to call the A-Team, but that is a different book. We are seeing a trend of hackers using their online skills to impact the real world based on their belief systems, thus Hactivism was born. There have been many examples of this sort of hacking motivation. Recently, the two most prevalent were when environmental researchers in England had their email server hacked. The resulting emails were shared far and wide, revealing some rather disturbing information that perhaps they altered climate change data to make it worse than it is. Another example is the recent denial-of-service attack against the World Trade Organization (WTO) website that coincided with street protests or the Wiki-Leak fans who targeted MasterCard. Although, it should be noted that one man's activism is another man's hate crime.
- **Cyberwarfare:** The newest and perhaps the most evolving motivation is Cyberwarfare. Simply put, this is using the Internet to conduct aggressive operations. One of the most recent attacks was when Chinese hackers (many suspect it was their government) tried to hack into Google, specifically Gmail users who were Chinese human rights activists. As expected, the Chinese government denied all involvement. Regardless, Cyberwarfare has come of age and is being used. Even earlier, in the Russian-Georgian war, there was a call to arms by Russia to its hackers who commenced to bring down the Georgian governmental website whose own hackers responded too. Although these are two of the most published examples, in security circles, what has been on page two is that many governments all over the globe are looking at making Cyberwarfare military units. The U.S. government sees Cyber as a newer domain to be treated equally like air, land, or water and to be protected just the same, keeping the nation's critical assets secure.

This section briefly looked at some of the more common hacker motivations and how we are seeing each of them in the world today. The next section deals with target selection by hackers.

Targets of Opportunity

I cannot keep track of the number of times I have been with customers who discuss their network and its security only to hear the following:

“We are a <Non-IT business> and there is nothing on our network that a hacker would want. Why should we be worried about making sure our network is secure?”

Wow! What a statement. It astounds me every time I hear it. There are many ways to reply to such a statement—some of which are politically correct, and some of which are not. Usually the person making this statement is a customer, so the focus here should be on the politically correct response.

This statement epitomizes an attitude known as Security Through Obscurity. In this book, you will see that when it comes to security, relying on obscurity is dangerous, regardless of the company's size or business, and it is rarely if ever effective. Just because you haven't been p0wned yet does not mean it won't happen to you or your corporation. Even if you have sophisticated monitoring, detection, and threat remediation tools and processes in place, how could you be sure the threats and exploits have not evolved past your current controls and countermeasures?

Perhaps the company in question might not be a financial institution, but its network certainly contains servers, hard drive space, bandwidth to the Internet, and personal employee information. Now, with the shift to private and public clouds, there could even be more of a challenge. Believing that this information is unimportant to a hacker can be fatal. An asset valuation and classification program is essential to categorize and identify what information your corporation has and associate an appropriate protection level. Consider what a hacker could do with such information:

- **Servers:** Hack a server, and you get a slave device that could potentially be used remotely to attack other, more important targets. Can you envision getting a call from men in dark suits that have no sense of humor regarding what your server might be doing? How does the shift to server virtualization and hypervisor or host change how you need to consider security controls? (I personally have assisted companies in ridding themselves of devices in their network that have become part of a botnet, which is using them for nefarious purposes.)
- **Hard drive space:** Every network has PCs with unused disk space. What if you were hacked and files of a questionable or perhaps even illegal nature were placed on them? Consider what the lawyers enforcing copyright laws or law enforcement might do if the files were to contain illegal types of pornography or terrorist material. In addition, most PC hard drives today are of the multihundred gigabyte variety or larger, the capacity of which is attractive to someone who needs to park a recently bootlegged movie or child pornography for a few hours or even days. If this happens in your network, would you know? If data were removed from these drives by a USB key, CD, DVD, or another method, how would you track data loss and have a viable digital or network forensic process in place to recover that data?
- **Bandwidth and bots:** A hacker can always use extra bandwidth and an alternative means of connecting to other companies to hack into them. If they gain access to your network, it is the PCs they want to control and make part of their botnet. A *botnet* is a collection of computers running malicious software (at Layer 4) enabling them to be controlled and used without the users' knowledge. Layer 4 botnet traffic visibility at the Web or firewall is critical to remediating these threats and visibility into the infected hosts in your network.

- **Personal employee information:** Armed with all the information an employer might need to verify employment and even pay its employees, a hacker could engage in identity theft. Consider the way in which corporate credit cards, Social Security numbers, addresses, and payroll information are stored—juicy information for a hacker.

These hacker activities could place IT personnel, management, or even the entire company in danger with legal or criminal ramifications, not to mention the bad press associated with being hacked to this degree. Consider a company's brand or reputation being destroyed and having to rebuild from there.

The more important question is not “Why (when) would someone hack us?” but “Am I vulnerable enough to be selected as a target?”

Targets of opportunity are clearly the easiest for a hacker to penetrate because something has happened, or not happened, that enables a hacker to easily identify and gain access to a corporate network that has nothing valuable, except all the PCs or virtual hosts.

Are You a Target of Opportunity?

In many cases, hackers prowl and crawl the Internet using a variety of tools (covered in Chapter 12, “Tools of the Trade”) and usually have an agenda in mind when they discover a potential target. In addition to hackers, there are a variety of individuals known as script kiddies.

Note A script kiddie (sometimes spelled “kiddy”) is a derogative term, originated by the more sophisticated hackers of computer security systems for the less skilled and not necessarily younger, but unfortunately often just as dangerous, exploiter of Internet security lapses. The typical script kiddie uses existing and frequently well-known and easy-to-find techniques and programs or scripts to search for and exploit weaknesses in other computers on the Internet—often randomly and with little regard or perhaps even understanding of the potentially harmful consequences. Hackers view script kiddies with alarm and contempt because they do nothing to advance the “art” of hacking, except sometimes unleashing the wrath of authority on the entire hacker community.

Although a hacker takes pride in the quality of an attack—leaving minimal to no trace of an intrusion, for example—a script kiddie might aim at quantity, seeing the number of attacks that can be mounted as a means of obtaining attention and notoriety. Script kiddies usually hack for the challenge and not for financial gain; although that can be a motivator. As novices, script kiddies often do not know what they are doing and can inadvertently cause a Denial of Service (DoS) attack. The word is that, in most cases, expert hackers were script kiddies at one time—makes sense because everyone has to start somewhere.

Determining whether you are a target of opportunity depends on your security infrastructure. A good rule is that if you do not have a firewall in place or your firewall has not been updated in a while, you are likely to be a target of opportunity. Because hackers employ

automated tools that look for vulnerabilities in your security, script kiddies are the most common threats to networks that are targets of opportunities. One of the easiest ways to ensure that you do not become a target of opportunity is to update your infrastructure (firewalls, IPS/IDS, secure routers, switches, servers, and PCs) with the latest patches. Do not get lulled into a false sense of security by patching only a server or two. A formal test and patch management process should be in place.

Remember that if you and your buddy are being chased by a hungry bear, you do not have to be faster than the bear, just faster than your buddy! You can easily protect yourself such that you might not be a target of opportunity because hackers will see easier targets elsewhere. If you are a target of a hacker, however, you're going to be thankful for taking action—hopefully, if you have not, this book can help you understand the importance of security.

Targets of Choice

Hackers often have a goal in mind when selecting a target. Consider the role the media has played in setting your internal vision of what a hacker is. Many people think that a hacker possesses the following characteristics:

- Disgruntled, negative, and angry at the world
- Bitter, with few friends and low self-esteem
- Extremely smart, yet not able to focus on making a living or having a career
- Has trouble maintaining relationships, friendships, or romance
- Disrespects authority; a social misfit, lone wolf
- Young and inept with women and others
- Enjoys junk food and pizza, ensuring the presence of acne (this was true, at least for me)

These stereotypes are true in some cases but not all; regardless, a subculture of hacking exists, and some hackers revel in it. However, believing that all the security threats against your network come from individuals like these would be a mistake.

Are You a Target of Choice?

The following scenarios can help you understand that your company—or perhaps even you—might be a target of choice by a hacker:

- Perhaps your company has a new product or solution that is going to revolutionize your area of business. What if it is a breakthrough?
- Perhaps you are engaged in a bitter dispute with a family member and you have information that the other party wants. A nasty divorce comes to mind as an example; your ex-wife might be going steady with a hacker checking your email and snail mail.

- Perhaps you have upset someone who knows a hacker.
- Perhaps you have a good credit rating or credit cards, making your identity very attractive—priming you for identity theft or botnet target.
- Perhaps your company is in a business that, if disrupted or left unavailable, would enable people with an agenda to make a point.
- Perhaps your company has information on another company that is important to someone such as a competitor, for example industrial sabotage.
- Perhaps an employee or former employee has become disgruntled and wants to make a point, which is often the case because most security threats come from employees.
- Perhaps you want to hide something from someone during a legal action.
- Perhaps your company is doing business in a part of the world that is in the middle of social or political upheaval—even hackers have geopolitical consciences nowadays.

In these cases and perhaps many others, you are now officially a target of choice because there is a reason why the hacker has chosen you.

Certainly the hacker could fit within the subculture described earlier, but perhaps he is not something out of a Hollywood movie. What about private investigators and lawyers—might they not be interested in information that you or your company might have?

As people wanting to know all sorts of things hire them, private investigators are learning new skills; therefore, to be successful, they could have turned to the Internet to find this information about you. What about the ex-military or those trained by the government as security specialists and business espionage? It is highly doubtful that they fit the Hollywood hacker stereotype. What about a spurned lover or spouse who has some computer skills, or an employee who knows all your partnering companies? These groups do not fit the hackers we see on Hollywood's silver screen, but they can certainly be viewed as a threat to your network.

Understand, as well, that a hacker might not do all the work himself, and it might not be electronic. For example, do you recall the term *dumpster diving*? Dumpster diving is legal and is an easy means of acquiring all kinds of information that could be helpful to a hacker because your trash is not your property anymore.

The following section covers how an attack begins and the process an attacker takes to begin compromising the target, which could be a person, software program, network, server, or the common Windows flaws. (Fortunately, this book was written and edited on a Mac.)

The Process of an Attack

An attacker can attempt to gain access to or exploit a system in many ways. This system can be as simple as a home computer connected to the Internet through a DSL connection, or a complex corporate network. Regardless of the kind of system an attacker targets, they typically employ the same fundamental steps:

1. Reconnaissance via social engineering or other methods
2. Footprinting/fingerprinting
3. Scanning (passive or active)
4. Enumeration
5. Gaining access
6. Escalating
7. Creating backdoors and covering tracks (cleanup)

The following sections discuss these steps in detail. You need to understand the concepts of what attackers might do in each step, and their goals, so you can detect and thwart their attacks.

Reconnaissance

Considering the introduction to this chapter, this discussion begins with hacking innocent information, which is also known as social engineering. Hacking innocent information from a person via social engineering is much easier than bypassing a firewall.

Fundamentally, people want to trust and help others, so they are more vulnerable to social engineering; combating this most basic hacking can be one of the biggest challenges to those who are responsible for security.

Although you might not think innocent information is worth protecting, it can be crucial to a social engineer attacker. When an attacker is armed with this information, he can use it to present himself as believable. In reality, this is where the hacker usually begins penetrating a company, by obtaining some document that might seem innocent and commonplace; be careful, however, because it could be useful to others.

Consider the following scenario, which I used once while performing a network assessment. To see what people would be willing to give up to someone who “sounded” official, I called the senior IT engineer, Daniel:

“Hello, this is Tom from WindWing Travel. Your tickets to San Jose are ready; would you like us to deliver them or arrange for you to pick them up as e-tickets at the airport?”

“San Jose?” Daniel says, “I do not have any travel plans there.”

“Is this Daniel Thomas?” I asked.

“Yes, but I do not have any trips scheduled until AppleCon in Las Vegas, later this year.”

“Well,” I chuckle, “are you sure you do not want to go check out San Jose?”

Daniel chuckles as well, responding to a humorous situation and a break in his normal routine by saying, “Sure, I’d be happy to go if you can convince my boss...”

“Sounds like another computer glitch,” I say and, while chuckling, I remark, “I thought computers were supposed to make our lives easier.”

Daniel laughs, too.

“In our travel system, we track travel arrangements under your employee number. Perhaps someone used the wrong number when booking the flight. What is your employee number?”

Daniel knows that several groups within his company have his employee number: security, human resources, his boss, and obviously finance, so why wouldn’t the travel company use a way to identify him that would fit with his company. There is no danger here, is there?

A competent hacker working on social engineering can take this simple piece of information and use it with some rather easily obtained data to take his hack to the next level. Imagine what access he might gain if he had an employee’s number, full name, telephone extension, department, work location, email address, and even his manager’s information. This information is innocent when viewed in pieces, but it paints a scary picture when compiled together.

Clearly, innocent information should be protected, and all employees should be made aware that mishandling information that should never be released to the public could truly endanger both the company and, more importantly, the employee. A strong security awareness program from all corporate employees and by a service-level agreement (SLA) for contractors should be tied and enforced by HR. For example, consider the following example:

“Daniel, I can’t find you by employee number. Let me try another way. What is your Social Security number?”

As you can see, a rapport was quickly established making my claims believable, and ultimately I got around to asking for the information I wanted, his Social Security number. A good rule is that all company data should be considered sensitive and not released unless an individual is explicitly authorized to do so. Remember that all calls and email are corporate property; with the move to IP convergence for voice and data, calls may be recorded for “quality purposes” and email may be archived and read later. The same applies to instant messaging communication. Security should be an enabler to the business, not a roadblock to progress.

Note For additional information on social engineering and how hackers gather information without ever alerting your network engineers, refer to the following enjoyable and well-written book, *The Art of Deception: Controlling the Human Element of Security*, by Kevin Mitnick and William Simon. This book also describes techniques and policies that you can use to defend against these types of attacks. I strongly recommend this enjoyable and well-written book.

Footprinting (aka Casing the Joint)

“Intelligence preparation” of the enemy and the battlefield is a military term used to define the methodology employed to reduce uncertainties concerning the enemy, environment, and terrain for all types of military operations. During military actions, this concept has been clearly demonstrated through the use of drone aircraft that enabled military commanders to see the battlefield and thus pick when and, more important, *how* they engaged the enemy. Understanding the battlefield and subsequently having the ability to choose *how* you engage the target is analogous to the choices hackers make. In network security terms, this intelligence preparation is known as reconnaissance and footprinting; in Hollywood movies, it is referred to as “casing the joint.” The network resources that security professionals are tasked with securing are analogous to a battlefield, and battlefield intelligence is critical to victory.

Hackers conduct these preparation operations against your company and network they need to understand “where” their target is and how it is put together. Footprinting is a continuous process used throughout all planned and executed operations. The myriad of attackers and intruders from the void are the aggressors who are constantly on the offense. The security professionals are the defenders, entrusted to preserve the confidentiality and integrity of data against these intruders and protect against disclosure, alteration, and destruction (DAD).

In the real world, many criminals perform this step, but they probably have not named it. For example, a criminal might review the security of a convenience store so that he can understand what the security is, where the money is kept, the location of security cameras, possible exits, and any other items that might help him succeed in his crime. As shown in Table 1-1, hackers look to gain information during this phase.

Table 1-1 *Goals of Reconnaissance and Footprinting*

Technology	What Is Learned
Your Internet Presence	<p data-bbox="390 301 1182 390">Ideally, a target would be connected to the Internet, and what network these days is not connected to the Internet? Attackers would therefore want to learn the following as they begin casing your network:</p> <ul data-bbox="390 399 1182 1349" style="list-style-type: none"> <li data-bbox="390 399 1182 453">• Information on individuals associated with the systems: name, phone number, position, address, what they know, and so forth. <li data-bbox="390 462 1182 516">• Develop any information that might make it easier to conduct social engineering. <li data-bbox="390 525 1182 614">• Where are these devices and systems physically located? You would be surprised what a simple traceroute can tell you about where your network is connected to the Internet. <li data-bbox="390 623 1182 677">• The target's domain names and DNS servers. Assigned blocks of public IP addresses. <li data-bbox="390 686 1182 740">• Which specific IP addresses (of those assigned) are accessible from the Internet? <li data-bbox="390 749 1182 802">• Of the IP addresses found to be accessible from the Internet, what services (www, FTP, email, and so on) are viable targets? <li data-bbox="390 811 1182 937">• Of the services found, what kind of computers—both hardware and operating system (including version/build so potential vulnerabilities can be known)—are they running on? For example, Windows, Linux, Sun, UNIX, and so on. Each of these has different vulnerabilities. <li data-bbox="390 946 1182 999">• Are there any mechanisms in place that control and track access to the network? <li data-bbox="390 1008 1182 1098">• What kinds of firewall, Intrusion Detection or Prevention Systems (IDS or IPS) are deployed to protect the target? Is there centralized logging and reporting with time sync to a Network Time Protocol (NTP) server? <li data-bbox="390 1107 1182 1232">• System enumeration allows for the specific identification of a system and some of the data available on it (user and group names, domain name, system banners, routing tables, and SNMP information are just a few examples). <li data-bbox="390 1241 1182 1295">• Network protocols (routed and routing) that are in use; for example, IP, OSPF, or BGP. <li data-bbox="390 1304 1182 1349">• Construct a simple network map with all the previous information, plus which company provides the target Internet access.
Intranet Characteristics	<p data-bbox="390 1385 1182 1501">Some network engineers understand that hackers try to gain access from the Internet; thus, many networks have duplicate infrastructure inside and outside their firewalls. As a result, thorough hackers repeat the footprinting steps they conducted from the Internet against the target's intranet.</p>

Table 1-1 *Goals of Reconnaissance and Footprinting*

Technology	What Is Learned
Remote Access Possible	<p>Many companies not only have normal Internet access through Frame Relay or broadband, but they also have dialup access. More commonly, dialup is going away for corporate backup and is being replaced with broadband or satellite, depending on the company's needs. This is yet another way for an attacker to enter the network, so a thorough hacker footprints these as well:</p> <ul style="list-style-type: none"> • What type of remote access is available and to whom? • Where does the remote access connect, and what is the connection's destination? • How is access to the network controlled? Are employees asked for a username and password or just a password (RADIUS, TACACS, and so on)? Is multifactor authentication possible or consider single sign on?

Impressive list, isn't it? The disturbing aspects of this list are twofold:

- Even the most inept hackers can figure these things out.
- Learning the answers to these questions is free and quite likely you will never know the threat until it's too late.

Hackers can take a lot of steps to learn about your network without your knowledge. Consider what simply looking at a Domain Name System (DNS) can reveal about your network through the use of a simple (and free) command known as **dig** (domain information groper), which has replaced **nslookup**. (See Example 1-1.)

Example 1-1 *Using DNS for Passive Reconnaissance via the dig Command*

```
Toms-iMac:~ ccie9360$ dig cisco.com any

; <<>> DiG 9.6.0-APPLE-P2 <<>> cisco.com any
;; global options: +cmd
;; Got answer:
;; ->HEADER<<- opcode: QUERY, status: NOERROR, id: 25065
;; flags: qr rd ra; QUERY: 1, ANSWER: 12, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;cisco.com.                IN      ANY

;; ANSWER SECTION:
cisco.com.                85877  IN     A      198.133.219.25
cisco.com.                86214  IN     MX     15 rtp-mx-01.cisco.com. <- MX =
    designates an email server
cisco.com.                86214  IN     MX     20 ams-inbound-a.cisco.com.
cisco.com.                86214  IN     MX     25 syd-inbound-a.cisco.com.
cisco.com.                86214  IN     MX     10 sj-inbound-a.cisco.com.
cisco.com.                86214  IN     MX     10 sj-inbound-b.cisco.com.
cisco.com.                86214  IN     MX     10 sj-inbound-c.cisco.com.
cisco.com.                86214  IN     MX     10 sj-inbound-d.cisco.com.
```

```

cisco.com.          86214  IN      MX      10 sj-inbound-e.cisco.com.
cisco.com.          86214  IN      MX      10 sj-inbound-f.cisco.com.
cisco.com.          86360  IN      NS      ns2.cisco.com. <- NS designates a
  DNS server
cisco.com.          86360  IN      NS      ns1.cisco.com.

;; Query time: 27 msec
;; SERVER: 208.67.222.222#53(208.67.222.222)
;; WHEN: Sat Jul 24 15:52:02 2010
;; MSG SIZE rcvd: 339

```

Note The `any` keyword asks for any DNS record. Many other more specific options are available. You can run the `man dig` command for the manual, which is the UNIX/Linux way to reference a command's manual.

As you can see, the output reveals Cisco.com's email, DNS, and web servers, which you can then use to determine more information. Now `traceroute` is used to the DNS A record for the domain to figure out where this domain is located on earth:

```

Toms-iMac:~ ccie9360$ traceroute 198.133.219.25
traceroute to 198.133.219.25 (198.133.219.25), 64 hops max, 52 byte packets
 1 172.16.17.1 (172.16.17.1) 1.957 ms 2.033 ms 1.090 ms
  <<<Output omitted for my security!>>>
 7 te-1-1-0-4-cr01.dallas.tx.ibone.comcast.net (68.86.91.253) 23.208 ms 28.560
  ms 34.533 ms
 8 pos-0-3-0-0-pe01.1950stemmons.tx.ibone.comcast.net (68.86.86.154) 30.395 ms
 26.313 ms 25.119 ms
 9 as7018-pe01.1950stemmons.tx.ibone.comcast.net (75.149.231.22) 23.717 ms
 29.184 ms 23.887 ms
10 cr2.dlstx.ip.att.net (12.122.195.242) 69.468 ms 70.911 ms 69.258 ms
11 cr2.la2ca.ip.att.net (12.122.28.178) 70.152 ms 72.476 ms 77.454 ms
12 cr2.sffca.ip.att.net (12.122.31.134) 83.423 ms 72.581 ms 70.189 ms
13 gar8.sffca.ip.att.net (12.122.114.65) 69.502 ms 70.041 ms 76.617 ms
14 12.91.205.14 (12.91.205.14) 71.215 ms 69.936 ms 69.898 ms
15 sjc5-dmzbb-gw1-ten4-5.cisco.com (128.107.224.250) 69.791 ms 74.528 ms
 69.920 ms
16 sjce-dmzbb-gw1-ten3-3.cisco.com (128.107.224.2) 69.690 ms 71.246 ms 71.310 ms
17 sjck-dmzdc-gw1-gig5-2.cisco.com (128.107.224.69) 69.856 ms 70.424 ms 70.198 ms
18 * * *
19 * * *
20 * * *
^C

```

The `traceroute` output shows that this domain lives on the AT&T Internet backbone because it owns the 12.x.x.x class A address range. Because hops 15+ are DMZs, it is likely

they are firewalls of some sort with descriptive PUBLIC DNS names. The first three letters are *sjc*, which must be San Jose California, which you can learn at www.cisco.com is where Cisco is headquartered. This also means that hop 14 must be a router because it is connecting to the Internet on behalf of the firewalls. Also somewhat interesting is hops prior to, consider that Comcast is plainly telling you where its routers are located. AT&T, on the other hand, is abbreviating cities, but its naming convention is apparent: three letters for the city and two letters for the state, thus *dlstx* is Dallas, TX, which then goes to *la2ca*, Los Angeles, California. You can learn a lot with what is available freely online, definitely compiling some good intel as you case the joint!

Now determine whether all the servers that **dig** reported to you are in the same location. First ping the Cisco DNS server to get its IP address, and then you can traceroute to it:

```
Toms-iMac:~ ccie9360$ ping ns1.cisco.com
PING ns1.cisco.com (128.107.241.185): 56 data bytes
64 bytes from 128.107.241.185: icmp_seq=0 ttl=111 time=71.388 ms
64 bytes from 128.107.241.185: icmp_seq=1 ttl=111 time=71.445 ms
64 bytes from 128.107.241.185: icmp_seq=2 ttl=111 time=70.924 ms
^C
--- ns1.cisco.com ping statistics ---
3 packets transmitted, 3 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 70.924/71.252/71.445/0.233 ms
Toms-iMac:~ ccie9360$

Toms-iMac:~ ccie9360$ traceroute 128.107.241.185
traceroute to 128.107.241.185 (128.107.241.185), 64 hops max, 52 byte packets
 1 172.16.17.1 (172.16.17.1) 1.584 ms 1.403 ms 1.052 ms
<<<Output omitted for my security!>>>
 7 te-1-1-0-4-cr01.dallas.tx.ibone.comcast.net (68.86.91.253) 23.504 ms 25.326
   ms 23.431 ms
 8 pos-0-2-0-0-pe01.1950stemmons.tx.ibone.comcast.net (68.86.86.150) 24.266 ms
   28.507 ms 28.004 ms
 9 as7018-pe01.1950stemmons.tx.ibone.comcast.net (75.149.230.162) 22.479 ms
   28.274 ms 23.051 ms
10 cr2.dlstx.ip.att.net (12.122.195.242) 71.920 ms 69.378 ms 69.901 ms
11 cr2.la2ca.ip.att.net (12.122.28.178) 69.274 ms 69.825 ms 76.966 ms
12 cr2.sffca.ip.att.net (12.122.31.134) 70.084 ms 69.025 ms 69.719 ms
13 gar8.sffca.ip.att.net (12.122.114.65) 69.148 ms 69.967 ms 68.863 ms
14 12.91.205.14 (12.91.205.14) 69.637 ms 72.182 ms 69.718 ms
15 sjc5-dmzbb-gw1-ten4-5.cisco.com (128.107.224.250) 71.831 ms 71.324 ms
   69.001 ms
16 sjce-dmzbb-gw1-ten3-3.cisco.com (128.107.224.2) 70.783 ms 73.464 ms
   71.846 ms
17 sjck-dmzdc-gw1-gig5-2.cisco.com (128.107.224.69) 71.295 ms 71.022 ms
   69.650 ms
18 * * *
^C
Toms-iMac:~ ccie9360$
```


The **traceroute** results are identical, so now you know that San Jose holds many important servers for Cisco and that they have at least three different ranges of public IP addresses, specifically the following:

- 198.133.219.25 = DNS A record
- 128.107.224.69 = DMZ firewalls
- 128.107.241.185 = DNS server

All these subnets are in San Jose, CA, the location of the headquarters campus of the target. I know this through confirmation at its website and an Internet search.

Consider that through using just simple and free DNS tools, the attacker can reveal the public IP address of the Cisco website and that of its DNS and email servers. Cisco.com is queried all the time and is, therefore, not alarmed by *passive reconnaissance*.

Whois is a tool that is again freely available in many applications and on the Internet at the following locations. Try it out on your domain:

- www.networksolutions.com/: Whois web interface
- www.arin.net/: ARIN Whois
- <http://whois.ripe.net/>: European Whois
- <http://whois.apnic.net/>: Asia Pacific IP address allocations
- <http://whois.nic.mil/>: U.S. military
- <http://whois.nic.gov/>: U.S. government

Do not forget the information available on a company's website and how useful it is to know the address, main phone number, fax number, mergers, press releases, and members (with bios) of the company's management team. A target's corporate website has become a well of useful information from which an attacker can learn quite a lot. The hacker could use this knowledge for social engineering, identity of network systems, system administrators, and so forth. Usenet and Web searches on the system administrators and technical contacts are found when running host queries. By taking the time to track down this information, the attacker might be able to gain greater insight into the target network, plus the attacker knows how to do some social engineering, as key names become known.

If you were a hacker, however, you could begin a more active reconnaissance process to determine what services you could see on these servers through their public IP address.

Unfortunately, most companies are not prepared to detect these types of scans or probes. It is not that they do not have some of the necessary tools; it is simply that the target devices most likely are not logging what is going on—or if they are, no one is looking at the logs. This is one of the problems facing security professionals these days—information overload from device logs!

Consider that one of the next steps could be a simple ping scan using any number of freely available tools on the Internet. Figure 1-1 shows one of the best FREE tools available for black- and white-hat hackers, NMAP. Specifically I used its front end GUI, known as Zenmap, to run a simple scan of a test subnet, 172.16.17.0 /24.

The output of this command also includes some interesting and helpful information. From the results in Figure 1-1, you now know that three known devices on the network exist: Cisco, Apple, and HP. Unfortunately, there are not any Windows machines, so this is going to be a bit tougher; perhaps you should try a different network?

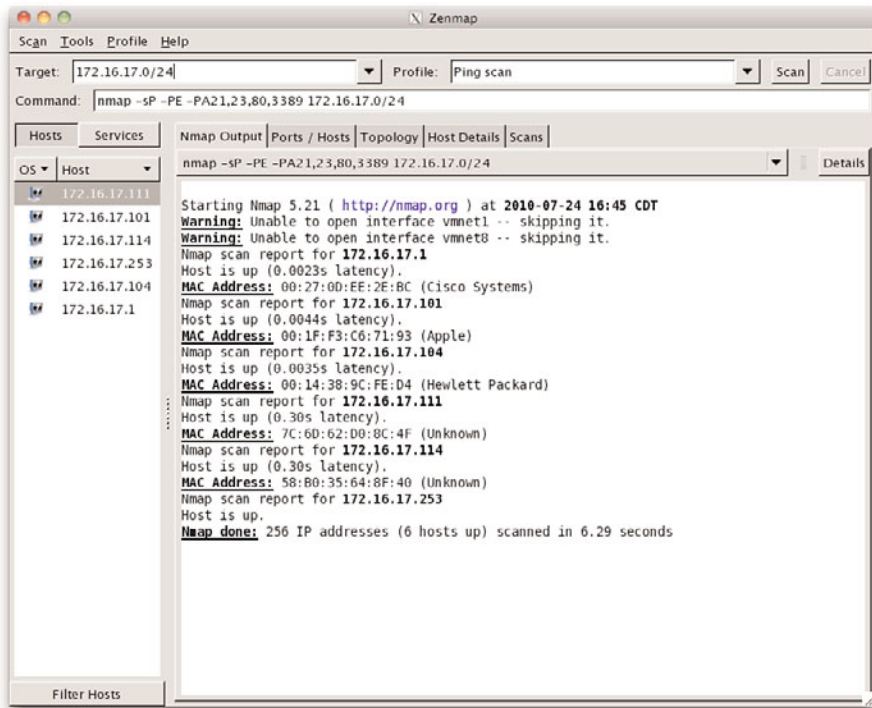


Figure 1-1 Ping Scan of a Class C Subnet

During this phase of an attack, the methods employed involve nonintrusive and standoff methods that hopefully do not enable the attacker's efforts to be detected. The attacker wants to determine the type of network with which he is dealing, and with whom he is dealing: system, network, and security administrators. Attackers know this and understand that the more active they are, the more likely their activities will be noticed. Attackers can, therefore, start active reconnaissance and allow it to continue until they learn enough information to launch an exploit against that system. If the exploit succeeds, the attackers move on to the next step; if not, they go back and gather more information.

Again, the intent is to develop a network map that uses information gathered during footprinting; the hackers figure out which devices are routers, firewalls, PCs, printers, and so on, and place them on the map, and identify key systems such as mail servers, domain name servers, file servers, and so on. The attackers also want to know where the target gets its Internet access in case they need to try to access the target through its ISP.

Scanning

At this point, the attackers have a good map of the machines on the network, their operating systems, who the system administrators are, any discussions posted to newsgroups, their office locations, and who their upstream intrusion prevention system (IPS) or IDS is. The attackers also know that, from this point forward, everything they do might be logged; at a minimum, they should assume that it is. The attackers have a map of the network and devices and are ready to move on to identifying listening services and open ports. The attackers also determine the acceptable risk. Can they afford to be logged during scanning? Are they behind a series of proxies outside the United States? Is compromise acceptable during the latter stages of the attack? Is concealment of the originating attack location necessary? What about exposure of the sponsor if he is working on behalf of another entity? There is a lot going on in the attacker's mind. Some attackers sketch things out, and others internalize these considerations as they move from step to step.

In Example 1-2, the attacker has initiated a more active set of scans against a target using NMAP (www.insecure.org), a free tool that both hackers and ethical hackers (good guys) commonly use. Because it is free, you will not find a script kiddie without it! Or better yet, Backtrack 4 (BT4) and its open-source complete compilation of wired, wireless, and forensic tools. Check out the software at www.backtrack-linux.org/.

Example 1-2 *Active Port Scan Results*

```
[AppleKick:/Users/topkick] topkick# nmap -sS -O 192.168.254.69

Starting nmap V. 3.00 ( www.insecure.org/nmap/ )
Interesting ports on (192.168.254.69):
(The 1579 ports scanned but not shown below are in state: closed)
Port      State      Service
7/tcp     open      echo
9/tcp     open      discard
13/tcp    open      daytime
17/tcp    open      qotd
19/tcp    open      chargen
25/tcp    open      smtp
42/tcp    open      nameserver
53/tcp    open      domain
80/tcp    open      http
119/tcp   open      nntp
135/tcp   open      loc-srv
139/tcp   open      netbios-ssn
443/tcp   open      https
445/tcp   open      microsoft-ds
563/tcp   open      snews
1025/tcp  open      NFS-or-IIS
1027/tcp  open      IIS
```

```

1031/tcp  open      iad2
1033/tcp  open      netinfo
3372/tcp  open      msdtc
3389/tcp  open      ms-term-serv
Remote operating system guess: Windows 2000/XP/ME

Nmap run completed — 1 IP address (1 host up) scanned in 5 seconds
[AppleKick:/Users/topkick] topkick#

```

If you refer to Figure 1-1, you can see that the ping scan revealed an active host at 192.168.254.69 and that the more detailed scan with NMAP, shown in Example 1-2, provided additional information. You might be wondering how accurate NMAP is. The answer is *very accurate*; specifically, the device scanned was, in fact, a Windows 2000 server.

For example, Figure 1-2 uses TigerSuite (www.tigertools.net) to see the other services that are accessible on the server. The figure shows that the server is readily identified, as are some extraneous services that are easily exploitable: SMTP, NTP, and FTP. Also notice that the scan revealed the server's name and domain, which is helpful to users and hackers in a Windows network! Not to mention the highly useable and open NetBIOS ports, typically 135–139 on Windows-based machines.

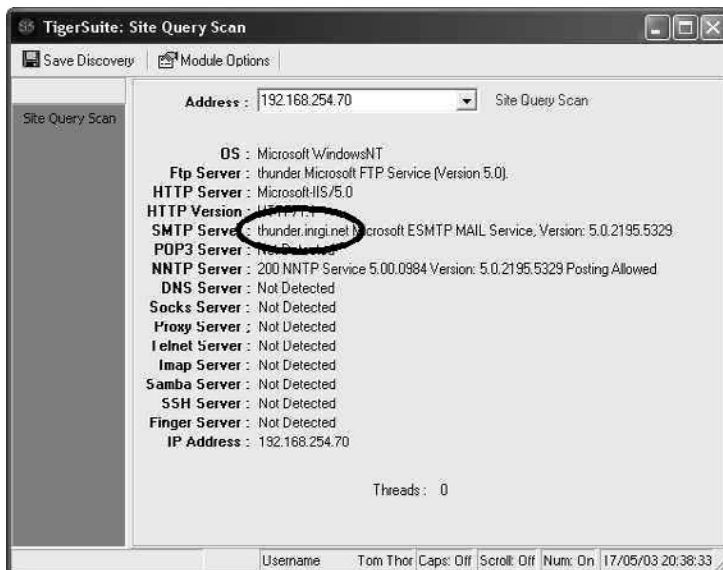


Figure 1-2 Server Query Scan

Being somewhat concerned about these services, I immediately shut them off and disabled them from starting again; they can also be audited and turned up and secured when

applications call critical services. This is a relatively new server, and I was more concerned about getting it functioning for my users than securing it. In their hearts, most OS vendors felt the need to be helpful and reduce the number of expensive (for them) technical support phone calls turning on every service and function from the beginning. Clearly, they were not thinking of security in this decision—only money. Apparently, they want their CEO to be the richest man in the world. Regardless of their irresponsible motives, Figure 1-3 shows that it does not take the IT professional long to correct this situation. As you can see by the circled server name, the scan has revealed whom they are and what the hacker found.




Figure 1-3 Secured Server Scan Results

You can Telnet to an open port 80 and do a simple `get` command, but SSH (22) should always be the default out of band (OOB) management option. The result should be a “banner,” which identifies web server type (IIS, Apache, and so on) and other interesting facts, as shown in Example 1-3.

Example 1-3 Telnet to Mail Server, Doing Some Reconnaissance

```
[AppleKick:~] topkick% telnet 192.168.254.69 80
Trying 192.168.254.69...
Connected to 192.168.254.69.
Escape character is '^]'.

get
```

```

HTTP/1.1 400 Bad Request
Server: Microsoft-IIS/5.0
Date: Tue, 20 May 2011 00:43:14 GMT
Content-Type: text/html
Content-Length: 87

<html><head><title>Error</title></head><body>The parameter is incorrect.
</body></html>Connection closed by foreign host.
[AppleKick::~] topkick%

```

The first thing to do is to open a connection from your computer to your mail server. For some odd reason, newer versions of Windows no longer have Telnet enabled by default (annoying Microsoft), so you need to manually enable it. Also, some email providers use content-aware firewalls to block Telnet, Cisco, and Hotmail, for example. In this example, we are going to find an executive's email address because we are going to fool the email server that we are trying to send a real email:

```
telnet mail.domain.ext 25
```

You should receive a reply like this from the email server:

```

Trying <mail.servers.public.ip.address>Connected to <mail.domain.ext>.Escape
character is '^]'.
220 mail.domain.ext Microsoft ESMTPL MAIL Service, Version: 6.0.3790.4675 ready at
Sun, 25 Jul 2010 15:31:21 -0400

```

You then need to tell the email server where you are sending the email from:

```
HELO www.hackerdomain.com
```

In this case, we are making up a domain because we do not want to tell anyone who we are or where we are coming from. In the real world, although you really should use your exact fully qualified domain name as seen by the outside world, the mail server has no choice but to take your word for it because those are the rules defined in RFC822 through RFC1123. Also, some email servers accept EHLO in place of HELO. The email server should reply as follows:

```
250 mail.domain.ext Hello [YOUR.ip.address], pleased to meet you!
```

Now give your email address; on many mail servers the space after the : is required rather than optional.

```
MAIL FROM: naughthacker@checkingyouout.com
```

The email server should reply as follows:

```
250 2.1.0 naughthacker@checkingyouout.com...Sender OK
```

Now give the recipient's address:

```
RCPT TO: mail@otherdomain.ext
```

When you discover the valid and real email address of the CEO, the email server responds as follows:

```
250 2.1.0 mail@otherdomain.ext... Recipient ok
```

To start composing your message, issue this command:

```
DATA
```

The email server should reply as follows:

```
354 Start mail input; end with <CRLF>.<CRLF>
```

If you want a subject for your email, type the following:

```
Subject: <type.subject.here>
```

Then press Enter twice (these are needed to conform to RFC 882); you will not see this reflected on the screen but do it anyway. You can now proceed to type the body of your message:

```
Hello CEO, there is some really important invoices for you to review please click here to ensure they are correct. You are being whaled so will need all account numbers for us to hack, thank you.
```

Note After phishing comes *whaling*, a sneaky attempt by scammers to hijack the personal computers of top-ranking business execs.

To tell the mail server that you have completed the message, press Enter a couple times; then enter a single “.” on a line on its own and press Enter. The mail server should reply with the following:

```
250 2.0.0 ???????? Message accepted for delivery
```

You can close the connection by issuing the following command:

```
QUIT
```

The mailserver should reply with something like:

```
220 2.0.0 <mail.domain.ext> Service closing transmission channel
Connection closed by foreign host.
```

Note Telnetting to target IP addresses on various port numbers can sometimes yield surprising results. Try Telnetting to some of the more commonly known ports (such as port 80) to see what kind of results you get. You will be surprised at the information readily available to you.

Another type of scanning, known as vulnerability scanning, is typically done from the Internet to find out how well a system is protected. This type of scanning and some of the

tools available are discussed in Chapter 12; until then, you can search online to find many scanners, and yes, many are free! As you have seen, each technique in the reconnaissance phase has value, but the true value (for the attacker) is gained when multiple techniques are combined to gain a complete picture of the target network or device. Nessus and Kismet are good examples of open-source tools in BackTrack4 that enable through vulnerability scanning.

As you move on to the next step in the attack process, remember that the scanning of a target allows the attacker to focus her efforts and attention on the most promising avenues of entry into your network. Attackers expect your IT professionals to be watching, but they doubt they will be seen; that assumption is subject to change, however, if you get involved and make a difference.

Enumeration

Defining the network environment involves footprinting, scanning, and enumeration as the hacker learns and prepares for the attack. Footprinting enables the attackers to limit the scope of their activities to those systems that are potentially the most promising targets to vulnerabilities they plan on running against the server. Scanning told the attackers what ports are open and what services are running.

Enumeration is the extraction of valid account information and exported resources. The key difference between the preceding scanning and footprinting techniques is that enumeration involves active connections to specific systems and directed requests to connect to these specific systems. Enumeration is succinctly defined at Wikipedia as follows:

Network enumerating is a computing activity in which user names, and info on groups, shares and services of networked computers are retrieved. It should not be confused with Network mapping, which only retrieves information about which servers are connected to a specific network and what operating system is run on them.

Note The previous section concluded by saying that attackers expect to be seen but ignored while they are footprinting. However, when enumeration begins, the attackers' attempts must be stopped or, at the minimum, logged and *acted upon!*

Like all steps in the attack, pulling the results together makes the difference in the success of the attack. Following are the four main categories within a network:

- Network resources and open shares
- Users and groups
- Applications
- Device logon banners and message of the day (MOTD) on network devices

As you can tell, the presence of each of these categories differs on every operating system. Consider that every major operating system enables shares, but each—Mac OS, Windows, Linux, and Novell—handles them in a different way. This means that, from an attacker’s perspective, each operating system must be handled differently.

The earlier example of a layered approach becomes apparent here because, through the use of NMAP, you have a good idea of what operating systems you are trying to enumerate and thus ultimately attack.

Enumerating Windows

As the industry leader in computer operating systems, Microsoft Windows is perhaps the most widely discussed; therefore, it makes sense to spend some time on it first. Windows operating systems still depend heavily on the use of NetBIOS (UDP Port 137), and many of the tools an attacker might use to learn more about a Windows-based network are built in to the operating system itself, as you see in the text that follows.

Example 1-4 shows the results of issuing a **net view** command from the command line of a Windows machine. In this case, the domain was known, so including it in the command revealed all the machines in that domain. Had the domain been omitted, all the LAN’s domains would have been displayed.

Example 1-4 *Using Windows Net View*

```
C:\>net view

Server Name          Remark
-----
\\APPLEKICK          Toms MAC
\\LIGHTNING
\\THUNDER
\\TOPKICK

The command completed successfully.

C:\>
```

This enumeration technique is even more useful when you combine it with the results of the earlier ping scan. You can use IP addresses and NetBIOS names interchangeably; so for example with NetBIOS, you might access another computer using \\COMPUTER-NAME. You can also use \\192.168.254.69. Attackers know this and modify their systems so that their machines “automatically” cache the NetBIOS names.

Another great built-in Windows tool is **nbtstat**, which enables you to query another computer for its NetBIOS name table. This means an attacker can query a server for its table, as shown in Example 1-5.

Example 1-5 *Query via nbtstat*

```
C:\>nbtstat -A 192.168.254.70

Local Area Connection:

Node IpAddress: [192.168.254.69] Scope Id: []

                NetBIOS Remote Machine Name Table

    Name                Type            Status
    -----
    THUNDER              <00>    UNIQUE        Registered
    THUNDER              <20>    UNIQUE        Registered
    INRGI                 <00>    GROUP         Registered
    INet-Services        <1C>    GROUP         Registered
    IS-THUNDER.....    <00>    UNIQUE        Registered
    INRGI                 <1E>    GROUP         Registered
    INRGI                 <1D>    UNIQUE        Registered
    .._MSBROWSE_.      <01>    GROUP         Registered
    MAC Address = 00-C0-9F-20-E4-F0

C:\>
```

In addition, if you do not know the IP address of the machine you have your sights set on, you can issue the command `nbtstat -c`; you are then provided with a listing of the NetBIOS names (in your cache) and their corresponding IP addresses, as demonstrated in Example 1-6. Don't you just love friendly operating systems that are eager to freely tell you about themselves?

Example 1-6 *Using nbtstat -c to Display NetBIOS Names*

```
E:\>nbtstat -c

Local Area Connection:

Node IpAddress: [192.168.1.101] Scope Id: []

                NetBIOS Remote Cache Name Table

    Name                Type    Host            Address            Life [sec]
    -----
    PRO200              <20>    UNIQUE         192.168.1.100    587
    PRO200              <00>    UNIQUE         192.168.1.100    95
```

The best way to stop an attacker from learning this kind of information from your network is to ensure that your router and firewall are blocking the entry *and* exit of NetBIOS

packets. Block at both points to prevent a layered approach to security. Specifically, block the following:

- TCP and UDP on ports 135 through 139
- TCP and UDP 445 for Windows 2000

Blocking these ports does not stop NetBIOS; it simply prevents it from entering your network. Ways exist to disable NetBIOS on a Windows PC; however, this might not be an option. Table 1-2 shows some of the common tasks and tools that attackers use.

As discussed at the beginning of this section, each operating system has associated techniques that enumerate against it. You have looked at a couple techniques that are just for Windows, and there are many more. Later in this chapter, you will see some recommended titles that discuss more about the other enumeration possibilities.

Gaining Access

Many people mistakenly believe that an attacker wants to take control of a target device and that is the ultimate goal of an attack. This is not entirely true. What is more likely is that an attacker wants to “only” gain access to a target PC. After enumeration identifies promising avenues of entry, more intrusive probing can begin as valid user accounts and poorly protected resource shares are exploited to gain access.

Ultimately, attackers must gain access to a system through some aspect of that system. There are typically four major types of exploits that reflect different aspects of a system that attackers target:

- Operating system attacks
- Application attacks
- Misconfiguration attacks
- Script attacks
- Broader DoS or DDoS attacks that might include the preceding attacks

Table 1-2 *Attacker Tasks, Tools, and Techniques*

Attacker Tasks	Tools and Techniques
List file shares	Being onsite
List usernames	NetBIOS and NetBUI
Identify applications	Using Telnet to see default banners
Identify operating systems	Null sessions

Within these different aspects of an attack, an attacker can proceed in two ways:

- **Automated attacks:** These types of attacks target one or more aspects of the target and are usually opportunistic by design. Automated attacks are opportunistic in the sense that they scan an entire block of IP addresses to look for vulnerability. For example, an automated attack might scan every IP address in a Class C block on port 80 looking for a known vulnerability that affects web servers. If the scan is successful, the attack proceeds; if not, the scan continues looking.
- **Targeted attacks:** These types of attacks might be more dangerous than automated attacks because your organization has been singled out for an attack; a good example is a government organization—local, state, or federal. In other words, attackers know that you have something they want, or that by succeeding in their attack on you, they can achieve a goal. Increasingly, the latter force drives attacks by using politics or social agenda as a rationale for an attack. Fortunately, targeted attacks seem to make up the minority of Internet activity. However, the bad news is that if you are targeted, the more skilled the attacker, the less likely you are to “see” or detect the attack.

Remember these two ways an attack might occur as you consider how an attack can affect the different aspects of a system.

Operating System Attacks

An operating system is designed to support what a user would like to accomplish and, in the context of this discussion, the operating system must enable networking to some degree. The more networking enabled on a system, the more services are activated to support these needs. This results in more open ports and active services being available and visible. Therefore, attackers have more opportunities to select an attack, thereby resulting in the access they want. Remember, other than financial motivation, one of the largest differences between the black and white hats is time. The black hats have all the time in the world to get their tasks at hand accomplished—you and I do not.

In addition, users and administrators often think that the job is finished when a server has its OS installed and its services configured. Alas, this is a mistake that results in a perfect target for attackers. Consider being a hacker and finding a server that has the original operating system installed without patches and with all default services activated. That server will be compromised within the hour! There should be a corporate template for OS hardening supported by an application security policy with an exact procedure for the team to follow to keep this consistent and up to date.

Application Attacks

I once worked in a business unit that wrote networking software for one of its products. The company was a large, international company with a strong history in telecommunications. I explain this background because with all the software being written these days, you would think that this company would take advantage of its understanding of the technology and security.

Alas, that was not the case; software programmers were under amazingly tight deadlines and were always asked for new features. I knew many of them—they inherently wanted to do the right thing, but outside factors drove their activities in many ways. Essentially, software was not being tested as it should have been. Add in its increasing level of functionality, and you have opportunities for attackers. This is all terrible, but consumers did not care about security several years ago—only whether the software had the features they wanted. Perhaps if consumers change what they spend money on, secure software will become more of a pressing issue.

Misconfiguration Attacks

Sometimes, system administrators work on the system when trying to secure a system or ensure that it provides the functionality users need. Usually, this means turning on several options, and the wanted feature starts working when you press the right option.

Did you clean up those options after yourself? Likely not. The problem is that the system administrator does not go back and research what fixed her issue and deactivate the unneeded options. This is perplexing because verifying that a system is not misconfigured is an easy precaution to ensure that your system is functioning correctly. A good rule of thumb is to turn unnecessary services off and concentrate on correctly securing and configuring those that are needed.

Keep a written record of what services and options you enable or disable; in the heat of the moment (especially when it is 3:00 a.m. and you are wondering what you did to deserve being hacked), the written record can help you reverse what you might have done earlier.

Another issue that fits under the misconfiguration umbrella is deploying a device and not changing the default administrator username and password programmed into the device. If you are wondering what I am referring to, look at the manual that came with your shiny, new firewall device that has all the blinky lights and whiz-bang security features. Have you looked at the “quick startup” section that almost all manuals have nowadays? Somewhere among those pages is a section about logging in for the first time and setting up the device. Most security devices either have no password, or the username/password combination is something such as “admin/admin.” Guess what? Hackers read manuals, too, and they are aware that default passwords are still active on routers, firewalls, and other Internet devices.

Note Your next step is to consider device configuration templates, automated deployment tools with role based access control (RBAC), and staging before deployment should be configured to reduce misconfiguration exposure for a corporation.

Scripted Attacks

UNIX, BSD, and Linux are undoubtedly the systems for which attackers will find scripts susceptible to their activities. Many of these operating systems come with sample scripts and programs available for use. These are a blessing in disguise and, if left activated or unchecked, they can result in successful attacks against your system.

Attackers try to execute some of the following attacks against your system during this phase of the attack:

- **Buffer memory overflows:** The information has to go somewhere, and the attacker can direct it to compromise a system. When the Buffer In Question (BIQ) blows, the OS might do things that the developer never intended.
- **Brute force guess passwords:** The attacker starts a program that tries every word in a dictionary. Webster's is fine, but it could also be a dictionary of names, movies, or sports teams/lingo, and so on. Tools such as John the Ripper are effective for dictionary and brute force attacks on passwords.
- **Try and sniff a password:** Everyone has to log in, and if the attacker can "see" a user's password, he is in! Can you imagine the number of captured passwords that could be seen in the morning when everyone is logging in to your network? Many passwords are plain and clear text to be uncovered with an open-source sniffer.
- **Capture the password:** In this case, the attacker wants to capture the password file, which can then be decrypted and cracked at the attacker's leisure and most likely *not* on the system that was compromised. In other words, the attacker copies this file and cracks it at his leisure (that is, sleeping or at his day job) so that the information it contains is useful.

The techniques, tools, and procedures vary according to the attacker's level of expertise and ability to code custom scripts and programs. Either way, a plethora of free open-source tools is available for use; the attacker will more than likely make use of some, if not all the following:

- **NMAP:** www.insecure.org/
- **NESSUS:** www.nessus.org
- **SATAN:** www.porcupine.org/satan/
- **ETTERCAP:** <http://ettercap.sourceforge.net/>

Do not discount the fact that commercial products are also available and can be used by hackers as well:

- **GFI LANGuard:** www.gfi.com/
- **SAINT:** www.saintcorporation.com/

- **METASPLOIT PRO:** www.rapid7.com
- **CORE IMPACT:** www.coresecurity.com/

The following section discusses how attackers work on escalating how much they are allowed to do (that is, privilege) after accessing a system.

Escalating Privilege

At this point in a hack, attackers might have gained access to a system. Perhaps the attackers learned/guessed/hacked a user's password because it was something simple, like the user's favorite sports team or movie. A regular user, however, might not have the privileges the attackers need for their goal. Thus, in this phase of the hack, the attackers must begin escalating their privilege level. They now understand the system a bit more, so they likely look for the following:

- Being "in" the system, the attackers can run the appropriate exploit code against the system to gain more privileges.
- Try to crack passwords using the many freely available password crack tools.
- Look for passwords that are not encrypted (that is, clear text).
- Evaluate the trusts that exist between the hacked system and others within the network. Perhaps there is another opportunity?
- Perhaps file or share permissions are incorrectly set.

These are the types of steps attackers take after they have gained rudimentary access. They would not likely go through all the risk and trouble to stop without ensuring that they can do whatever it is they intend to do.

If all else fails, or if the attackers want to implement a denial of service (DoS) attack, they use specialized tools called *exploit code* to disable a system. The use of these exploits is operating system-specific and can also depend on the patch level of the system state. Specifically, this means that system X is vulnerable to exploit 666, but if it has been patched with service patch 5, it is not vulnerable. Some of the exploits that could be used are SYN flood, ICMP techniques, overlapping fragments/offset bugs, and out of buffer. Again, the effectiveness largely depends on the system's patch level. The attackers knows that when an exploit becomes public it can quickly become useless against systems where the system administrators stay on top of things; however, attackers also know that new exploits are found daily, and that research and experimentation are required to find the most effective tools and techniques.

The remaining steps are rather straightforward and obvious. After the hackers gain administrator/root access (that is, ownership), they complete the reason behind the attack, begin concealing their activities, and almost always leave a way for the hackers to get back into the system now that it is compromised.

Covering Tracks

After the attackers accomplish ownership of the target system, they must hide this fact from the system administrator and blend in if they intend to stay a while and siphon off information. This is one of the most fundamental rules of hacking; however, it is also one of the hardest for attackers to accomplish. For Windows-based systems, event log and registry entries are cleared/cleaned. For a UNIX-based system, attackers clear the history file and execute a log wiper to clean entries from UTMP, WTMP, and Lastlog.

Note The attackers clear the logs—not delete them. When log files are deleted or cleared, a notification occurs that might draw attention to the fact that the system was compromised.

If the attackers want to maintain access to the system after achieving initial access, they create backdoors for future access. The methodology, tools, and techniques are system-dependent, but the intent is to create accounts, schedule batch/cron jobs, infect startup files, enable remote control services/software, and replace legitimate applications and services with Trojans. Possible tools include the following:

- **Netcat:** A simple UNIX utility that uses TCP or UDP protocol to read and write data across network connections (<http://nmap.org/ncat/>).
- **VNC (Virtual Network Computing):** A remote display system that enables you to view a system's desktop environment—not only on the machine where it is running but also from anywhere on the Internet and from a wide variety of machine architectures. Many programs do this—VNC just happens to be free and rather popular. Plus, it works on Windows, Linux, and UNIX (www.uk.research.att.com/vnc/).
- **Keystroke loggers:** Hundreds are available on the Internet, and they can be either hardware- or software-based. Keystroke loggers record every keystroke pressed for a computer and can even email you what they record—anywhere—including all your or your company's banking transactions.
- **Customized programs:** Add them to the Windows startup folder or configuration files (system.ini, win.ini, autoexec.bat, config.sys, and so on). For UNIX-based systems, you can employ entries in the /etc/rc.d directory. A good example of this would be ZEUS as a massive botnet with a command and control screen for easy management and deployment.

Note Zeus is perhaps one of the largest botnets in the world, and it's an interesting read: [http://en.wikipedia.org/wiki/Zeus_\(trojan_horse\)](http://en.wikipedia.org/wiki/Zeus_(trojan_horse)).

There are cases when the attacker does not want to have a backdoor placed in the target system. Usually, this is in the case of corporate espionage, in which an attacker gains access to acquire a certain piece of information and leaves. In a situation involving corporate espionage, the attackers know what they want and have no interest in regaining access

to the system at a later time. In these types of attacks, the attackers' main goal is to cover their tracks so that no one will ever know what happened.

Where Are Attacks Coming From?

It is clear by now that the bad guys are out there on the Internet using all kinds of tools, from automated to those that target you specifically. Everyone knows that a public IP address is required to connect to the Internet. These addresses are allocated across the globe, so you should be able to find out where these attacks are going.

“Each quarter, Akamai publishes a quarterly “State of the Internet” report. This report includes data gathered across Akamai’s global server network about attack traffic, average & maximum connection speeds, Internet penetration and broadband adoption, and mobile usage, as well as trends seen in this data over time.”

(www.akamai.com/stateoftheinternet/)

Figure 1-4 shows the results for Q4 & Q3 for 2009 in a top 10 format by country originating the attacks. This report is available free upon registering on the Akamai website; it is an interesting report and worth reading.




Figure 1-4 Q4 2009 - Attack Traffic Top 10 Originating Countries

The interesting aspect of this data is that Akamai maps out the ports that it is seeing generating these attacks, as shown in Figure 1-5. This top 10 list of attacked ports is dominated by Microsoft-DS (tcp port 445) used by its Server Message Block (SMB) running on TCP to enable a variety of things, most notably file sharing. It is speculated that this domination of list is because of the global infection of PCs running Windows that have been infected by Confiker.




Figure 1-5 Q4 2009 - Top 10 Attack Ports

Conficker is a computer worm that was first detected in 2008. It uses TCP port 445 aka Microsoft-DS by targeting the Windows operating system. This worm is believed to have the largest number of infected PCs since the inception of the Internet; estimates vary, but 6 million plus is generally agreed to be the low estimate. Although Microsoft has announced a \$250,000 reward for information leading to the arrest of the Conficker creator, this author feels that the money would be better spent making its software more secure given how many of that attacks relate to it. Cisco also has free global correlation for IPS intelligence to provide additional visibility into attack mitigation. OWASP, owasp.org, also maintains its top 10 application threats as it evolves—remediation tools including test tools such as Web Goat and well-documented application security guides.

Common Vulnerabilities, Threats, and Risks

This section reviews some of today's more common vulnerabilities, threats, and risks that you will face. As a general rule, imperfect people create imperfect software, and they make mistakes unintentionally allowing vulnerabilities to be exploited by hackers. This list provides a brief synopsis and examples to help increase awareness, enabling you to protect and educate your users. The next section deals with attack examples.

- **Antivirus software:** A software program dedicated to protecting your computer from viruses. As threats are evolving, so are these programs. With the decreasing occurrences of virus and increased exploits and attacks, these programs have developed into suites of programs designed to protect you while browsing the Internet and from threats you might not see. Unfortunately, these programs are anything but perfect; they, too, have limitations and bugs. For example, a bug may enable the program to be stopped or not to update. Of course, without updates the software cannot recognize attack variations and changes. Industrywide, the move is from signature to anomaly-based antimalware and host IPS so that the end user is not left waiting for the latest signature, much like anomaly detection on network-based IPS to be used with signature and correlation services. Initial attacks known as *zero-day threats*, which are much more common today, exist where no signature is available initially to remediate such a threat or exploit.
- **Media players:** It is common to have links in websites or files (music or movies) that when clicked start your system's media player, thus allowing the content to be played. Hackers have learned to embed in these links or files means to exploit vulnerabilities within media players, yet another example to only click links you trust.
- **Adobe Flash:** One of the de facto standards of web content, Flash is quickly becoming the playground of hackers because they love to find vulnerabilities in Flash and exploit them. The worrisome aspect of these vulnerabilities is that hackers can infect Flash on servers and PCs; thus, you may go to a reputable site that has been compromised already. These concerns and others have caused a backlash against Adobe, resulting in its developers and users slowly reconsidering its use for more open standards such as HTML5.

- **Adobe Reader and Acrobat:** Adobe has another winning piece of software on its hands here. These products enable us to read and create PDF files, yet another de facto standard online about how to protect and share documents. This success also means it is drawing interest by hackers, who go where the users are. The risk here is that many companies by default permit PDFs easily throughout their network, whereas they are actively scanning or blocking other document formats, making these new vulnerabilities a serious concern.
- **Backup software:** Backups are critical, and as computing moves to virtual machines stored on storage arrays, the ability to cause damage by finding exploits in this type of software is rather serious. Past vulnerabilities have allowed entire servers to be hacked rather easily.
- **Database software:** Whole books have been written on databases and securing them, which isn't discussed here. The vulnerabilities and threats are not unique; several programs dominate the market and run most database applications. These applications are often web-based or have a web interface, enabling exploits to allow DoS attacks and deep exploitation when initially compromised.
- **Email clients:** Hackers have changed tactics and use email vulnerabilities to corrupt and compromise email clients. They are preying on user ignorance, which unfortunately has given them a fertile field of growth because users believe almost any email they get and click away. User awareness, security education, and regular software patching are important as server-based technologies are evolving to protect users and their email. You can put web and email security point security solutions in place in addition to succinct policy to support those controls.
- **Excessive user rights:** Have a single sign-on domain? Most do, and the threats here are based around users with rights to areas and things they do not need; they might want, but they do not need—a rather important distinction. Hackers then rely on weak password policies or nonexistent policies that enable users to never change passwords or use extremely weak ones; thus, the user gets hacked, and the hacker has sufficient privileges to hack again.
- **Instant messaging and social messaging:** Another communication mechanism that is seeing users adopt it like crazy, with hackers following suit by researching a variety of different ways to attack systems. One thing that astounds me is that when people randomly message you; typically they are hot women or at least claim to be. The hacker's hope is that you click on a shared link or otherwise behave in a way you normally would not. User education and awareness is key as the old axiom of "If its too good to be true, it is, besides all the money in those accounts people in Africa have is all mine!" Encryption should be considered when it is required for businesses to use instant messaging.
- **Office software:** This sort of threat is one that hackers are not more reactive in; they will infect a document or spreadsheet, say, via a macro vulnerability. They make the document useful and place it online or replace a good copy with an infected copy. Users come along, download and open the file, and poof, they become infected,

enabling the hacker access to your computer without you even being aware of what has occurred.

- **Removable media:** Perhaps not a traditional vulnerability but still much has been made of allowing USB keys to flow freely between users' homes and corporate resources. In many high-security environments, they are not allowed. The belief is that after a USB key gets infected it can be spread unwittingly between physical PCs in different locations. Unfortunately, manufacturers have been slow to place security onto these sorts of removable devices, further adding to the problem. Encryption should be considered for removable media and whole disk if possible managed by a comprehensive PKI.
- **DNS:** The Domain Name System (DNS) is a distributed resource used by most network applications, especially the newer versions of Windows. DNS data is generally trusted implicitly. Considering that DNS is the lynchpin of the corporate enterprise, the impact of these vulnerabilities is significant, and a successful attack could jeopardize the integrity of any network. In my experience, patching your DNS application security and hardening is critical.
- **SSL:** Perhaps the most common data security protocol on the Internet, you can find SSL in use at every e-commerce site to protect data during transactions. Current threats enable a hacker to get in the middle of the information flow between the user with the credit card and the server that processes his order after he hands over all the credit card information. Granted, it is difficult to get in the middle like that, but it is not impossible. SSL decryption is possible, as mentioned earlier, but this is a time- and resource-consuming effort.
- **Phishing:** Pronounced *fish-ing*, this term describes a hacking technique that attempts to acquire sensitive information such as credit card numbers, bank accounts, usernames, passwords, and so on. Hackers do this via electronic communications such as email or impersonating a website to trick people into revealing this information without knowing they did so. Phishing is perhaps one of the most advanced criminal techniques and successful use of social engineering in use on the Internet today. Phishing has been extremely successful in tricking people to reveal their secrets through the hacker's use of these bait (email) and catch (fake website) techniques.
 - **Spear phishing or whaling:** These terms are used to describe a specific type of phishing, such as when these email techniques target high-profile individuals, such as corporate executives.
 - **Vishing:** This attack sends users an email claiming to be a financial institution that needs the victims to call a phone number about some fictitious problem with their account. These phones are owned by the hacker and are provided via a Voice over IP (VoIP) service so that when the victims dial the account, it prompts them for their account numbers and PIN, recording it all.
- **Peer-to-Peer (P2P):** These types of networks are quite common these days and have continued to be risky to users. Typically, attacks here come in several forms; there is a passive attack, where a hacker creates a desirable file that people can download and

access, thereby infecting their PCs. The more active version is when hackers take advantage of a P2P network's design allowing them to execute a man-in-the-middle attack.

- **Web browsers:** As the primary tool of people accessing the Internet, everyone should be well aware of the security issues surrounding every web browser.

Overview of Common Attacks and Exploits

This section reviews some of the more commonly used attacks and exploits available to attackers. It should by no means be considered complete because new attacks are discovered at an alarming rate every day. For a more complete list or more information on the exploits listed here, refer to any of the organizations presented in the previous section:

- **Denial of service (DoS):** A DoS attack attempts to force the target into a failure condition, thereby denying its services to others. There are several ways in which a failure condition can be induced, such as flooding the target with attempts to connect (http://en.wikipedia.org/wiki/Denial-of-service_attack).
- **Distributed denial of service (DDoS):** This type of attack uses a collection of unknowing accomplices to attack a target from multiple locations at once. The accomplices are compromised machines spread out in many different places.
- **Zero day attacks:** A security term used to describe when a new attack is launched; the key is this is before security professionals have detected it—hence, zero day. When detected, it becomes day one.
- **Botnets:** A grouping of compromised machines running malicious software under control of a single controller or bot master. This malicious software is stealthily run and communicates extremely securely with the command and control server; these communications can typically be done via chat and instant messaging. Botnets are rented out to third parties for them to send spam or join in as part of a DDoS. The largest botnet as of this writing is Conficker, with an estimated 10+ million machines under its control.
- **SYN flood attack:** A SYN flood attack occurs when a network becomes so overwhelmed by SYN packets initiating incomplete connection requests that it can no longer process legitimate connection requests (thereby causing high CPU, memory, and network usage) and resulting in a DoS.
- **UDP flood attack:** Similar to the ICMP flood, UDP flooding occurs when UDP packets are sent with the purpose of slowing down the system to the point that it can no longer handle valid connections. Port 53–DNS flooding is the hallmark *modus operandi* of this kind of attack.
- **Port scan attack:** Port scan attacks occur when packets are sent with different port numbers with the purpose of scanning the available services, in hopes that one port will respond. When a port is detected as open (because it responded), the hacker can begin looking for ways to compromise the system through that port.

- **IP spoofing:** Spoofing attacks occur when an attacker attempts to bypass the firewall security by imitating a valid client IP address, email address, or user ID. This becomes important when an attacker decides to exploit trust relationships that exist between computers. Usually, administrators set up trust relationships between multiple computers; one of the side benefits to this is a single login for all.
- **Land (C) attack:** Combining a SYN attack with IP spoofing, a land attack occurs when an attacker sends spoofed SYN packets that contain the victim's IP address as both the destination and source IP address. The receiving system responds by sending the SYN-ACK packet to itself, thereby creating an empty connection that lasts until the idle timeout value is reached. Flooding a system with such empty connections can overwhelm the system, resulting in a DoS condition on the target system.
- **Tear drop attack:** Tear drop attacks exploit the reassembly of fragmented IP packets. In the IP header, one of the options is offset. When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash.
- **Ping scan:** Similar to a port scan attack, a ping scan attack occurs when an attacker sends ICMP echo requests (or pings) to different destination addresses in hopes that one will reply and, therefore, uncover a potential target's IP address.
- **Java/ActiveX/ZIP/EXE:** Malicious Java or ActiveX components can be hidden in web pages. When downloaded, these applets install a Trojan horse on your computer. Similarly, Trojan horses can be hidden in compressed files such as .zip, .gzip, .tar, and executable (.exe) files. Enabling this feature blocks all embedded Java and ActiveX applets from web pages and strips attached .zip, .gzip, .tar, and .exe files from email.
- **Smurf:** The little blue folks are not coming back to make your day; rather, **ping** (ICMP) is being used to target devices via an intermediate device, thus hiding the attacks from the true source. You can read more about Smurf attacks at www.cert.org/advisories/CA-1998-01.html.
- **Brute force:** In a brute force attack, an attacker tries to guess passwords through techniques such as repeatedly trying to log in to an account by using a dictionary of potential passwords.
- **Source routing:** Source routing is an option in an IP packet's header that defines how packets are routed. When this option is on many firewalls, rules are bypassed, thereby allowing access to your network. For example, the IP header information can contain routing information that can specify a different source IP address than the header source. This causes the packets to be routed in a different direction. Following are several other ways to control the routing of ICMP packets:
 - **Record route:** An attacker sends packets where the IP option is 7 (Record Route). This option is used to record the route of a packet. A recorded route is composed of a series of Internet addresses that an outsider can analyze to learn details about your network's addressing scheme and topology.

- **Loose source route:** An attacker sends packets where the IP option is 3 (Loose Source Routing). This option provides a means for the source of a packet to supply routing information for the gateways to use to forward the packet to the destination. This option is a loose source route because the gateway or host IP is allowed to use any route of any number of other intermediate gateways to reach the next address in the route.
- **Strict source route:** An attacker sends packets where the IP option is 9 (Strict Source Routing). This option provides a means for a packet's source to supply routing information for the gateways to use to forward the packet to the destination. This option is a strict source route because the gateway or host IP must send the datagram directly to the next address in the source route, and only through the directly connected network indicated in the next address to reach the next gateway or host specified in the route.
- **ICMP flood:** An ICMP flood occurs when ICMP pings overload a system with so many echo requests that the system expends all its resources responding until it can no longer process valid network traffic. Several different types of ICMP messages exist, each with its own purpose, and attackers can use them:
 - **ICMP Echo Reply:** (Code 0, Echo Reply) A response to a ping. Many firewalls enable ping responses so that internal people can gain access to external resources. Therefore, they are an effective flooding technique.
 - **ICMP Host Unreachable:** (Code 3, Destination Unreachable) An error message from a host or router indicating that a packet you sent did not reach its destination.
 - **ICMP Source Quench:** (Code 4, Source Quench) A response indicating congestion on the Internet. Someone might be trying to flood your network with these packets in an attempt to convince your machines to slow down data transmission.
 - **ICMP Redirect:** (Code 5, Redirect) A message advising to redirect traffic; for example, for network X directly to gateway G2 because this is a shorter path to the destination. Someone might be trying to redirect your default router. This could be from a hacker trying to execute a man-in-the-middle attack against you by causing you to route through his own machine.
 - **ICMP Echo Request:** (Code 8, Echo Request) These are commonly used ping request packets. They might indicate hostile intent of someone trying to scan your computer, but they might be part of the normal network functionality.
 - **ICMP Time Exceeded for a Datagram:** (Code 11, Time Exceeded in Transit) A message indicating that a packet never reached its target because something timed out.
 - **ICMP Parameter Problem on Datagram:** (Code 12, Parameter Problem on Datagram) A message advising that something unusual is going on; this probably indicates an attack.
 - **Large ICMP Packet:** An ICMP packet with a length greater than 1024 can cause trouble for some devices because ICMP packets are not normally this size.

- **Sniffing packets:** The use of a sniffer is a passive attack that allows a network interface card to be placed into a special mode: promiscuous. Do not be fooled into thinking that there is no danger because it is a passive attack. For an attacker to get a sniffer on your LAN, serious security issues have already occurred. Now that the attacker can see most of the packets on your LAN with a sniffer, there is a definite threat.

This is simply a short list of the thousands of vulnerabilities known today. Now imagine the effectiveness of a coordinated attack using some of these vulnerabilities. It puts it in a different perspective, doesn't it?

Network Security Organizations

This section primarily examines some of the exploits and vulnerabilities available to attackers. Prior to that though, it is important to look at where you can go to learn about vulnerabilities and other security-related information; several organizations are covered with the descriptions “in their own words” direct from their websites.

At one time, each vendor or manufacturer was responsible for tracking all the vulnerabilities that affected its products. The result was that different companies would report that same vulnerability, thereby causing some confusion—or perhaps they would *not* acknowledge the vulnerability until it became public. The network security industry realized that this was not efficient, and it created common vulnerabilities and exposures (CVE). Do not misunderstand; CVE is not a database of vulnerabilities, but a dictionary that defines its role as follows:

Common Vulnerabilities and Exposures (CVE [www.cve.mitre.org/]) is a dictionary of common names (i.e., CVE identifiers) for publicly known information security vulnerabilities, while its Common Configuration Enumeration (CCE) provides identifiers for security configuration issues and exposures.

CVE's common identifiers makes it easier to share data across separate network security databases and tools, and provide a baseline for evaluating the coverage of an organization's security tools. If a report from one of your security tools incorporates CVE Identifiers, you may then quickly and accurately access fix information in one or more separate CVE-compatible databases to remediate the problem.

CVE is

- One name for one vulnerability or exposure
- One standardized description for each vulnerability or exposure
- A dictionary rather than a database
- How disparate databases and tools can “speak” the same language
- The way to interoperability and better security coverage
- A basis for evaluation among tools and databases

- Free for public download and use
- Industry-endorsed via the CVE Editorial Board and CVE-Compatible Products

CERT Coordination Center

The CERT Program (www.cert.org/) is part of the Software Engineering Institute (SEI), a federally funded research and development center at Carnegie Mellon University in Pittsburgh, Pennsylvania. Following the Morris worm incident, which brought 10 percent of Internet systems to a halt in November 1988, the Defense Advanced Research Projects Agency (DARPA) charged the SEI with setting up a center to coordinate communication among experts during security emergencies and to help prevent future incidents. This center was named the CERT Coordination Center (CERT/CC).

While CERT continues to respond to major security incidents and analyze product vulnerabilities, our role has expanded over the years. Along with the rapid increase in the size of the Internet and its use for critical functions, there have been progressive changes in intruder techniques, increased amounts of damage, increased difficulty of detecting an attack, and increased difficulty of catching the attackers. To better manage these changes, the CERT/CC is now part of the larger CERT Program, which develops and promotes the use of appropriate technology and systems management practices to resist attacks on networked systems, to limit damage, and to ensure continuity of critical services.

SANS

The SANS (SysAdmin, Audit, Network, Security) Institute (www.sans.org/) was established in 1989 as a cooperative research and education organization. Its programs now reach more than 165,000 security professionals around the world. A range of individuals from auditors and network administrators to chief information security officers are sharing the lessons they learn and are jointly finding solutions to the challenges they face. At the heart of SANS are the many security practitioners in varied global organizations from corporations to universities working together to help the entire information security community.

Many of the valuable SANS resources are free to all who ask. They include the popular Internet Storm Center (the Internet's early warning system), the weekly news digest (NewsBites), the weekly vulnerability digest (@RISK), flash security alerts, and more than 1,200 award-winning, original research papers.

Center for Internet Security (CIS)

The mission of the Center for Internet Security (CIS) is to establish and promote the use of consensus-based standards to raise the level of security and privacy in Internet-connected systems, and to ensure the integrity of the business, government, and private Internet-based functions and transactions on which society increasingly depends. CIS (<http://cisecurity.org>) is an independent organization governed by a volunteer board of directors; it is not owned or controlled in full or part by any corporation or government entity.

CIS develops and distributes the following:

- Security configuration benchmarks describing consensus best practices for the secure configuration of target systems. Configuring IT systems in compliance with these benchmarks has been shown to eliminate 80 percent to 95 percent of known security vulnerabilities. The benchmarks are globally used and accepted as the de facto user-originated standard for IT security technical controls.
- Benchmark audit tools for assessing compliance with CIS benchmarks.
- Security metrics that offer enterprise IT and security teams insight into their own security process outcomes.

SCORE

SCORE is a cooperative effort between SANS/GIAC and the Center for Internet Security (CIS). SCORE (www.sans.org/score/) is a community of security professionals from a wide range of organizations and backgrounds who work to develop consensus regarding minimum standards and best-practice information. It essentially acts as CIS's research engine. After consensus is reached and best practice recommendations are validated, CIS can formalize them as best practice and minimum standards benchmarks for general use by industry at large.

SCORE objectives are as follows:

- Promote, develop, and publish security checklists.
- Build these checklists via consensus and through open discussion via SCORE mailing lists.
- Use existing references, recruit GIAC-certified professionals, and enlist subject matter experts where and whenever possible.

Internet Storm Center

Internet Storm Center (<http://isc.sans.org/>) defines itself as a center that gathers more than 3,000,000 intrusion detection log entries every day. It is rapidly expanding in a quest to do a better job of finding new storms faster, isolating the sites that are used for attacks, and providing authoritative data on the types of attacks that are being mounted against computers in various industries and regions around the globe. Internet Storm Center is a free service to the Internet community. The SANS institute supports the work with tuition paid by students attending SANS security education programs.

National Vulnerability Database

The National Vulnerability Database (NVD) is the U.S. government repository of standards-based vulnerability management data represented using the Security Content Automation Protocol (SCAP). This data enables automation of vulnerability management, security measurement, and compliance. NVD (<http://nvd.nist.gov/>) includes databases of

security checklists, security-related software flaws, misconfigurations, product names, and impact metrics.

Security Focus

Since its inception in 1999, SecurityFocus has been a mainstay in the security community. From original news content to detailed technical papers and guest columnists, it strived to be the community's source for all things security-related. SecurityFocus was formed with the idea that community needed a place to come together and share its collected wisdom and knowledge. At SecurityFocus, the community has always been the primary focus. The SecurityFocus website now focuses on a few key areas of greatest importance to the security community:

- BugTraq is a high-volume, full-disclosure mailing list for the detailed discussion and announcement of computer security vulnerabilities. BugTraq serves as the cornerstone of the Internet-wide security community.
- The SecurityFocus Vulnerability Database provides security professionals with the most up-to-date information on vulnerabilities for all platforms and services.
- SecurityFocus Mailing Lists enable members of the security community from around the world to discuss all manner of security issues. There are currently 31 mailing lists; most are moderated to keep posts on-topic and to eliminate spam.

Learning from the Network Security Organizations

These organizations did not always exist, but the increase in threats across the Internet from attackers of all types has supported their birth and growth. You should explore each website because there is a wealth of information that takes you beyond what is presented here. The following section reviews some of the ways vulnerabilities and exploits are used in attacks.

One of the useful things that manufacturers are doing these days is setting methods for users and white-hat hackers (good guys) to report security issues with their products. For example, Cisco has provided this information to you online:

Cisco Security Advisories & Notices:

www.cisco.com/en/US/products/products_security_advisories_listing.html (PSIRT)

Cisco Security Intelligence Operations: <http://tools.cisco.com/security/center/>

Cisco produces a report on cyber risks every few months and can be found on the preceding website—it is worth looking at.

Chapter Summary

This chapter examined the ways an attacker selects his targets, as those of opportunity or those of choice. Ultimately, you learned that everyone is a target, and the true differentiator comes when attackers either stumble across an unprotected target or when there is perhaps a deeper and more malicious intent in the attacker's selection.

After attackers determine that you are a target, they employ six common steps, which form the components of the attack whose goal is the ultimate compromise of a system.

This chapter also discussed online places to learn more about network security. These places were the “good guys,” and it is important to point them out because most locations on the Internet are the bad guys; be careful visiting these websites! Instead, read the last part of this chapter, where a few of the attacks and possible exploits were discussed. The following chapter discusses the next step in understanding network security—security policies, which are the first step in protection.

Chapter Review

Each chapter concludes with a “Chapter Review” section. In a question-and-answer format, the “Chapter Review” section tests the basic ideas and concepts covered in each chapter. In tandem with the “Chapter Objectives” and “Chapter Summaries,” the “Chapter Review” section builds upon and reinforces key ideas and concepts. Each “Chapter Review” section is composed of a series of topical questions and answers to the “Chapter Review” section are included in Appendix A, “Answers to Chapter Review Questions.”

1. What is a target of opportunity?
2. What is a target of choice?
3. What is the purpose of footprinting?
4. Which of the following are ways by which an attacker can gain access?
 - a. Operating system attacks
 - b. Application attacks
 - c. Misconfiguration attacks
 - d. Script attacks
 - e. DoS or DDoS
 - f. All the above
5. List four of the network security organizations.
6. Briefly explain why it is important for an attacker to cover his tracks.
7. Social engineering can be damaging to a corporation without an overt attack ever happening. Explain why.

- 8.** What kind of information might be found if an attacker dumpster dives at your place of work?
- 9.** DNS information gained through WHOIS is used for what kind of reconnaissance?
- 10.** What two free reconnaissance tools are available with most versions of the Windows operating system?

Security Policies

“...Being defeated is often a temporary condition. Giving up is what makes it permanent...”—Marlene vos Savant

By the end of this chapter, you should know and be able to explain the following:

- What role does a security policy play in my network?
- How do I create a security policy?
- How do I deal with any security policy violations?
- What security policies are appropriate for my organization?
- What are the Security Standards and do they apply to my organization?

Being able to answer these key questions will enable you to understand the overall characteristics and importance of a network security policy.

Having clear, definable, enforceable, and up-to-date security policies is the most essential first step in protecting and securing your people, property, network, and data. Policies provide the foundation for defining acceptable and appropriate behavior within your organization and network. At the most fundamental level, policies form the “rule of law,” which is the legal maxim stating that no one is immune to the law, or in this particular circumstance, no one is immune to the policy. Short- and long-term contractors and consultants can be tied to policy via service-level agreements (SLA) with similar verbiage.

Consider a security policy that is analogous to rules and laws found in your neighborhood. What would life be like to live within these boundaries? If you wanted to accomplish anything worthwhile, it would be unbearable. Viewed in this light, a security policy defines what is acceptable, or not, inside and outside your network. This is a fundamental definition of the role of security policy, yet there are many additional reasons that define a security policy’s usefulness:

- Establishes expectations for standards, procedures, and guidelines
- Defines appropriate behavior

- Communicates an operational and business consensus
- Provides a foundation for HR action if unacceptable behavior occurs
- Defines roles and responsibilities of each group in securing the company
- Assists in prosecuting legal action if unacceptable behavior occurs
- Provides definitions for concepts and ideas crucial in securing your network
- Allows for required tools to be defined by justifying funds for network security

Having a security policy enables everyone within a company to clearly understand who is responsible for what and establishes a foundation for the policies and processes of each department within your organization, with an emphasis on “protect” from a corporate perspective. For example, customer service understands its roles and responsibilities in protecting sensitive customer information; human resources understands what is expected of employees; and manufacturing and development knows how to protect the results of expensive research and development. Of course, the greatest achievement of a security policy is what a security policy means to an IT department when trying to manage users and assets. From the security policy, the IT staff knows what to configure on servers, the tools it needs, rules for firewalls, virtual private network (VPN) secure remote access settings, and so on—the list is endless.

You might wonder what some of the most commonly used security policies are and what areas of IT should consider using a policy. The System Administration, Networking, and Security (SANS) Institute Security Policy Project and SANS certification (GIAC Fundamentals of Security Policy [GFSP]) provide a starting point to policy creation, templates, and process (www.sans.org/security-resources/policies/). SANS provides a variety of security policies, some of the more common of which are described in Table 2-1.

Table 2-1 *Common Security Policies*

Policy Name	Description
Acceptable Encryption	Provides guidance that limits the use of encryption to those algorithms that have received substantial public review and have been proven to work effectively. Additionally, provides direction to ensure that applicable laws and regulations are followed.
Acceptable Use	Outlines who can use company-owned computer equipment and networks. It covers company computers located on company premises and in employees' homes.
Analog/ISDN Line	Explains the analog and ISDN line acceptable use and approval policies and procedures. Separate rules apply to lines that are to be connected for the sole purpose of faxing and receiving and lines that are to be connected to computers.

Table 2-1 *Common Security Policies*

Policy Name	Description
Anti-Virus Process	Defines guidelines for effectively reducing the threat of malicious code on your network.
Application Service Providers (ASP) Standards	Describes the company's requirements of Application Service Providers (ASP). (ASPs combine hosted software, hardware, and networking technologies to offer a service-based application.) It refers to and incorporates the separate ASP Standards Policy.
Acquisition Assessment Policy	Defines responsibilities regarding corporate acquisitions, and defines the minimum requirements of an acquisition assessment to be completed by the information security group.
Audit Vulnerability Scanning	Provides the authority for members of the information security department team to conduct a security audit on any system owned by the company or installed on the company's premises.
Automatically Forwarded Email	Prevents the unauthorized or inadvertent disclosure of sensitive protected company information.
Bluetooth Device Security	Provides for more secure Bluetooth device (79 bands, 2.4 GHz, and typically weak security) operations. It protects the company from loss of personally identifiable information (PII) and proprietary company data.
Database Credentials Coding	States the requirements for strong cryptography and securely retrieving database usernames and passwords (that is, database credentials) for use by a program that accesses a database running on one of the company's networks.
Dial-in Access	Establishes rules that protect electronic information from being inadvertently compromised by authorized or unauthorized personnel using a dial-in analog connection.
DMZ Security	Provides definable standards for all networks and Internet-facing equipment located in the demilitarized zone or external network segments that might also be dictated by a solution architecture.
E-mail & E-mail Retention	Defines standards to prevent tarnishing of the public image of the organization. The E-mail Retention policy is intended to help employees determine what information sent or received by email should be retained and for how long.
Extranet	Defines the requirement of third-party organizations requiring access to the organization's networks must sign a third-party connection agreement. An example might be a corporate partner or service provider.

Table 2-1 *Common Security Policies*

Policy Name	Description
Information Asset Sensitivity	Helps employees determine what information can be disclosed to nonemployees, and the relative sensitivity of information that should not be disclosed without proper authorization.
Information System Audit Logging	Attempts to address the problem most organizations have concerning the vast amounts of information gathered during a typical day/week/month in a log file integration. It identifies specific requirements information systems must meet to generate appropriate audit logs and integrate those logs with an enterprise's log mgmt function.
Internal Lab Security	Establishes information security requirements for labs to ensure that confidential information and technologies are not compromised, and that production services and other interests are protected from lab activities.
Internet Usage	This document provides clear guidelines on what is, and is not, appropriate usage of Internet access on company time, such as social media sites, dating services, pornography, and so on.
Password	Establishes a standard for creation of strong passwords, the protection of those passwords, and the frequency of change.
Personal Communication Device	Describes information security's requirements for personal communications devices (PDAs/smartphones/iPad) and how and whether those devices are allowed in your organization or secure labs.
Remote Access	Defines standards for connecting to a company's network from any host. These standards are designed to minimize the potential exposure from damages such as the loss of sensitive or company confidential data, intellectual property, damage to public image, damage to critical internal systems, and so on. This typically includes secure remote access via VPN with AH and ESP but can still include malicious code.
Removable Media	Defines coverage of all computers and servers in an organization. The use of USB flash drives (that is, thumb drives), external hard drives, and CD/DVD burners.
Risk Assessment	Empowers the information security department to perform periodic information security risk assessments for the purpose of determining areas of vulnerability, and to initiate appropriate remediation.

Table 2-1 *Common Security Policies*

Policy Name	Description
Router and Switch Security	Describes a required minimal security configuration including hardening templates for all routers and switches connecting to a production network or used in a production capacity.
Server Security	Establishes standards for the base configuration of internal server equipment that is owned and/or operated on company premises or at web hosting locations.
Virtual Private Network	Provides guidelines for remote access IPsec, L2TP, SSL, EZ-VPN, or DVPN connections to the company's corporate network and protected resources.
Wireless Communication	Establishes standards for access of the company's network via secured wireless communication mechanisms conforming to any required specific regulatory and compliance related standards and bodies.

In addition to knowing what is expected, every person or department in a company is affected by a security policy and can make an impact individually protecting corporate reputation, brand, and assets. As you can see from the following list, each group within an organization is affected:

- **Generic user:** Because users access network resources, your policy impacts them the most.
- **Management team:** This group (and the executive sponsor for the information security policy program) is ultimately concerned with the protection of corporate resources and data while monitoring the financial impact.
- **Accountants, legal, and investors:** Understand that the company's responsibility to protect itself depends on such policies while recognizing the positive impact a security policy can have enabling the business and tying together the technical and business requirements.
- **Security management team:** This group's role is defined in the policy to pinpoint what group is tasked with security policy enforcement.

The following section discusses what could be viewed as the critical first question for designing a security policy: who and what to trust.

Responsibilities and Expectations

In an organization, the question of responsibility is a big one. Is the organization responsible if the end user misuses his company-owned laptop and gets caught with illicit material, or is the user responsible? What about for use of items such as PDAs/Blackberrys/smartphones, and so on? What about the use of sites such as Wikileaks, or social networking sites such as Twitter, Facebook, Linked-In, or MySpace? The proliferation of electronic media and the tendency for people to talk too much are not a good combination. Ultimately, it is the responsibility of the organization to protect itself, the information housed in its databases, and its users. To do so, you must implement policies, standards, procedures, and guidelines to ward off potential lawsuits, loss of intellectual property (IP), and loss of resources, and set forth the expectations you have of your personnel. It won't be a cure-all, but it will give the organization a leg to stand on when the need arises to protect corporate assets.

A Real-World Example

In 2009, Heartland Payment Systems,¹ which processes card payments for restaurants, retailers, and other merchants, was attacked by intruders who hacked into the system used to process 100 million payment card transactions per month for 175,000 merchants. Essentially, the hackers wormed their way into the system and recorded Heartland's system for weeks in late 2008. The CISP then PCI-DSS security standards were created to prevent examples like this from happening.

Who Is Responsible? You Are!

It is the responsibility of the organization to protect its personnel, the organization, the data entrusted to that organization (such as credit card, bank accounts, and Social Security numbers), and the IP resources (designs, plans, and code) of the organization, if applicable. You should be aware of legal precedence, ISO certifications, and security standards that pertain to your organization, or type of organization, such as a medical office and the Health Insurance Portability and Accounting Act (HIPAA). Ignorance is not bliss when your organization gets levied with fines; it's expensive—and depending on your tier level when it comes to PCI-DSS, which is quite prescriptive from a technology perspective, can run into the millions of dollars of fines to the corporation at fault.

Legal Precedence

The United States Department of Justice maintains a Computer Crime and Intellectual Property website (www.justice.gov/criminal/cybercrime/index.html) providing you with news releases for computer crime, current and archived cases, policy and programs, legal resources, and so on—there is a wealth of information out there.

¹ www.usatoday.com/money/perfi/credit/2009-01-20-heartland-credit-card-security-breach_N.htm

Internet Lawyers

There is new legal precedence everyday concerning such things as intellectual property, Internet law, and domain disputes. If you are in need of protection, you can find a good Internet lawyer. These individuals specialize in Internet laws and information technology company representation. They specialize in Internet law and understand better than anyone the challenges that occur when legal issues arise in cyberspace. They deal in specialties such as areas of jurisdiction, venue, breach of contract, e-commerce, trademarks, copyrights, and patents, and other issues often out of the realm of typical corporate lawyers.

Evolution of the Legal System

As we move forward the legal system is evolving. Many states and countries now have specific laws concerning cyber security and the protection of people and their assets.

Information Technology Law (IT Law) is a set of recent legal enactments that digitally govern the process and dissemination of information. These legal enactments cover a broad range of different aspects relating to computer software, protection of software code, access, and control of digital information, privacy, security, Internet access and usage, and electronic commerce.

- Florida Electronic Security Act
- Illinois Electronic Commerce Security Act
- Texas Penal Code - Computer Crimes Statute
- Maine Criminal Code - Computer Crimes
- Singapore Electronic Transactions Act
- Malaysia Computer Crimes Act
- Malaysia Digital Signature Act
- UNCITRAL Model Law on Electronic Commerce
- Information Technology Act 2000 of India
- Computer Misuse Act of 1990 (Great Britain)

A prime example of this is the September 2009² case in which an Indiana couple was allowed to sue their bank for its alleged failure and negligence to implement the latest security measures. The judge ruled that a "...reasonable finder of fact could conclude that the bank breached its duty to protect Plaintiff's account against fraudulent access..." It was pointed out that the authentication methods were inadequate. That the bank relied on usernames and passwords to control access to accounts whereas other banking institutions had begun using two-factor or multifactor authentication, including token-based authentication (hardware and software tokens). The couple highlighted the fact that a 2005 document authored by the Federal Financial Institutions Examinations Council (FFIEC), called single-factor authentication inadequate and recommended the use of two-factor authentication by banks.

2 www.computerworld.com/s/article/9137451/Court_allows_suit_against_bank_for_lax_security

Criminal Prosecution

With new laws and new legal precedents being made everyday, prosecuting individuals and corporations is becoming more and more common. Your organization can be sued for breach of contract because of insufficient security, insider trading, or just loss of your clients' PII—all the more reason to know what data needs to be secured and how to best do so without hindering productivity. You run into the continual situation in which maintaining a secure platform hinders the research and development aspect of your job. There needs to be security, but there also needs to be an understanding of why you protected those assets.

Real-World Example

In 2007, TJX Companies (T.J. Maxx, Marshalls, and Bob's Stores) revealed that some 45.6 million credit and debit card numbers were stolen from one of its systems over a period of more than 18 months by an unknown number of intruders. In addition to that, the personal data provided with the return of merchandise without receipts by an estimated 451,000 individuals in 2003 was also stolen. Now you might be asking yourself how this happened. In addition to poor wireless network security (the WEP key was easily cracked—the current standard PCI-DSS 2.0 10/28/10 does not enable the use of WEP—www.pcisecuritystandards.org), "...poorly secured in-store computer kiosks are partly to blame...." The kiosks that enabled individuals to electronically apply for jobs were not isolated on the network and enabled direct access to the company's network infrastructure. The people who started the breach opened up the back of those terminals and used USB drives to load software onto those terminals.

Attorneys are suing retailer TJX citing TJX failed to comply with 9 of 12 applicable PCI requirements and that the data thief managed to walk away with 80 gigabytes of data on TJX customers. Following are some of the security issues:

- An improperly configured/secured wireless network
- Failure to isolate and secure cardholder data devices from the rest of the network
- Failure to properly securely manage the systems used to store, process, and transmit cardholder data
- Insecurely storing prohibited cardholder data
- Using usernames and passwords that were easy to crack or guess
- Weak or nonexistent security software and systems

The most heinous allegation in the court filings are charges that TJX was aware of the security problems and failed to disclose the risks or remedy those problems; those inactions have increased the company's liability under the law.

Individuals Being Prosecuted

This isn't just concerning larger corporation getting sued or being held liable. Even the people who work on the systems are not immune from prosecution; take for example the case of Terry Childs.³ Terry Childs was the network engineer for the city of San Francisco. At a meeting he was asked by his boss to reveal the passwords to the FiberWAN and essentially relinquish control. Terry refused. The city of San Francisco then slapped him with four criminal charges and set his bail at \$5 million dollars; only one charge stuck and that was specifically "...Childs violated a California statute regarding illegal denial of service for the San Francisco FiberWAN..." Ultimately, Terry was found guilty and sentenced to 2 to 5 years for what amounts to not having given up the passwords to the network. Once again, role-based access control and segregation of duties is critical to prevent this type of criminal and job protection behavior.

International Prosecution

Furthermore, prosecution is no longer limited to the United States. In Sept 2009, Computer World ran a story about a long-running cybercrime operation.⁴ New York prosecutors indicted five eastern European men in an extensive credit-card fraud operation that saw the theft of more than \$4 million using nearly 95,000 stolen credit card and debit card numbers. This was the third phase of a four-year investigation involving law enforcement agencies in the U.S., the Czech Republic, Greece, and Ukraine. A total of 17 defendants have been indicted—these are very serious crimes.

We don't bring these incidences to your attention to scare you but as a spotlight on the importance of establishing, implementing, and maintaining a strong, sound security strategy. The best way to protect yourself and the sensitive information you are entrusted with as an organization is to begin by establishing sound policies and procedures where security is concerned. Information security and policy are not a nice-to-have, they are a must-have and need to be effectively communicated, kept up to date, and built in to the overall SDLF (software development lifecycle) for an organization.

Corporate Policies and Trust

Trust is a central theme in many aspects of security and must be foremost in your mind when discussing security policies. In a perfect world, there would be no issues with trust; you would trust everyone, and they would always do the right thing. Unfortunately, that is not realistic, nor does it take into account other factors, such as bugs in network resources. Again, trusting the resources on your network would be great, but remember that buggy hardware and software is commonplace in networking.

A security policy can be written with the belief that no one in an organization is to be trusted; however, that would not likely work. It is a well-known fact that users circumvent policies that are too restrictive. For example, in my organization we run an R&D lab, and

³ www.ktvu.com/news/23283217/detail.html

⁴ www.computerworld.com/s/article/9137403/Five_indicted_in_long_running_cybercrime_operation

the overarching posture is one of restriction and lockdown. It just so happens that the individuals doing our R&D are Ph.D.s and big types who continually circumvent the policy so that they can do things such as testing and downloading of the latest and greatest and coolest “gadgets.” It is frustrating from an information assurance (IA) position and from the position of having to enforce and reiterate to these developers why they cannot do what they want. There needs to be a balance between trust and securing the network. This balance is different for each organization, but the need for security does not change. To begin this journey into trust and balance and defining what is acceptable from a risk perspective, let us first look at the importance of having policies that reflect current threats. That is, make your security policy pertinent and continually update it. The second aspect you need to look at is the importance of user awareness and education.

Relevant Policies

Many times I have sat down in a new organization and asked the “CXO” in charge (that is, COO, CFO, CEO) to review his/her current security and IT policy letters, and the answer is, nine times out of ten, “...we haven’t reviewed it since we wrote it X years ago or we are working on that—it’s in flux...” Not having a current policy set is unacceptable and opens your corporation to significant risk from a business and technical perspective. If you are reading this and do not know what the date of your last policy is, make a note to yourself to check it in the morning. With the number of new standards coming out, the threat of new network incursion, and the potential threat of lawsuits, you cannot afford to let this be a nonissue. If worse comes to worst and you need to be held accountable, like the TJX corporation you just read about, and you don’t have a current *enforced* policy, you will find yourself on the losing end of the lawsuit.

Notice I say *enforced*. If you have a policy and do not enforce the application of the policy, the policy does you no good. It is but clanging brass and cymbals. The first part of enforcing a policy is end-user education.

User Awareness Education

Congratulations! You and your IT staff have spent hours cooped up in a hot conference room, wading through hours of arguments, jockeying for position, and fighting for industry best practices to meet in the middle...you have achieved balance, the zen of IT security policy, your opus! You and your team pat each other on the back, wrap it up neatly in a three-ring binder, and hand it off to the boss; and there it sits, flaccid and dormant, doing nothing more than gathering dust on a shelf. You need to implement the policy, and you begin by having training. You must educate your end users about proper security awareness. If you use a PKI smartcard for logging in to the organization’s computer systems and you are implementing a policy in which the users are required to remove the card from their keyboards when they leave their office/cube, you need to tell them; otherwise, they won’t know, and how are you to enforce the policy if something happens?

Part of your security policy should reflect training; if it doesn’t, look at your IT security policy again. Make yearly refresher training compulsory, and follow up the yearly training

with a publication. The time and overhead that you spend tending to this detail will be offset by the confidence that your personnel, IP, and networks are protected. Find creative ways to get the word out. A few good examples are a security awareness day, voice messages, video messages, and succinct email campaigns and posters.

And lest I forget, being a victim of this oversight myself...your IT staff needs training, too. Don't expect to implement an IT security policy, expect your employees to adhere to it, and expect your IT staff to know and understand what the latest threats and newest designs for protection are. They will not. The industry and threats evolve too rapidly. Too many times the staff implementing the policy cannot enforce it, but not due to apathy. No! This inability is due to ignorance. Don't expect your staff to accurately implement your security requirements if they do not understand them.

Coming to a Balance

When considering the level of trust to write into a security policy, consider the following items and keep them in mind as your policy is being developed:

- Determine who receives access to each area of your network based on their roles by using role-based access control (RBAC).
- Determine what they can access and how (RBAC).
- Balance trust between people and resources (segregation of duties).
- Allow access based on the level of trust for users and resources (RBAC).
- Use resources to ensure that trust is not violated and the risk is managed and measured.
- Define the appropriate uses of your network and its resources (Network Admission Control [NAC] or Group Policy Object [GPO]).

You need to consider many other things beyond this short list, including your company's politics and users' reactions. A security policy cannot account for every consideration, but you need to understand the reactions that a security policy brings out in people—your team needs to be business enablement focused, and there will be compromises for both of your teams along this long road and partnership.

Corporate Policies

According to the SANS Security Policy Project, security policies should emphasize what is allowed, not what is prohibited; where appropriate, supply examples of permitted and prohibited behavior. This way, there is no doubt—if not specifically permitted by the security policy, the behavior in question is prohibited. The security policy should also describe the ways to achieve its goals. Table 2-2 lists the security policy sections and describes their content.

Table 2-2 *Generic Description of a Security Policy's Contents*

Section Name	Content Guide
1.0 Overview	Justifies the reason for the policy and identifies the risks that the policy addresses.
2.0 Purpose	Explains why the policy exists and the goal that it is written to accomplish.
3.0 Scope	Defines the personnel that the policy covers. This might range from a single group in a department to the entire company.
4.0 Policy	This is the policy itself. It is often broken down into several subsections. Examples often illustrate points or facilitate the user's understanding.
5.0 Enforcement	Defines the penalty for failure to follow the policy. It is usually written as "everything up to and including..." so that a series of sanctions can be applied. Dismissal is typically the most severe penalty, but in a few cases, criminal prosecution should be listed as an option.
6.0 Definitions	Any terms that might be unclear or ambiguous should be listed and defined here.
7.0 Revision History	Dates, changes, and reasons are listed here. This ties in to enforcement in that the infraction should be measured against the rules in place at the time it occurred, not necessarily when it was discovered.

Your security policy defines the resources or assets that your organization needs to protect and the measures you must take to protect them. In other words, it is, collectively, the codification of the decisions that went into your security stance. Policies must be published and distributed to all employees, owners, and other users of your systems. Management should ensure and champion by support that everyone reads, understands, and acknowledges their role in following them and in the penalties that violations can bring.

As stated previously, you can trust everyone or trust no one; neither option works effectively when trying to balance productivity and security. Users all have differing views as to a network's security needs, and they all have a certain level of inherent fear. Users fear that their jobs might be more difficult as a result of security or that they might be punished if they make a mistake or forget to do something. Ultimately, people at any level do not like to feel restricted when they are trying to work. These kinds of attitudes are normal, emotional reactions, so they must be understood and appropriately managed for the security policy to provide balanced protection for your company. Building involvement, or "buy-in," in security policy development by including representatives from the areas listed in Table 2-3 is highly recommended.

Table 2-3 *Members of the Policy Review Team, Council, or Board (Partially from the SANS Security Policy Project)*

Representative From	Duties
Management	Someone who can enforce the policy. This is often a senior member of the HR staff.
Information Security Department	Someone who can provide technical insight and research.
User Areas	Someone who can view the policies the way a user might view them.
Legal Department	Possibly part time, but someone who can review policies with respect to applicable laws. For multinational firms, this review is exponentially more complicated.
Publications	Someone who can make suggestions on communicating the policies to the members of the organization and getting their buy-in. Also, a good writer is always helpful.

You can avoid the personal minefield if you ask the involved groups for their input as part of the policy development process. This enables you to do a little social engineering for the good folks by allowing these groups to participate in the process; they will more readily accept increased security restrictions in this case.

The following section reviews some actual security policies that we've used in the past and helps define how we write a security policy.

Acceptable Use Policy

SANS (www.sans.org) provides a wide range of security policies freely available on its website. These policies are based on these publicly available policies. Visiting SANS can complement what you learn from and implement based on this chapter. We will use a fictitious company called Granite Systems and show how it based its policies on those recommended by SANS.

In this policy, the company's IT security department is known simply as the *Corporate Security Team* for *Granite Systems*. *Granite Systems* and other Granite Systems-specific departments appear in *italics* throughout the policy; if you want to reuse this policy, you can replace these designations with your own.

Policy Overview

The *Corporate Security Team's* intentions for publishing an Acceptable Use Policy are not to impose restrictions contrary to *Granite Systems'* established culture of openness, trust, and integrity. *Corporate Security* is committed to protecting *Granite Systems'* employees, partners, and the company from illegal or damaging actions by individuals, either knowingly or unknowingly.

Internet/intranet/extranet-related systems, including but not limited to computer equipment, software, operating systems, storage media, network accounts providing electronic mail, WWW browsing, and FTP, are the property of *Granite Systems*. These systems are to be used for business purposes that serve the interests of the company, its clients, and its customers.

Effective security is a corporatwide team effort involving the participation and support of every *Granite Systems* employee, contractor, business partner, or any affiliates who deal with information and information systems. It is the responsibility of every computer user to know the guidelines contained within this security policy and to conduct their activities accordingly.

Purpose

The purpose of this security policy is to outline the acceptable use of computer equipment at *Granite Systems*. These rules are in place to protect the employee and *Granite Systems*. Inappropriate use exposes *Granite Systems* to risks, including but not limited to virus attacks, compromise of network systems and services, and legal issues.

Scope

This security policy applies to employees, contractors, consultants, temporaries, and other workers at *Granite Systems*, including all personnel affiliated with third parties. This policy applies to all equipment that is owned or leased by *Granite Systems*, to include personal equipment that might come in contact with the corporate IT infrastructure.

General Use and Ownership

1. Although *Granite Systems' Corporate Security Team* wants to provide a reasonable level of privacy, users should be aware that the data they create on the corporate systems remains the property of *Granite Systems*. Because of the need to protect *Granite Systems' network*, management cannot guarantee the confidentiality of information stored on any network device belonging to *Granite Systems*.
2. Employees are responsible for exercising good judgment about the reasonableness of personal use. Individual departments are responsible for creating guidelines concerning personal use of Internet/intranet/extranet systems. In the absence of such policies, employees should be guided by departmental policies on personal use and, if there is any uncertainty, employees should consult their supervisor or manager.
3. The *Corporate Security Team* recommends that any information that users consider sensitive or vulnerable be encrypted. For guidelines on information classification, see the *Corporate Security Team's Information Sensitivity Policy*. For guidelines on encrypting email and documents, go to Security Team's Awareness Initiative.

Note In many cases, you will see a security policy that references other policies within an organization. This is considered reasonable and considered a best practice. This enables you

to keep a policy specific to the topic at hand. Consider the preceding points, which reference encryption of data. Realistically, everyone within an organization must read and sign an acceptable use security policy; however, compare that to those who would be expected to encrypt data, a vastly different list and type of person. Thus, these policies are kept separate, thereby allowing or preventing confusion on the part of the user.

4. For security and network maintenance purposes, authorized individuals within *Granite Systems* may monitor equipment, systems, and network traffic at any time, per the *Corporate Security Team's Audit Policy*.
5. *Granite Systems* reserves the right to audit any and all networks and related systems on a periodic or ad hoc basis to ensure compliance with this policy.

Note Items 4 and 5 are chief. They enable your organization to notify all personnel that you can and will monitor and audit the network in all ways and on a regular, as-needed basis. It is crucial for these statements to be present because this enables employees to know that they will be watched in some fashion.

Security and Proprietary Information

1. The user interface for information contained on Internet/intranet/extranet-related systems should be classified as either confidential or not confidential, as defined by corporate confidentiality guidelines, the details of which can be found in the *Granite Systems* Human Resources policies. Examples of confidential information include but are not limited to the following:
 - Company private or confidential
 - Corporate strategies or projections
 - Competitor-sensitive or competitive analyses
 - Trade secrets, patents, test results
 - Specifications, operating parameters
 - Customer lists and data
 - Research data

Employees should take all necessary steps to prevent unauthorized access to this information. If an employee suspects that such information has been released outside the company, he should notify *Corporate Security* immediately.
2. Keep passwords secure and do not share accounts. Authorized users are responsible for the security of their own passwords and accounts. System-level passwords should be changed quarterly; user-level passwords should be changed every six months, but this might vary by organization requirements.

3. All PCs, laptops, and workstations should be secured with a password-protected screensaver with the automatic activation feature set at 10 minutes or less, or by logging off (Ctrl-Alt-Delete for WinXP users) when the host will be unattended.

Note The items discussed in 2 and 3 presuppose that best practices are being used. This means there is a dependency that servers require users to change passwords and that these passwords follow specific guidelines, as you will see later in the section, “Password Policy.”

4. Use of strong encryption of information in compliance with Corporate Security Acceptable Encryption Use policy.
5. Because information contained on portable computers is especially vulnerable, special care should be exercised. Protect laptops in accordance with the “Laptop Security Tips.”
6. Postings by employees from an *Granite Systems* email address to newsgroups should contain a disclaimer stating that the opinions expressed are strictly their own and not necessarily those of *Granite Systems*, unless posting is in the course of business duties.
7. All hosts used by the employee that are connected to the *Granite Systems* Internet/intranet/extranet, whether owned by the employee or *Granite Systems*, shall be continually executing approved virus-scanning software with a current virus database.

Note This portion of the policy reflects the strong trend of people checking email from multiple PCs and different physical locations. Consider an employee who might check his free web mail service at work and download a file that contains a virus without realizing it. The goal here is to ensure that, when at work, an approved virus checker catches this virus. However, if an employee accesses the same email from a home PC that she uses to connect to the corporate network, the vulnerability and ramifications should be closely considered.

8. Employees must use extreme caution when opening email attachments received from unknown senders that might contain viruses, email bombs, or Trojan horse code (malicious code). When in doubt, employees are advised to manually scan showing the original headers of the document and contact *Corporate Security* before opening them.

Unacceptable Use

The following activities are, in general, prohibited. Employees can be exempted from these restrictions during the course of their legitimate job responsibilities. (For example, systems administration staff might have a need to disable the network access of a host if that host is disrupting production services.)

Under no circumstances is an employee of *Granite Systems* authorized to engage in any activity that is illegal under local, state, federal, or international law while using *Granite Systems*-owned resources.

The lists that follow are by no means exhaustive, but they attempt to provide a framework for activities that fall into the category of unacceptable use. If an employee has any questions about the appropriateness of an action, he should contact *Corporate Security* for clarification.

System and Network Activities

The following activities are strictly prohibited, with no exceptions:

1. Violations of the rights of any person or company protected by copyright, trade secret, patent, or other intellectual property or similar laws or regulations, including, but not limited to the installation or distribution of “pirated” or other software products that are not appropriately licensed for use by *Granite Systems*.
2. Unauthorized copying of copyrighted material including, but not limited to digitization and distribution of photographs from magazines, books, or other copyrighted sources, copyrighted music, and the installation of any copyrighted software for which *Granite Systems* or the end user does not have an active license is strictly prohibited.
3. Exporting software, technical information, encryption software, or technology in violation of international or regional export control laws is illegal. The appropriate employee manager should be consulted prior to export of any material that is in question.

Note These first several instances are imperative for a security policy and an organization on many different levels. Consider probably the most vocal and legally active organizations on the Internet:

Recording Industry Association of America (www.riaa.org)

Report Cable Theft (www.cabletheft.com/)

Business Software Alliance (www.bsa.org/)

These organizations monitor theft, pirating, copyright violations, and so on, and prosecute those who engage in these activities. Individuals and businesses have been the primary legal targets of those engaged in this activity; they have been successful and are set to tackle educational institutions and the pirating that goes on from their campuses.

4. Introduction of malicious programs into the network or server (for example, malicious code including viruses, worms, Trojan horses, email bombs, and so on).
5. Revealing your account password to others or allowing use of your account by others. This includes family and other household members when work is done at home.

Note No one in the company will ever ask for your password. If a technical difficulty occurs, they will reset the password. Never reveal your password to anyone and, if asked, report the request to corporate security immediately.

6. Using a *Granite Systems* computing asset to actively engage in procuring or transmitting material that is in violation of sexual harassment or hostile workplace laws in the user's local jurisdiction.
7. Making fraudulent offers of products, items, or services originating from any *Granite Systems* account.
8. Making statements about warranty, expressly or implied, unless it is a part of normal job duties.
9. Effecting security breaches or disruptions of network communication. Security breaches include, but are not limited to, accessing data of which the employee is not an intended recipient or logging in to a server or account that the employee is not expressly authorized to access, unless these duties are within the scope of regular duties. For purposes of this section, "disruption" includes, but is not limited to, network sniffing, ping floods, packet spoofing, any denial of service, and forged routing information for malicious purposes.
10. Port scanning or security scanning (vulnerability assessment or penetration testing in wired or wireless networks) is expressly prohibited unless prior notification to *Corporate Security Team* or *authorized company executive* is made.
11. Executing any form of sanctioned network monitoring that will intercept data that is not intended for the employee's host, unless this activity is a part of the employee's normal job/duty.
12. Circumventing user authentication or security controls of any host, network, or account.
13. Interfering with or denying service to any user other than the employee's host (for example, any denial of service attack).
14. Using any program/script/command, or sending messages of any kind with the intent to interfere with or disable a user's terminal session via any means, locally or via the Internet/intranet/extranet.
15. Providing information about or lists of *Granite Systems* employees to parties outside *Granite Systems*.

Email and Communications Activities

1. Sending unsolicited email messages, including the sending of "junk mail" or other advertising material to individuals who did not specifically request such material (email spam).

2. Any form of harassment via email, telephone, or paging, whether through language, frequency, or size of messages.
3. Unauthorized use or forging of email header information and email encryption to obscure data in some cases.
4. Solicitation of email for any other email address, other than that of the poster's account, with the intent to harass or to collect replies.
5. Creating or forwarding "chain letters," "Ponzi," or other "pyramid" schemes of any type.
6. Use of unsolicited email originating from within *Granite Systems*' networks of other Internet/intranet/extranet service providers on behalf of, or to advertise, any service hosted by *Granite Systems* or connected via *Granite Systems*' network.
7. Posting the same or similar nonbusiness-related messages to large numbers of Usenet newsgroups (newsgroup spam or social networking site).

Enforcement

Any employee found to have violated this policy might be subject to disciplinary action, up to and including termination of employment and law enforcement inclusion if necessary.

Conclusion

Every security policy should end with a few common elements to clear up any potential miscommunication and confusion on the part of the users now that they understand what is permitted and what is not:

1. **Enforcement:** The main element is the enforcement and the ramifications to an employee if these policies are violated.
2. **Definitions:** Not every employee or user will understand some of the terminology used in a policy; therefore, it is a good idea to provide yet another level of clarification by defining industry-specific terms.
3. **Revisions:** Changes are always applied to policies such as these. The source of these changes alter with time; however, it might be a change in management, new laws, or perhaps a clarification of older laws, new threats against your network's security, your company has decided it wants to become certified (for example, ISO), or perhaps your company has new technology that needs to be covered. All these factors might require a policy change, and it is wise to document the changes.

Although these kinds of policies have a tendency to upset people who think they are *entitled* to something from their employer, they are not; they are there to contribute to the company's business goals. This fundamental truth enables the policy to protect the company, its employees, and everyone associated with it. Quoting from *Star Trek II: The Wrath of Khan*, "The needs of the many outweigh the needs of the few." Being one of a

few power users in my organization, I do not look forward to approving policies; however, it is the right thing to do for the company.

Password Policy

SANS (www.sans.org) provides a wide range of security policies freely available on its website. These policies are based on these publicly available policies. You should visit SANS and use discussions in this chapter to spark your ideas. Granite Systems (www.granitesystems.net) based these policies on those recommended by SANS and allowed me to present them here.

In this policy, the company's IT security department is known simply as the *Corporate Security Team* for *Granite Systems*. *Granite Systems* and other Granite Systems–specific departments will appear in *italics* throughout the policy; if you want to reuse this policy, you can replace these designations with your own.

Overview

Passwords are a crucial aspect of computer security. They are the first line of protection for user accounts. A poorly chosen password might result in the compromise of *Granite Systems*' entire corporate network. As such, all *Granite Systems* employees (including contractors and vendors with access to *Granite Systems* systems) are responsible for taking the appropriate steps for selecting and securing their passwords, as outlined in the following sections.

Purpose

The purpose of this policy is to establish a standard for the creation of strong passwords, the protection of those passwords, and to define how often you should change them.

Caution Passwords should be changed on a regular basis because user passwords are the first thing an attacker will try to crack. Most systems automatically prompt a user to change a password after a set amount of time has elapsed. Many of the newer operating systems apply some intelligence to a user's password, thus forcing the user not to use words that can be guessed or found in a dictionary. If you are not using these features or are not sure whether they are a part of your systems, it is a good idea to research the matter and activate them.

Scope

The scope of this policy includes all personnel who have or are responsible for an account (or any form of access that supports or requires a password) on any system that resides at any *Granite Systems* facility, has access to the *Granite Systems* network, or stores any nonpublic *Granite Systems* information.

Note An account can be defined and expanded to include email, keypad locks, FTP, shared drives, and so on. All passwords used to access these kinds of resources should follow some sort of password policy, as discussed in other portions of this policy.

General Policy

All system-level passwords (for example, root, enable, Windows admin, application administration accounts, and so on) must be changed on at least a quarterly basis.

- All production system-level passwords must be part of the Corporate Security Team's administered global password management database.

Note Not every organization has such a grandiose sounding "global password database" way of tracking passwords and, frankly, it is not necessary for most organizations. However, you must track passwords and how often they are changed in some manner. This enables you to ensure that your policy is being followed. Of course, ensure that you restrict access to whatever tool you put in place.

- All user-level passwords (for example, email, web, desktop computer, and so on) must be changed at least every six months. The recommended change interval is every four months.
- User accounts that have system-level privileges granted through group memberships or programs such as administrator or root must have a unique password from all other accounts held by that user.
- Passwords must not be inserted into email messages or other forms of electronic communication.
- Passwords must never be given out to anyone, regardless of their position in the company by email, voice, text or instant message. Employees are instructed to directly contact Corporate Security if anyone asks for your password, before giving it out.
- Where SNMP is used, the community strings must be defined as something other than the standard defaults of "public," "private," and "system," and must be different from the passwords used to log in interactively. A keyed hash must be used where available (for example, SNMPv2 and later).

Note This last part means changing the default passwords for the device in question. It is amazing how many organizations have never changed the default passwords. When I run across a device for which I do not know the default password, I always consult this site: www.cirt.net/cgi-bin/passwd.pl. At the time of this writing, there are more than 162 vendors with a total of 1132 default passwords and an ever-growing list for wireless devices and their passwords (SSID).

All user-level and system-level passwords must conform to the guidelines described in the following section.

General Password Construction Guidelines

Passwords are used for various purposes at *Granite Systems*. Some of the more common uses include user-level accounts, web accounts, email accounts, screensaver protection, voicemail password, and local router logins. Because few systems have support for one-time tokens (that is, dynamic passwords only used once), everyone should be aware of how to select strong passwords.

Poor, weak passwords have the following characteristics:

- The password contains less than eight characters.
- The password is a word found in a dictionary (English or foreign).
- The password is a common usage word such as the following:
 - Names of family, pets, friends, coworkers, fantasy characters, and so on.
 - Computer terms and names, commands, sites, companies, hardware, software.
 - The words “Granite Systems,” “energy,” “Granite Systems,” or any derivation.
 - Birthdays and other personal information such as addresses and phone numbers.
 - Word or number patterns such as aaabbb, qwerty, zyxwvuts, 123321, and so on.
 - Any of the previous words spelled backward.
 - Any of the previous words preceded or followed by a digit (for example, secret1, or 1secret).
 - Sports teams or famous players.
 - Do not post passwords anywhere.

Note Chapter 12 “Tools of the Trade,” discusses word lists and dictionaries; however, while discussing passwords, it is also appropriate to mention word lists and dictionaries in this chapter. A word list is simply a list of words, such as words from the dictionary, sports teams, industry terms, slang words, names; or all these lists are available in many different languages on the Internet. A good online source is <http://wordlist.sourceforge.net/>.

Attackers use these word lists as the basis of an attack, hoping someone would use a derivation of a word found on one of these lists. Just to be sure, they also inject numbers. This capability of attackers is the basis for the preceding portion of this policy.

Strong passwords have the following characteristics:

- Contain both upper- and lowercase characters (for example, a–z, A–Z)

- Have digits and punctuation characters and letters (for example, 0–9, !@#%&^&*()_+!~=-\`{}|:; '<>?,./)
- Are at least eight alphanumeric characters in length
- Are not words in any language, slang, dialect, jargon, and so on
- Are not based on personal information, names of family, and so on

Note Passwords should never be written down or stored online. Try to create passwords that you can remember easily. One way to do this is create a password based on a song title, affirmation, or other phrase. For example, the phrase might be: “This May Be One Way To Remember,” and the password could be “TmB1w2R!” or “Tmb1W>r-,” or some other variation. NOTE: Do not use either of these examples as passwords.

Password Protection Standards

Do not use the same password for *Granite Systems* accounts as for other non-*Granite Systems* access (for example, personal ISP account, option trading, benefits, and so on). Where possible, do not use the same password for various *Granite Systems* access needs. For example, select one password for the engineering systems and a separate password for IT systems. Also, select a separate password to be used for an NT account and a UNIX account.

Do not share *Granite Systems* passwords with anyone, including administrative assistants or secretaries. All passwords are to be treated as sensitive, confidential *Granite Systems* information.

Following is a list of “don’ts”:

- Do not reveal a password over the phone to ANYONE.
- Do not reveal a password in an email message.
- Do not reveal a password to the boss.
- Do not talk about a password in front of others.
- Do not hint at the format of a password (for example, “my family name”).
- Do not reveal a password on questionnaires or security forms.
- Do not share a password with family members.
- Do not reveal a password to coworkers while on vacation.

If someone demands a password, refer him to this document or have him call someone on the *Corporate Security Team*.

Do not use the “Remember Password” feature of applications. (For example, Outlook, web browsers, and so on—clear caches often dictated by procedures and guidelines.)

Again, do not write passwords down and store them anywhere in your office. Do not store passwords in a file on *any* computer system (including Palm Pilots or similar devices) without encryption that has been approved by the *Corporate Security Team*.

Change passwords at least once every six months (except system-level passwords, which must be changed quarterly). The recommended change interval is every four months.

If an account or password is suspected to have been compromised, report the incident to the *Corporate Security Team* and immediately change all passwords.

The *Corporate Security Team* or its delegates can perform password cracking or guessing on a periodic or random basis. If a password is guessed or cracked during one of these scans, the user will be required to change it.

Enforcement

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

Conclusion

Every security policy ends with a few common elements. These elements clear up all potential miscommunication and confusion on the part of the users, now that they understand what is permitted and what is not.

- 1. Enforcement:** The most essential element is the enforcement and the ramifications to an employee if these policies are violated.
- 2. Definitions:** Not every employee or user understands some of the terminology used in a policy; thus, it is always a good idea to provide yet another level of clarification by defining industry-specific terms.
- 3. Revisions:** Changes are always applied to policies such as these. The source of these changes alter with time, however; it might be a change in management, new laws, clarification of older laws, new threats against your network's security, your company has decided it wants to become certified (for example, ISO), or perhaps your company has new technology that needs to be covered. All these factors might require a policy change, and it is wise to document the changes.

Users always try to get around the restrictions placed on them via a password policy—no one likes to remember the cryptic passwords required in such a policy. If a user does not remember a password that meets these guidelines, do not worry—he will have to change it soon!

Unfortunately for users, they will have to remember and follow this policy. Password security is the first step in protecting your network; as such, beginning with the right expectations of your users helps to ensure that the overall security of your organization is preserved.

The next section examines a security policy targeted at virtual private networks (VPN) and what to look for to ensure their security.

Virtual Private Network (VPN) Security Policy

Chapter 9, “IPsec Virtual Private Networks (VPN),” covers VPNs in more detail; however, because this chapter covers security policies, the growth of VPNs in use today demands inclusion of a sample policy for VPNs here. This policy is prefaced by a brief definition of what a VPN is, but you should refer to Chapter 9 for the full scope of this technology.

VPNs are becoming popular and have matured considerably in the last several years. Many companies use them as a means of securely connecting small remote offices or users of every description. The connections can be made secure through the use of IPsec (IP Security) and L2TP (Layer 2 Tunneling Protocol) and with the increasing prevalence of high-speed Internet connections such as DSL or cable VPNs becoming affordable. Therefore, it becomes imperative to have a security policy to regulate their use so that all traffic is properly secured.

SANS (www.sans.org) provides a wide range of security policies freely available on its website. These security policies are based on these publicly available policies. I strongly encourage you to visit SANS and use the discussions in this chapter to spark your ideas. *Granite Systems* (www.granitesystems.net) based these policies on those recommended by SANS and have allowed me to present them here.

In this policy, the company’s IT security department is known simply as the *Corporate Security Team* for *Granite Systems*. *Granite Systems* and other Granite Systems–specific departments appear in *italics* throughout the policy; if you want to reuse this policy, you can replace these designations with your own.

Purpose

The purpose of this policy is to provide guidelines for Remote Access IPsec or L2TP virtual private network (VPN) connections to the *Granite Systems* corporate network.

Note VPNs based on IPsec are preferred over those using L2TP because they are generally considered more secure.

Scope

This policy applies to all *Granite Systems* employees, contractors, consultants, temporaries, and other workers, including all personnel affiliated with third parties that use VPNs to access the *Granite Systems* network. This policy applies to implementations of VPN that are directed through a VPN concentrator or VPN-aware firewall.

Policy

Approved *Granite Systems* employees and authorized third parties (customers, vendors, and so on) can use the benefits of VPNs, which are a “user-managed” service. This means that the user is responsible for selecting an Internet service provider (ISP), coordinating installation, installing any required software, and paying associated fees.

Note Although some companies might provide (that is, pay for) broadband or dial-up Internet connections for some of its employees, this is usually on a case-by-case basis. In general, companies leave that responsibility up to its employees, and that is, therefore, expressed in the corporate security policy.

In addition

1. It is the responsibility of employees with VPN privileges to ensure that unauthorized users are not allowed access to *Granite Systems* internal networks. VPNs may be used for site-to-site connectivity or remote access to systems or networks.
2. VPN use is to be controlled using either a one-time password authentication, such as a token device, or a public/private key system with a strong passphrase.
3. When actively connected to the corporate network, VPNs force all traffic to and from the PC over the VPN tunnel; all other traffic is dropped.
4. Dual (split) tunneling is NOT permitted; only one network connection is allowed.
5. *Split-tunneling* is a method of configuring a VPN and is either on or off. Essentially, if split-tunneling is on, users are allowed to simultaneously connect to the corporate network and the Internet. This presents a danger to the corporate network’s security because if an attacker were to take control of the computer creating a VPN to the corporate network, the attacker could also gain access to the company’s network via the VPN. It is therefore considered best practice to disable split-tunneling.
6. VPN appliances are set up and managed through *Granite Systems* network operational groups.
7. All computers connected to *Granite Systems* internal networks through VPN or any other technology must use the most up-to-date antivirus software that is the corporate standard and can be downloaded through the corporate intranet. This also includes personal computers.
8. VPN users are automatically disconnected from *Granite Systems*’ network after 30 minutes of inactivity. The user must then log in again to reconnect to the network. Pings or other artificial network processes are not to be used to keep the connection active.
9. Users of computers that are not *Granite Systems*–owned equipment must configure the equipment to comply with *Granite Systems*’ VPN and Network Security policies.

10. Only VPN clients approved by the *Corporate Security Team* can be used.
11. By using VPN technology with personal equipment, users must understand that their machines are a de facto extension of *Granite Systems'* network and, as such, are subject to the same rules and regulations that apply to Granite Systems–owned equipment; that is, their machines must be configured to comply with all Corporate Security Policies.

Conclusion

Every security policy should end with a few common elements; these elements clear up all potential miscommunication and confusion on the part of the users now that they understand what is and is not permitted:

1. **Enforcement:** The element that is most critical is the enforcement and the ramifications to an employee if these policies are violated.
2. **Definitions:** Not every employee or user understands some of the terminology used in a policy; thus, it is always a good idea to provide yet another level of clarification by defining industry-specific terms.
3. **Revisions:** Changes are always applied to policies such as these. The source of these changes alter with time, however; it might be a change in management, new laws, or perhaps a clarification of older laws, new threats against your network's security, your company has decided it wants to become certified (for example, ISO), or perhaps your company has new technology that needs to be covered. All these factors might require a policy change, and it is wise to document the changes.

VPN technology is ever-evolving, faster than most from a network security perspective. As discussed, businesses are deploying VPNs in ever-increasing numbers; therefore, it is crucial that all organizations have policies governing their use. If there is a mistake with a VPN, the consequences can be costly from both a security and financial perspective. Auditing VPN access should be a critical part of your process and larger governance policy.

Wireless Communication Policy

The purpose of this policy is to secure and protect the information assets owned by *Granite Systems*. *Granite Systems* provides computer devices, networks, and other electronic information systems to meet missions, goals, and initiatives.

Granite Systems grants access to these resources as a privilege and must manage them responsibly to maintain the confidentiality, integrity, and availability of all information assets preventing disclosure, alteration, and destruction.

This policy specifies the conditions that wireless infrastructure devices must satisfy to connect to *Granite Systems* network. Only those wireless infrastructure devices that meet the standards specified in this policy or are granted an exception by the Information Security Department are approved for connectivity to the *Granite Systems* network.

Scope

All employees, contractors, consultants, and temporary and other workers at *Granite Systems*, including all personnel affiliated with third parties that maintain a wireless infrastructure device on behalf of *Granite Systems*, must adhere to this policy. This policy applies to all wireless infrastructure devices connected to a *Granite Systems* network or reside on a *Granite Systems* site that provides wireless connectivity to endpoint devices including, but not limited to, laptops, desktops, smart phones, and personal digital assistants (PDAs). This includes any form of wireless communication device capable of transmitting packet data in 802.11 and Bluetooth networks, spectrums, and channels.

The Information Security Department must approve exceptions to this policy in advance.

Policy Statement

General Network Access Requirements

All wireless infrastructure devices that reside at a *Granite Systems* site and connect to a *Granite Systems* network, or provide access to information classified as *Granite Systems* Confidential, *Granite Systems* Highly Confidential, or *Granite Systems* Restricted must

- a. Abide by the standards specified in the Wireless Communication Standard
- b. Be installed, supported, and maintained by an approved support team
- c. Use *Granite Systems* approved authentication protocols and infrastructure
- d. Use *Granite Systems* approved wireless encryption protocols
- e. Maintain a hardware address (MAC Address) that can be registered and tracked aiding in MAC spoofing controls
- f. Not interfere with wireless access deployments maintained by other support organizations

Lab and Isolated Wireless Device Requirements

All lab and wireless infrastructure devices that provide access to *Granite Systems* Confidential, *Granite Systems* High Confidential, or *Granite Systems* Restricted information must adhere to the guidelines specified in the section “General Network Access Requirements” and isolated wireless devices that do not provide general network connectivity to the *Granite Systems* network must

- a. Be isolated from the corporate network (that is, it must not provide any corporate connectivity) and comply with the DMZ Lab Security Policy or Internal Lab Security Policy.
- b. Not interfere with wireless access deployments maintained by other support organizations.

Home Wireless Device Requirements

- a. Wireless infrastructure devices that provide direct access to the *Granite Systems* corporate network must conform to the Home Wireless Device Requirements as detailed in the Wireless Communication Standard.
- b. Wireless infrastructure devices that fail to conform to the Home Wireless Device Requirements must be installed in a manner that prohibits direct access to the *Granite Systems* corporate network. Access to the *Granite Systems* corporate network through this device must use standard remote access authentication.

Enforcement

An employee found to have violated this policy may be subject to disciplinary action, up to and including termination of employment. A violation of this policy by a temporary worker, contractor, or vendor might result in the termination of their contract or assignment with *Granite Systems*.

Definitions

Term	Definition
Granite Systems network	A wired or wireless network including indoor, outdoor, and alpha networks that provide connectivity to corporate services.
Corporate connectivity	A connection that provides access to a Granite Systems network.
Enterprise Class Teleworker (ECT)	An end-to-end hardware VPN solution for teleworker access to the Granite Systems network.
Information assets	Information that is collected or produced and the underlying hardware, software, services, systems, and technology that is necessary for obtaining, storing, using and securing that information that is recognized as important and valuable to an organization.
MAC Address	The MAC address is a hardware number that uniquely identifies each node on a network and is required for every port or device that connects to the corporate network.

Revision History

Date of Change	Responsible Owner	Summary of Change(s)

The next section covers the security policy that is necessary when corporate business partners or other third parties need to connect to your organization's network—a sensitive situation, indeed.

Extranet Connection Policy

This security policy deals with “how to handle” and “the requirements” necessary for those not affiliated with your organization to connect to and access resources on the network.

The “who’s” and “why’s” behind such a request vary greatly and, when considering them, you should review the section on trust in Chapter 1, “There Be Hackers Here,” before making a decision. Requests will come to you from the following parties:

- Contractors/consultants trying to do legitimate work with your company
- Business partners of all sorts
- Customers, usually large and requiring special handling

This security policy provides the necessary guidelines for answering such requests and the requirements to be placed on the requestor. It also enables the members of the IT staff to deal with pushy and insistent people, making this policy a virtual panacea.

SANS (www.sans.org) provides a wide range of security policies freely available on its website. These policies are based on these publicly available policies. You should visit SANS and use the discussions in this chapter to spark your ideas. Granite Systems (www.granitesystems.net) based these policies on those recommended by SANS and allowed the policies to be presented here.

In this policy, the company's IT security department is known simply as the *Corporate Security Team* for *Granite Systems*. *Granite Systems* and other Granite Systems–specific departments appear in *italics* throughout the policy; if you want to reuse this policy, you can replace these designations with your own.

Purpose

This document describes the policy under which third-party organizations or consultants connect to the *Granite Systems* network for the purpose of conducting business related to *Granite Systems*.

Scope

Regardless of whether a dedicated telecommunications circuit (such as frame relay or ISDN), broadband, or VPN technology is used for the connection, connections between third parties that require access to nonpublic *Granite Systems* resources fall under this policy. Connectivity to third parties, such as Internet service providers (ISP) that provide Internet access for *Granite Systems* or to the Public Switched Telephone Network (PSTN) do *not* fall under this policy.

Note Some clarification is warranted for that last part, where the policy seems to make an exception for the corporate Internet access and telephone usage through the PSTN. These

are excepted because they are commodities purchased by your company; as such, if you requested that the phone company follow this policy prior to getting telephones, trust me, you would never get any results.

Security Review

All new extranet connectivity will go through a security review with the *Corporate Security Team*. The security review ensures that all access matches the business requirements in the best possible way, and that the principle of *least access and privilege* is always followed.

Third-Party Connection Agreement

All new connection requests between third parties and *Granite Systems* require that the third-party and *Granite Systems* representatives agree to and sign the *Third-Party Agreement*. This agreement must be signed by the *Senior Vice President* of the *Sponsoring Organization* and a representative from the third party who is legally empowered to sign on behalf of the third party. The signed document is to be kept on file with the company's *Legal Department* and *Corporate Security Department*.

Business Case

All production extranet connections must be accompanied by a valid business justification, in writing, that is approved by the *Senior Director of Corporate Security/CI/ISO*. Included in this business case is the identification of the network resources that are requesting to be accessed.

Point of Contact

The *Granite Systems Sponsoring Organization* must designate a person to be the point of contact (POC) for the extranet connection. The POC acts on behalf of the *Sponsoring Organization* and is responsible for those portions of this policy and the *Third-Party Agreement* that pertain to it. If the POC changes, the relevant extranet organization must be informed promptly.

Establishing Connectivity

Sponsoring Organizations within *Granite Systems* that want to establish connectivity to a third party are to file a new site request with the *Corporate Security* team. The sponsoring organization engages the *Corporate Security Team* to address security issues that are inherent in the project. The *Sponsoring Organization* must provide full and complete information as to the nature of the proposed access to the extranet group and *Security Team*, as requested.

All established connectivity must be based on the least-access principle, in accordance with the approved business requirements and the security review. In no case does *Granite Systems* rely upon the third party to protect *Granite Systems'* network or resources.

Modifying or Changing Connectivity and Access

All changes in access must be accompanied by a valid business justification and are subject to security review. Changes are to be implemented via corporate change management process. The *Sponsoring Organization* is responsible for notifying the *Corporate Security Team* when there is a material change in their originally provided information so that security and connectivity evolve accordingly.

Terminating Access

When access is no longer required, the *Sponsoring Organization* within *Granite Systems* must notify the extranet team responsible for that connectivity; this terminates the access. This might mean a modification of existing permissions up to terminating the circuit, as appropriate. The *Corporate Security Teams* must conduct an audit of their respective connections annually to ensure that all existing connections are still needed and that the access meets the needs of the connection. Connections that are deprecated and are no longer being used to conduct *Granite Systems* business are terminated immediately. Should a security incident or a finding that a circuit has been deprecated and is no longer being used to conduct *Granite Systems* business necessitate a modification of existing permissions or termination of connectivity, the *Security Team* notifies the POC or the *Sponsoring Organization* of the change before taking any action.

Conclusion

Every security policy should end with a few common elements. These elements clear up all potential miscommunication and confusion on the part of the users now that they understands what is and is not permitted:

- 1. Enforcement:** The most important element is the enforcement and the ramifications to an employee if these policies are violated.
- 2. Definitions:** Not every employee or user understands some of the terminology used in a policy; thus, it is always a good idea to provide yet another level of clarification by defining industry-specific terms.
- 3. Revisions:** Changes are always applied to policies such as these. The source of these changes alter with time, however; it might be a change in management, new laws, or perhaps a clarification of older laws, new threats against your network's security, your company has decided it wants to become certified (for example, ISO), or perhaps your company has new technology that needs to be covered. All these factors might require a policy change, and it is wise to document the changes.

It is always a touchy subject to grant such access to those outside your company. One of the things that happens is that employee A works with business partner Z, who needs to access some resource on your network; to complete the business, employee A promises partner Z access. Alternatively, it is someone in management that makes a promise.

These scenarios are common, and this policy helps ensure that, if such a requirement is needed, the proper due diligence is taken before making any promises given this established process.

Perhaps the fastest growing certification authority is the International Standards Organization (ISO). The following section briefly discusses how ISO has entered into the security arena. It is fitting to bring it to your attention because more and more companies are becoming ISO-certified to one degree or another.

ISO Certification and Security

Compliance with any internationally recognized standards is becoming more necessary. As a result, and because standards relevance is a common currency of instant legitimization, many companies are pursuing such a course. The ISO offers many standards, and all are valuable in their own right. You can find a lot of useful information on ISO standards, the processes used, and the implementation of those standards at www.iso.org. For purposes of this discussion, the concern lies with standard ISO/IEC 27002: Information Technology Security Techniques Code of Practice for Information Security Management.

Delivery

When delivering the security policy to users, you must then determine the most effective communication manner in which to present them to help facilitate compliance and support from your users. This is often much easier said than done.

Many discussions on the concepts and goals of security policies always seem to gloss over the delivery of these policies, especially when they are business-focused, nontechnical users. Yet it is crucial for everyone to understand and support these policies. To not reach for this goal and to make the effort dooms the policy to failure and backlash from users because they will resent the policy from the beginning.

Handling these types of situations is similar to handling interpersonal relationships. Beyond good interpersonal skills, consider the following additional suggestions:

- Ensure that all policies are presented clearly during new employee orientation.
- Always allow a sample of the personnel affected by a security policy to review it and provide input comment before implementing.
- Provide a security policy refresher course and delivery methodology.

In general, you should keep policies short, fewer than two pages. There is no need to complicate the situation. Occasionally, you might have to go over, but not usually. In closing,

ensure that your policies are updated annually, if not sooner, to reflect the changes of the past year.

ISO/IEC 27002

ISO/IEC 27002 is an information security standard published by the International Organization for Standardization (ISO) and by the International Electro-technical Commission (IEC). It was originally published as ISO/IEC 17799:2005. In July 2007, it was redesignated ISO/IEC 27002:2005, bringing it inline with other 27000-series standards.

This de facto standard is extremely comprehensive in its security coverage, providing you with best practice recommendations on information security management for implementing and maintaining an Information Security Management System (ISMS).

ISO/IEC 27002 contains a significant number of information security controls arranged into 12 different areas. The information security controls are considered best practice means of achieving those objectives:

- **Risk assessment:** Includes risk management
- **Security policy:** Management direction and support
- **Organization of information security:** Governance of information security
- **Asset management:** Inventory and classification of information assets
- **Human resources security:** Security aspects for employees joining, moving, and leaving an organization
- **Physical and environmental security:** Protection of the computer facilities
- **Communications and operations management:** Management of technical security controls in systems and networks
- **Access control:** Restriction of access rights to networks, systems, applications, functions, and data
- **Information systems acquisition, development, and maintenance:** Building security into applications
- **Information security incident management:** Anticipating and responding appropriately to information security breaches
- **Business continuity management:** Protecting, maintaining, and recovering business-critical processes and systems
- **Compliance:** Ensuring conformance with information security policies, standards, procedures, guidelines, laws, and regulations

As the title suggests, these are international standards and as such they have equivalent standards across the globe: AS/NZS ISO/IEC 27002:2006 in Australia, JIS Q 27002 in Japan, and BS ISO/IEC 27002:2005 in the United Kingdom, just to name a few.

The ISO certification is briefly discussed here, but the standard is perhaps one of the most comprehensive and will be growing in use. To learn more, visit the ISO website at www.iso.org.

Sample Security Policies on the Internet

The policies presented here are simply one means to meet an organization's needs; what works well for one organization might not be ideal for another. Thus, you should refer to the following additional resources on security policies:

- www.sans.org/reading_room/whitepapers/policyissues/: This site contains articles and papers written by GIAC-certified professionals.
- www.ietf.org/rfc/rfc2196.txt: The Site Security Policies Procedure Handbook.
- www.assurityriver.com/securityalerts-05052005.shtml: A discussion on why security policies fail.

Some general websites with information security policies include the following:

- www.security.kirion.net/securitypolicy/
- www.utoronto.ca/security/documentation/policies/policy_5.htm
- <http://doit.missouri.edu/security/>
- www.windowsecurity.com/whitepapers/ - for Microsoft specific security related items
- <https://security.berkeley.edu/policies.html>
- www.ruskwig.com/security_policies.htm

If you want to be overwhelmed, go to your favorite search engine and search on security policy templates. When I did it I got more than 20.9 million results. The information is out there, reader. All you have to do is look.

Industry Standards

After you get out of the general corporate security policy doldrums, you can now begin to focus on the standards set forth by other governing bodies, such as DISA, NIST, or the PCI-DSS/SSC. We have focused on just a few here that seem to be hot-button topics for clients and lawyers alike. The first question someone will ask if there is a problem or an issue is, "Were you conforming to industry standards and best practices?" Your answer had better be a resounding YES!

Following are specific regulations addressed by industries:

- **Financial Services:** Gramm-Leach-Bliley Act (GLBA), Sarbanes-Oxley Act (SAR-BOX), USA Patriot Act, PCI Data Security Standard (PCI DSS), and the Basel II Accord (EU)

- **Healthcare and Pharmaceuticals:** Health Insurance Portability and Accountability Act of 1996 (HIPAA) and FDA 21 CFR Part 11
- **Infrastructure and Energy:** Guidelines for FERC and NERC Cybersecurity Standards, the Chemical Sector Cyber Security Program, and Customs-Trade Partnership Against Terrorism (C-TPAT)
- **Federal Government:** Compliance with FISMA and related NSA Guidelines and NIST Standards
- **Security Methodologies:** Security and control frameworks such as ISO 1-7799, COSO, and COBIT
- **Consumer Protection and Data Privacy:** Children’s Online Privacy Protection Act (COPPA), Children’s Internet Protection Act (CIPA), CAN-SPAM - Federal law about unsolicited electronic mail, Bill C-6: personal information protection and electronic documents Act (Canada), California Individual Privacy Senate Bill - SB1386, and MA State Law CMR 17.99

Payment Card Industry Data Security Standard (PCI DSS)

This is a worldwide information security standard defined by the Payment Card Industry Security Standards Council (PCI SSC). It was put in place to prevent credit card fraud through increased controls around data and its exposure to external threats.

PCI DSS began as five separate but similar programs from the “Big Five”: Visa, MasterCard, American Express, Discover, and JCB Data Security Program. The PCI SSC was formed to standardize the industry security practices and on December 15, 2004, the PCI DSS was released.

In July 2009, the PCI SSC published the wireless guidelines for PCI DSS recommending the use of Wireless Intrusion Prevention Systems (WIPS) to automate wireless scanning for large organizations. These guidelines apply to the deployment of wireless LAN in cardholder data environments.

The current version of the standard is version 2.0; as of October 2010, it sets forth 12 requirements for compliance, organized into six logically related groups called control objectives.

To learn more on the PCI standard go to <https://www.pcisecuritystandards.org/> as referenced earlier in this chapter.

Sarbanes-Oxley Act of 2002 (SOX)

Enacted July 30, 2002, this is also known as the “Public Company Accounting Reform and Investor Protection Act” and “Corporate and Auditing Accountability and Responsibility Act.” It set new, or better defined, standards for all U.S. public company boards, management, and public accounting firms. It was enacted as a reaction to a number of major corporate scandals (Enron, Tyco International, Adelphia, WorldCom, and so on). It does not apply to privately held companies.

Sarbanes-Oxley contains 11 titles that outline specific mandates and requirements for financial reporting: 1) Public Company Accounting Oversight Board, 2) Auditor Independence, 3) Corporate Responsibility, 4) Enhanced Financial Disclosure, 5) Analyst Conflicts of Interest, 6) Commission Resources and Authority, 7) Studies and Reports, 8) Corporate and Criminal Fraud Accountability, 9) White Collar Crime Penalty Enhancement, 10) Corporate Tax Returns, and 11) Corporate Fraud Accountability.

Health Insurance Portability and Accounting Act (HIPAA) of 1996

The HIPAA Act was put in place to protect you and your family during times of crisis when you lose your job, and it put in place (in Title II) Administrative Simplification (AS) provisions. This is the requirement to establish national standards for electronic health care transactions and national identifiers for providers, health insurance plans, and employers. This AS provision also addresses security and privacy of health care data.

The Final Rule on Security Standards was issued on February 20, 2003, taking effect on April 21 of that same year and a compliance date of no later than 2005. The security rule deals specifically with Electronic Protected Health Information (EPHI) and it lays out three types of security safeguards required for compliance: 1) administrative, 2) physical, and 3) technical. Each safeguard contains various standards, and for each standard it lists required and addressable implementation guidelines:

- **Administrative safeguards:** Policies and procedures designed to clearly show how the entity will comply with the act
- **Physical safeguards:** Controlling physical access to protect against inappropriate access to protected data
- **Technical safeguards:** Controlling access to computer systems and enabling covered entities to protect communications containing Protected Health Information (PHI) transmitted electronically over open networks from being intercepted by anyone other than the intended recipient

Massachusetts 201: Standards for the Protection of Personal Information of Residents of the Commonwealth

This regulation establishes minimum standards to be met in connection with the safeguarding of personal information contained in both paper and electronic records. I want to focus on one section in particular, specifically section 17.04: Computer Systems Security Requirements.

Section 17.04 establishes eight elements that each computer system containing personal information must have, as follows:

1. Secure user authentication protocols.
2. Secure access control measures.

3. Encryption of all transmitted records and files containing personal information that will traverse public networks, and encryption of all data containing personal information to be transmitted wirelessly.
4. Reasonable monitoring of systems, for unauthorized use of or access to personal information.
5. Encryption of all personal information stored on laptops or other portable devices.
6. For files containing personal information stored on a system that is connected to the Internet, there must be reasonable up-to-date firewall protection and operating system security patches, reasonably designed to maintain the integrity of the personal information.
7. Reasonably up-to-date versions of system security agent software, which must include malware protection and reasonably up-to-date patches and virus definitions, or a version of such software that can still be supported with up-to-date patches and virus definitions, and is set to receive the most current security updates on a regular basis.
8. Education and training of employees on the proper use of the computer security systems and the importance of personal information security.

Compliance with MASS 201 is mandatory for every person who owns or licenses personal information about a resident of the Commonwealth on or before March 1, 2010.

If you'd like to read the regulation in its entirety, you can find it here:
www.mass.gov/Eoca/docs/idtheft/201CMR1700reg.pdf.

SAS 70 Series

Statement on Auditing Standards (SAS) No. 70, or SAS 70, is an industry-recognized standard published by the American Institute of Certified Public Accountants (AICPA). SAS 70 provides third-party validation of the internal controls of service organizations, enabling them to disclose control activities and processes to their customers and auditors in a constant and uniform format.

The SAS 70 standard does not specify a required set of control objectives. So, making sure you have well-written, concise security policies will be a boon when your organization gets audited. A significant component of the SAS 70 audit involves the evaluation of an organization's information security controls.

Chapter Summary

This chapter discussed what many view as simply paperwork when, in reality, a security policy reflects your company's commitment to security. It included key concepts in writing a security policy, such as determining who and what to trust and who to involve in the writing and crafting of a security policy.

This chapter also presented a variety of sample security policies. These security policies reflect the current trends and major areas upon which companies can improve.

Specifically, these areas include what is considered acceptable use of corporate IT resources, how to ensure that you have effective passwords, when and how to use VPNs, and what restrictions to use when connecting your corporate network to a business partner's network.

Chapter 3, "Processes and Procedures," discusses the use of technologies that have evolved to support and enhance network security. Many of these technologies are used today without you understanding when or where they operate. After reading this chapter, you should understand the benefits of these technologies, where they operate, and some of their associated risks.

Chapter Review

1. How important is it to involve other departments and employees in the crafting of security policies?
2. True or false: It is a well-known fact that users circumvent security policies that are too restrictive. Explain your answer.
3. What are three things that you should keep in mind when writing or reviewing a security policy?
4. Why is it important to include an enforcement section in every security policy?
5. An Acceptable Use Policy defines what kind of expectations for users?
6. When and under what circumstances should you reveal your password to someone?
7. Which of the following sample passwords would be considered effective when checked against the corporate password policy?
 - a. wolfpack
 - b. thomas67
 - c. simonisnot4
 - d. sj8Dtt&efs
 - e. Missing\$4u
8. Define VPN and the role it can play within a company's network infrastructure.
9. VPNs support a technology called *split tunneling*. Define this technology and explain whether it should be used in a network.
10. How frequently should security policies be updated or reviewed?

This page intentionally left blank

Processes and Procedures

“...There’s a time for daring and there’s a time for caution, and a wise man understands which is called for...” —Dead Poets Society

“...All he’d wanted were the same answers the rest of us want. Where did I come from? Where am I going? How long have I got?” —Blade Runner

By the end of this chapter, you should know and be able to explain the following:

- The processes for managing and responding to security advisories within your organization
- Which organizations produce security advisories
- What a zero-day alert is and how you should respond
- Best practices for handling updates
- Define an Access Control List (ACL) and how to use one within a networking environment

Answering these key questions will enable you to understand the overall characteristics and importance of the processes and procedures used in the day-to-day life of a network security technician. By the time you finish this book, you will have a solid appreciation for network security, its issues, how it works, and why it is important.

How do you relate these quotes to security process management? Every user within an organization needs to easily see within a process where they are, how they got there, what they are supposed to do, what is going to happen next, and how long they have to complete the step. Furthermore, you need to delicately balance the implementation of the workflow (that is, processes and procedures) between being functional and secure. You could easily secure your computer systems too much, making research and development nonfunctional. You can also go the other direction and not secure your systems and be harassed by viruses and malware.

This chapter covers the options available to you as a network/technology security specialist within your organization, establishing change control boards, responding to threats, and then finally touching on some of the best practices out there in the technology field today.

Security Advisories and Alerts: Getting the Intel You Need to Stay Safe

You need to ensure a high level of operational security for the many assets within your organization: routers, switches, servers, laptops and desktops (Apple and Microsoft), BlackBerrys, smartphones, and so on. Pervasive threats exist, from a disgruntled employee to corporate espionage, as shown in Figure 3-1. This chapter deals with a select segment, information assurance, and how updates are managed within your organization.

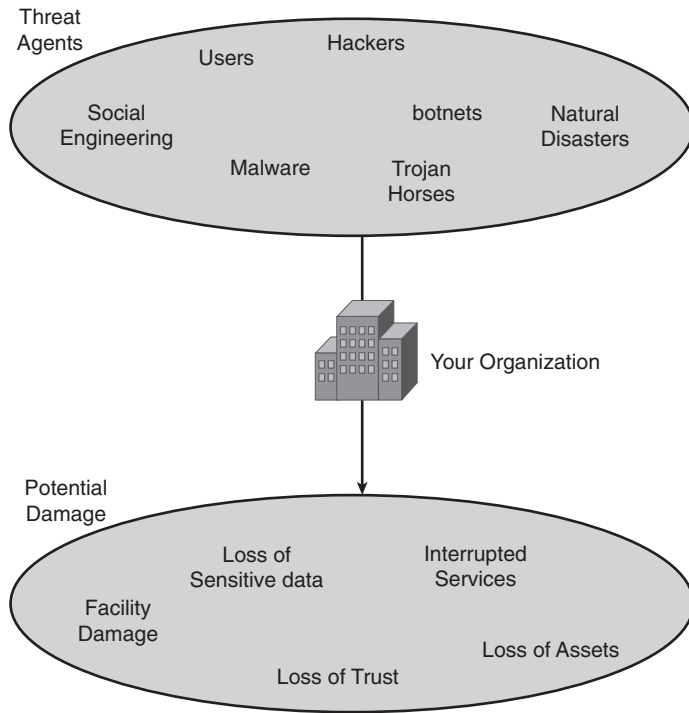


Figure 3-1 *Threat Agents*

Before you respond to security threats, you must be able to identify them. A threat agent exploits a vulnerability in an effort to cause harm to a computer, network, or company, ultimately in an attempt to impact your organization's capability to do business. Many types of threat agents can take advantage of several types of vulnerabilities. The potential damage listed in Figure 3-1 represents only a sampling of the risks many organizations should address in their risk management programs. Some threats are easier to identify than others. For instance, there may be a coding problem in a newly developed program your

company is creating whereby the application uses complex equations to produce results; however, if the equations are incorrect or if the application incorrectly uses the data, this could cause a cascading error as invalid results are passed from one process to another. These types of issues lie within the application's code and are hard to identify. Other threats, such as user error (intentional or accidental), are much easier to spot. However, you still must do your legwork. You must monitor and audit user activity on a continual basis. You must conduct audits and reviews to discover whether employees are misbehaving. Their malicious activity just doesn't show up in your email inbox on a Friday morning. You must put actionable policies and procedures in place to be proactive.

After you identify the vulnerabilities and threats, you must consider the results of those vulnerabilities. What are the risks and what is the loss potential of those risks? Following is the definition for loss potential:

“What the company would lose if a threat agent was actually to exploit a vulnerability (Harris, 2008).”

These losses may manifest themselves as corrupted data, destruction of systems, loss of confidential information, that is, corporate espionage, and unproductively of employees. This may result in destruction, alteration, and disclosure (DAD) of sensitive information that could damage a corporate brand. Asset valuation to determine the single loss expectancy (SLE) and annual loss expectancy (ALE) will help tag assets and help an organization classify its value.

Some of those threats (malware, botnets, viruses, Trojan horses, and cyber attacks from hackers or insiders) are more easily combated than others. You fight these nefarious agents by keeping your computing environment up to date and your users educated through consistent user awareness training on current policy, standards, procedures, and guidelines. When a threat is identified, you must respond. The how, the why, and the where to look are all addressed next. Keep reading...this gets good!

Responding to Security Advisories

So, how should you respond when you are notified there is potential risk to your organization? What are the procedures? Whom do you contact? Will the risk jeopardize all, or just a few, of the individuals in your organization? Can you afford the risk? These are all questions that senior management is going to want answers to. You could establish a Group Policy (GPO) in your organization (assuming you are running a Microsoft environment) whereby each system updates itself by going out to the Internet and contacting the Microsoft security updates website and downloading all critical updates. But now you have a different risk, don't you? This is a tricky topic. It is hard to establish a default answer to how you should respond to security advisories as they come out. I'll try to provide you with a basic framework you can use within your organization. You may add to, or remove from, it as you want and as it fits within your organization.

The environment or things that you should have in place include the following:

- A dedicated and up-to-date security policy (see Chapter 2, “Security Policies”).

- A chief information security officer (CISO) who is more than a paper tiger. He/she needs to have a budget, authority, and support of the executive team.
- A change control board (CCB) and procedures; these are discussed later in this chapter.
- A test bed or lab consisting of routers, switches, servers, and client workstations. These must match your current and future corporate technical environment.
- A Windows Server Update Services (WSUS) to manage critical updates.

Now that you have your ideal environment, consider five steps that need to take place. This may not fit you and your organization; cut and choose to make it fit. As you cut to streamline the processes, you may cut too much. Adjusting this response framework is similar to cutting hair...just take a little at a time; you can't put it back after it's off:

Step 1. Awareness

Step 2. Incident response (protected immediately, or can it wait?)

Step 3. Imposing your will

Step 4 and 5. Test patches and push patches

The steps of this framework contain a lot of information and procedures that we touch on but not go into too much depth because that is outside the scope of this book. Volumes are written on risk management and change management procedures. I just want to educate you on their existence, what these things are, and briefly discuss their importance within your organization.

Step 1: Awareness

To fix any problem, whether it's personal or business-related, you first must know there is a problem. For example, when I was a young airman in the USAF, I got into a mess during an inspection. I was supposed to be following a procedure that I didn't know existed until the inspector came and asked me about a program I had been running. The excuse I told my commanding officer was, "Sorry sir...I didn't know." His response to that was I should have known. It was my responsibility to know. Luckily, despite my failure, our unit did manage to pass the inspection. My point is this—ignorance is never an excuse. If you are in charge of a program, building a house, or running a corporation's information security team, you need to know what your job is about. You need to be better informed than the bad guys. That is what this portion is about: being aware. Several useful sites are available to help you stay abreast of what's going on.

The sections that follow describe the security advisories of Cisco, Apple, and Microsoft because each is a leader in their industry: routing, switching/Internet connectivity, servers, and desktop computing. That doesn't mean that the other organizations are any less informed; a lot of information is available from the Common Vulnerabilities and Exposures (CVE) database to the Defense Information Systems Agency (DISA) Information Assurance web page. (See the "Chapter Review" section.) We encourage you to begin there.

Cisco Security Advisories

The following is from the Cisco website.

“...Cisco releases bundles of IOS Security Advisories on the fourth Wednesday of the month in March and September of each calendar year. This does not restrict us from promptly publishing an individual IOS Security Advisory for a serious vulnerability which is publicly disclosed or for which we are aware of active exploitation. All other non-IOS Cisco security vulnerabilities will continue to be announced per the Cisco standard disclosure policy. You can receive them for free being more proactive just sign up for the email notifications and find the link here:

http://www.cisco.com/en/US/products/products_security_advisories_listing.html.”

“...Starting in January 2011, Cisco will be providing additional information, available through the Cisco Bug Toolkit, on all bugs reviewed by the Cisco Product Security Incident Response Team (PSIRT)...”

Apple Security Advisories

The following is from Apple’s website:

Apple provides multiple ways for the end-user/administrator to keep on top of vast amount of security updates required. For the protection of their customers, Apple does not disclose, discuss or confirm security issues until a full investigation has occurred and any necessary patches or releases are available. Apple usually distributes information about security issues in several ways: 1) its products, 2) through the Apple security website, and 3) a mailing list.

1. General product information

General information about Apple products is made available at Apple’s website (<http://www.apple.com>). This information includes product documentation, as well as technical papers, hints, tips, and questions and answers. Notifications developed by Apple are signed with the Apple Product Security PGP key. Apple encourages their customers to verify the signature to ensure that the document was indeed written by Apple staff and has not been changed.

You can verify the signature by going to the following website,

<https://www.apple.com/support/security/pgp/>

2. Updates

Check the Apple Security Updates page for released updates by going to <http://support.apple.com/kb/HT1222>.

3. Mailing list

The Security-Announce mailing list provides another way Apple provides customer support and information to its customers allowing them to obtain product security information from Apple. You can subscribe via <http://lists.apple.com/>, also available via RSS Feed at <http://rss.lists.apple.com/>.

Microsoft Security Bulletins

Microsoft’s security-focused website (www.microsoft.com/security/default.aspx) provides links to all things security whether you are a casual home user or a developer. The section you are probably most familiar with is the page for downloading the critical

updates. But there is so much more. Microsoft products provide a means to let the end user know when a critical update is ready for downloading and installation by means of Windows Update service. Windows Update is a free service built in to Windows. It is designed to help you keep your computer more secure, reliable, compatible with devices, and able to run new features that might enhance your computing experience. Windows Update enables you to easily get what your computer needs, such as the following:

- The latest security updates to protect against malware and other potentially unwanted software
- Updates that improve reliability and performance
- Upgrades to Windows features
- Drivers from Microsoft partners

Although Windows Update needs to check your computer to determine which updates it needs, it does not collect your personal information. Windows Update simply checks to see what software and hardware is installed so that it knows what updates you need. For home users this is a valid option. Home users don't have specific development environments that must be maintained; however, I do not suggest you, as the security guru of your organization, allow Microsoft to run rampant and have your systems updated in a kind of all-sizes-fit-all solution. An update may break a third-party app used corporate-wide, and you won't know which one it was.

Microsoft also provides an email or IM to those individuals who have subscribed to their notification system so that whenever major security updates are released, the subscribers are instantly notified.

NIST Security Documents

The following is from the NIST website:

The National Institute of Standards and Technology (NIST) maintain a Computer Security Division website (<http://csrc.nist.gov/>) that keeps up with the latest trending topics within the information security world.

Information security is an integral element of sound management. Information and computer systems are critical assets that support the mission of an organization. Protecting them can be as important as protecting other organizational resources, such as money, physical assets, or employees. However, including security considerations in the management of information and computers does not completely eliminate the possibility that these assets will be harmed.

Step 2: Incident Response

When you become aware of updates to firmware and OS software, you then need to be able to respond timely to them. Some third-party applications notify the user of a new update for things such as Java, Adobe Reader, or Adobe Flash player. Other software, mainly for your operating systems, have a system that informs the end user and the system administrator of critical updates that, if not installed in a timely manner, could jeopardize

the integrity and security of your systems. This is where that test bed comes into play as mentioned earlier.

In an ideal environment, the system administrator would have a WSUS server that would download all the updates to it; then the system administrator could push the patches to the client PCs on the test bed before pushing it to the general populace PCs of the organization. Pushing an untested patch to the CEO and crippling his/her email even for two hours makes for a very bad day. The old mantra of “test test test” is still valid today!

This isn't just about how you as a chief information security officer and your team (if you are lucky enough to have one) handle an incident, but also how well your people handle an incident. Does John from accounting forward you the email he just opened and his screen went black, or does he disconnect from the network and call you? These are the types of things you need to consider when you are putting together an incident response plan. Most likely John is not going to unplug his computer from the network, in case you were wondering.

Providing a framework for incident response is a challenge because for obvious reasons, we cannot presume to know what your organization is about, or how it falls within your internal guidelines, but the openness of the process is the feature that makes it most useful to individuals searching for effective security practices. As a result, we have come up with the following compromise that we hope proves effective:

To begin, you need to do the following:

1. Establish roles and responsibilities.
2. Define what a security incident is.
3. Establish procedures for reporting a security incident.
4. Establish guidelines for reporting an incident to an outside agency (such as DISA, NSA, or the Department of Homeland Security).
5. Establish procedures for responding to a security incident.

Establishing Roles and Responsibilities

This section will help you define who has ultimate governance of your organization's security policy—typically, the CEO, CIO, CISO, or CSO, or some other overarching office within the organization. The security policy should establish a security response team (SRT) and who those people are and how to reach them in case of an incident. The Roles and Responsibility section should also list the managers, the system administrators, and the users, and the responsibilities of each, for example:

Users: All employees and other systems users are responsible for reporting security incidents. They must immediately notify their manager or LAN administrator. If the manager or LAN administrator is not available, the user must immediately report the incident to the network service center and notify the CSO or RSO.

Defining a Security Incident

A good but fairly general definition of an incident is “...*the act of violating an explicit or implied security policy.*” Unfortunately, this definition relies on the existence of a security policy that, although generally understood, varies among organizations.

For the federal government, an incident, defined by NIST Special Publication 800-61, *is a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard computer security practices.* You can find federal incident reporting guidelines, including definitions and reporting timeframes at <http://www.us-cert.gov/federal/reportingRequirements.html>.

In general, types of activity commonly recognized as being in violation of a typical security policy include but are not limited to the following:

- Attempts (either failed or successful) to gain unauthorized access to a system or its data, including PII-related incidents (link to the following description)
- Unwanted disruption or denial of service
- The unauthorized use of a system for processing or storing data
- Changes to system hardware, firmware, or software characteristics without the owner’s knowledge, instruction, or consent

Security incidents might involve suspected threats to persons, attempted systems intrusions, unauthorized release of Privacy Act information, theft of government or personal property, or any other suspicious situation. This policy provides the procedure for reporting those incidents. It should be broken down into three or more subsections:

1. Information systems security incidents
2. Physical security incidents
3. Misuse or abuse

1. Information Systems Security Incidents This policy subsection can be divided between malicious software called malware (which consists of viruses, worms, Trojans, spyware, bad adware, botnets, and most rootkits) or systems intrusion.

Malicious code can be a virus, worm, or Trojan horse, and all are designed to do damage to data.

A *virus* is a specifically programmed set of instructions intended to destroy, alter, or cause loss of data. It can spread from one program to another, from one system to another, or from one computer to another. A typical computer virus copies itself into the operating software and executes instructions to erase, alter, or destroy data.

Worms are similar to viruses in that they make copies of themselves, but differ in that they need not attach to particular files or sectors. After a worm executes, it seeks other systems to infect, and then copies its code to them.

Trojan horses are not viruses. However, they are programs that contain destructive payloads, which pretend to be legitimate programs. They are spread when the user executes the program.

Botnets are a collection of infected computers (bots) that have been taken over by hackers and are used to perform malicious tasks or functions. A computer becomes a bot when it downloads a file (for example, an email attachment) that has bot software embedded in it. A botnet is considered a botnet if it takes action on the client via IRC channels without the hackers logging in to the client's computer. A botnet consists of many threats contained in one. The typical botnet consists of a bot server (usually an IRC server) and one or more botclients.

Note An Internet Relay Chat Channel (IRC Channel) is a form of real-time Internet text messaging (chat) or synchronous conferencing used primarily for group communication in discussion forums (channels), but also enables one-to-one communication via private message and chat and data transfer, including file sharing.

IRC was created in 1988. Client software is now available for every major operating system that supports Internet access. As of April 2011, the top 100 IRC networks served more than half a million users at a time, with hundreds of thousands of channels operating on a total of approximately 1500 servers out of roughly 3200 servers worldwide.

You might have controls in place to protect your data from alteration, destruction, and disclosure; however, there still might be attempts to gain access to your systems. Systems intrusions can take various forms. They may include denial of Internet or email services, unauthorized control or modification of web pages, vulnerability scanning, password cracking, sniffing, social engineering to gain system access, and others. All suspected systems intrusions, or attempts, must be immediately reported to management.

2. Physical Security Incidents Your organization's physical security program should be designed to protect personnel and facilities, materials, equipment, and information against threats both natural and man-made. Your corporate internal policies and instructions should contain the policies for your employees and procedures for reporting incidents. Some examples of what we are talking about can be as simple as posting the fire evacuation route and designating areas where employees are to meet to as critical as dealing with theft and threats. Most organizations consider these physical security controls guns, guards, and gates.

3. Allegations of Fraud, Abuse, or Misuse Many state laws, industry standards, and federal statutes require you to protect the integrity, privacy, and confidentiality of all personal data and ensure the integrity of your customers. We've gone over just a few in Chapter 2. You, as the security professional within your organization, should take precautions by using policies, procedures, standards, guidelines, and controls to ensure that personal data entrusted to your organization is not misused and that the programs are safe from abuse by the public and the employees who administer them.

You should list the applicable manuals or regulations you use and ensure your people are aware of their responsibilities as users, employees, and contractors. Also make known to the employees how to report suspected fraud cases to their supervisors or through the fraud hotline.

Establishing Procedures for Reporting a Security Incident

Employees, contractors, and members of the public may report any suspicious incidents involving information systems. You need to have procedures in place that define

- Whom to notify
- What to report

Typically, you would report the incident up the chain until it gets to security. (That is, you go from the end user to the immediate supervisor, and then to the IT coordinator and information assurance officer sitting in the security office.) Security will do a risk assessment and make decisions as to its response. You might have the end users disconnect their systems from the network to prevent any further contamination, or you might just have them turn off their systems immediately. Either way the risk is mitigated until you and your team can get a full assessment of what has happened.

The *what* to report is an easier list:

1. Employee name, number, and email address
2. Alternative point-of-contact (POC)
3. Location of the affected machine
4. Hostname and IP address of affected machine
5. Data or information at risk
6. Hostname and IP address of source of the attack (if known)
7. Any other information you can provide that assists in analyzing the incident

Establishing Guidelines for Reporting an Incident to an Outside Agency: What Are You Required to Report?

For the federal government, an incident, defined by NIST Special Publication 800-61, *is a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard computer security practices*. You can find federal incident reporting guidelines, including definitions and reporting timeframes at www.us-cert.gov/federal/reportingRequirements.html.

If you are a member of a federal agency and a security breach occurs in the information systems realm, you are required to report security incidents to FedCIRC. These reports are used by FedCIRC to build a governmentwide snapshot of attacks against government cyber resources and to assist in developing a governmentwide response to those incidents.

This URL takes you to the new federal online incident reporting website:
<http://www.us-cert.gov/federal/>.

Step 3: Imposing Your Will

Every company should have a policy indicating how changes take place within an organization, who can make those changes, how they are approved, and how the changes are documented and communicated to employees. Heavily regulated industries, such as finance, pharmaceuticals, and energy, have strict guidelines about what can be done, exactly what times, and under which conditions.

Historically, *change management* was a software development term that referred to a committee that made decisions about whether proposed changes to a software project should be implemented. It is composed of a selection of personnel that make up the decision-making head of the corporations, (that is, department heads) and this monster is called a change control board (CCB) The CCB is made up of project stakeholders or their representatives. The authority of the change control board may vary from project to project, but decisions reached by the CCB are often accepted as final and binding. A typical CCB consists of the development manager, the test lead, and a product manager.

Change management aims to ensure that standardized methods and procedures are used for efficient handling of all changes. The main goals of change management include the following:

- Minimal disruption of services
- Reduction in back-out activities
- Economic utilization of resources involved in the change

Change management would typically be composed of the raising and recording of changes; assessing the impact, cost, benefit, and risk of proposed changes; developing business justification and obtaining approval; managing and coordinating change implementation; monitoring and reporting on implementation; and reviewing and closing change requests.

Change management is responsible for managing change process involving the following:

- Hardware
- Communications equipment and software
- System software
- All documentation and procedures associated with the running, support, and maintenance of live systems

Any proposed change must be approved in the change management process. Although change management makes the process happen, the decision authority is the change advisory board (CAB), which is made up for the most part of people from other functions within the organization.

Now let's put that into perspective and focus on security and incident response and patch management. It is the same concept whereby any changes you are bringing forth into the organization should meet this CCB so that it is understood what is being patched, and

why, and what happens if the patch and upgrade should fail. Typically, this CCB would consist of a member of the SRT, the CISO, and department heads.

Security Response Team (SRT)

The SRT is responsible for making decisions on tactical and strategic security issues within the enterprise as a whole and should be tied to one or more business units. The group should be made up of people from all over the organization so that they can view risks and the effects of security decisions on individual departments and the organization as a whole. The CEO should head this team, and the CFO, CIO, CISO, department managers, and chief internal auditor (if applicable) should all have a seat. This team should meet quarterly, at a minimum, and have a well-defined agenda. Its responsibilities include the following:

- Defining the acceptable risk level for the organization
- Developing security objectives and strategies
- Determining priorities of security initiatives based on business needs
- Reviewing risk assessment and auditing reports
- Monitoring the business impact of security risks
- Reviewing major security breaches and incidents
- Approving any major change to the security policy and program

It should also have a clearly defined vision statement in place set up to work with and support the organizational intent of the business.

Note When comparing CSO versus a CISO, the CSO and the chief information security officer might have overlapping responsibilities. It is up to the organization to define the roles and whether one or both will be used. The CSO role usually has a farther reaching list of responsibilities compared to that of the CISO. The CISO is typically focused more on the hands-on technical aspect and has an IT background, whereas the CSO is more focused on the business risks, including physical security. In an organization, the CISO reports directly to the CSO.

Steps 4 and 5: Handling Network Software Updates (Best Practices)

How software updates are tested and applied to your working environment is a matter of balancing the need for security and the need for functionality. Many organizations allow the users (in a Windows environment—and let's face it, most office environments are Microsoft based) to use Microsoft's built-in security update tool, which in theory is a good idea but in practice might not be the smartest thing to do. Microsoft's Security

Update tool sees the users' environment as living in a vacuum. It sees you have a certain operating system, Microsoft Office products, maybe a database, maybe Visio, and so on, and then it pulls down the most up-to-date patches for that environment. It cares less about testing the patch on your system than your end user does. Apple's latest OS, Mac OS X, has a tool whereby it checks weekly for updates or patches and enables you to check which ones you want to have installed. This is better but not perfect. You might have a system running a bit of third-party-middleware that cannot be upgraded for whatever reason. Your end users download the latest patch, update their system, break it, and leave it more vulnerable than when it started.

To complicate matters even more, you might have virtual machines in your server farm. How can updating the host affect the virtual machines?

You need to have a plan, and this plan needs to consist of testing, planning a CCB, and then pushing the patches and having a rollback plan if the patch fails.

Some generic best practices apply to all updates regardless of whether they are service packs, hotfixes, or security patches. Then there are some specific items that need to be performed depending on what kind of update it is (security patch, hotfix, or service pack). First, let's define what these are aside from a general acknowledgment that they are updates to products to resolve a known issue or workaround.

Note Refer to *Best Practices for Applying Service Packs, Hotfixes, and Security Patches*, by Risk Rosato.

- **Service Pack:** A collection of updates, fixes, and enhancements to a software program delivered in the form of a single installable package. Many companies, such as Microsoft or Autodesk, typically release a service pack when the number of individual patches to a given program reaches a certain limit.
- **Hotfix:** A single, cumulative package that includes one or more files used to address a problem in a software product. In a Microsoft Windows context, hotfixes are small patches designed to address specific issues, accessibility service, and freshly discovered security exploits and other concerns of vulnerability. Other companies define it differently; the game company Blizzard Entertainment has a different definition for the term hotfix in its game *World of Warcraft*:

“...a hotfix is a change made to the game deemed critical enough that it cannot be held off until a regular content patch. Hotfixes require only a server-side change with no download and can be implemented with no downtime, or a short restart of the realms.”
- **Security Update:** A change applied to an asset to correct the weakness described by a certain vulnerability. This corrective action will prevent successful exploitation and remove or mitigate a threat's capability to exploit a specific vulnerability in an asset. If that sounds like a definition, that's because it is. It's taken right from the CISSP study material. What this means in a less convoluted manner is that security updates

or patches do just what their name implies. They are the primary source for fixing security vulnerabilities in software. Microsoft releases all its security patches once a month; other companies have different dates throughout the month.

Industry Best Practices

Security is no different from any other industry. Steps and techniques are expected as a baseline best practice. This section looks at some of them.

Use a Change Control Process

As mentioned earlier in this chapter, a good change control procedure has an identified owner, a path for customer input, an audit trail for any changes, a clear announcement and review period, testing procedures, and a well-understood back-out plan. Change control manages the process from start to finish. We should also mention that changes are typically applied only during nonwork hours. If your current procedure lacks any of these, reconsider carefully before using it for deployment of updates.

Read All Related Materials

Before applying any service pack, hotfix, or security patch, it is imperative you read all relevant documentation and have it peer reviewed. The peer review process is critical because it mitigates the risk of a single person missing critical and relevant points when evaluating the update. There is no worse feeling than pushing out an update and watching as your mail server reboots for the last time because you didn't do due diligence.

Reading all associated documentation is the first step in assessing whether

1. The update is relevant and will resolve an existing issue.
2. Its adoption won't cause other issues resulting in a compromise of the production system.
3. There are dependencies relating to the update. (That is, certain features being enabled or disabled for the update to be effective.)

Potential issues will arise from the sequencing of the update because specific instructions might state or recommend a sequence of events or updates to occur before the service pack, hotfix, or security patch is applied.

Documentation released with the updates is usually in the form of web pages, attached Word documents, and README.TXT files. These should be printed and attached to change control procedures as supporting documentation.

You can find specific write-ups on the Microsoft KB Articles; for instance, KB Article 892843 (<http://support.microsoft.com/kb/892843>) for a security update; Microsoft Outlook has a detailed description in the write-up.

Apply Updates as Needed

Apply updates only on an as-needed basis. One of the common misconceptions about any updates, be they from Microsoft, Apple, or Adobe, is that they are mandatory and urgent.

All updates, regardless of their type (whether they are service packs, hotfixes, or security patches), are to be applied on an as-needed basis. They need to be individually evaluated and treated as important optional updates.

Especially with security patches, the expectation is that it must be an urgent issue and must be quickly deployed. Without trying to detract from the urgency, security patches are very much a relative update; for example, customers using solely Windows XP (SP3) can ignore a patch for a security vulnerability in Windows 2007. However, if the issue is relevant and does plug a security hole, it should be urgently evaluated.

Only when it addresses or fixes an issue being experienced by the customer should it be considered. Of course, it still needs to be evaluated before being installed. Don't read this section and think you can just use a Windows NT 4 machine on every desktop and that mitigates the risk and therefore you do not need to worry about security updates. No! There is a reason Microsoft, Linux, HP-UX, Apple, and so on evolved throughout the years. To be the most secure you can be, you need to update to the most current version of the most current operating system of your choosing. Stagnating at a lower version doesn't help your security posture...it negates your security posture.

Testing

The prior points assist in giving you a feel (before installing) for the potential impact; however, testing allows for the "test driving" and eventual signing off of the update.

Service packs and hotfixes must be tested on a representative nonproduction environment prior to being deployed to production. This will help to gauge the impact of such changes.

Uninstall

Where possible, service packs, hotfixes, and security patches must be installed such that they can be uninstalled, if required.

Historically, service packs have enabled uninstalling, so verify there is enough free hard disk space to create the uninstall folder.

Consistency

Service packs, hotfixes, and security patch levels must be consistent on all servers and workstations throughout your computing environment. Inconsistent update levels across your company can lead to synchronization and replication-related problems. Anyone who has maintained a Microsoft domain knows it is extremely difficult to trap errors caused by domain controllers (DC) being out of sync, so it's critical that you maintain consistency.

Backup and Scheduled Downtime

Server outages should be scheduled and a complete set of backup tapes and emergency repair disks should be available, in case a restoration is required.

Make sure that you have a working backup of your system. The only supported method of restoring your server to a previous working installation is from a backup. For more information on backup and recovery procedures, refer to your operating system forums for best practices for backing up.

Have a Back-Out Plan

A back-out plan enables the system and enterprise to return to their original state prior to the failed implementation. These procedures must be clear, and contingency management must test them because in the worst case a faulty implementation can make it necessary to activate contingency options.

Enterprises might need to exercise their back-out plan if the update does not have an uninstall process or the uninstall process fails. The back-out plan can be as simple as restoring from tape, or may involve many lengthy manual procedures.

Forewarn Helpdesk and Key User Groups

You need to notify helpdesk staff and support agencies of the pending changes so that they are ready for arising issues or outages.

To minimize the user impact, it is also a good idea to prepare key user groups of proposed updates; this can assist in managing user expectations.

Don't Get More Than Two Service Packs Behind

Schedule periodic service pack upgrades as part of your operations maintenance, and try never to be more than two service packs behind. As mentioned before, service packs are composed of large security updates and hotfixes bundled together. If you do not keep up with the service packs as they are issued, you leave your computing environment vulnerable.

Target Noncritical Servers/Users First

If all tests in the lab environment are successful, start deploying on noncritical servers first, if possible, and then move to the primary servers after the service pack has been in production for 10–14 days. There is nothing worse than upgrading the CEO's system first only to find out the upgrade crashed his system, and he not only lost the time it takes for you to have him back up and running, but also the potential loss of critical business plans and his recipe for white chicken chili.

Service Pack Best Practices

Great Microsoft TechNet articles reference service pack best practices. All you need to know can be found in the following documents:

- *Steps to Take Before Installing Windows XP - SP3* (<http://support.microsoft.com/kb/950717>).
- *How to Install the Latest Service Pack or Update Rollup for Exchange 2010* (<http://www.microsoft.com/downloads/en/details.aspx?FamilyID=50b32685-4356-49cc-8b37-d9c9d4ea3f5b&displaylang=en>).
- *Learn How to Install Windows 7 SP1* (<http://windows.microsoft.com/en-US/windows7/learn-how-to-install-windows-7-service-pack-1-sp1>).

Apple products don't have service packs per se; however, they do have annual updates that typically include new functionality and features to the system.

Hotfix Best Practices

The following sections outline some best practices for hotfixes.

Service Pack Level Consistency

Now take a moment to reflect on one of the major bullets listed under the general guidelines: consistency. There is a reason you should have the same service pack deployed to all machines running the same operating system. Don't deploy a hotfix until you have all current service packs installed. A hotfix is related to a service pack and should be deployed with this in mind. If a hotfix is for post-Windows 2000 SP2, for example, you need to ensure that the server has SP2 installed. And if you are constant throughout your computing environment, you won't need to do extra work in bringing a system up to level to patch it.

Latest Service Pack Versus Multiple Hotfixes

This last one is common sense. I don't know about you, but my time at the office has me spread thin. If I don't need to spend three hours uploading 200 hotfixes to a system, I won't. My time can be better used elsewhere. If multiple hotfixes need to be applied to a system, and these are in the latest released service pack, apply the latest service pack instead of applying several hotfixes unless issues relating to the latest service pack might cause the server to break.

Security Update Best Practices

The following sections outline some best practices for security updates.

Apply Admin Patches to Install Build Areas

It is crucial that not only systems deployed to the desk are retrospectively updated with security patches, but also the client build areas are updated for any new clients. For example, I do not have a “client build” area in my organization—unless you count my office—however, I have a procedure established that I have a fast patch disk and a fast secure disk. The fast patch has all the latest updates on it post SP3 (for my Windows XP builds). So after I install the base operating system and Service Pack 3 (SP3), I run the latest fast patch disk. When this is done, I run the fast secure disk on the system. This then enables me to safely put the newly built system on the network or domain to go out to Microsoft’s update services and get the updates listed from the end of the fast patch disk to present.

Apply Only on Exact Match

Apply fixes only if you encounter exactly the issue the fix solves or if the circumstances relate to your environment.

Subscribe to Email Notification

Subscribe to the notification alias to receive proactive emails on the latest security patches, or use newsgroups or other forums to help you stay apprised of the ever-changing risks. Security is not passive. You must be aggressive and continually read, upgrade, and plan for the worst.

Following are links to various online resources to assist you in maintaining a Microsoft, Novell, and Cisco environment:

- **Microsoft:**
<http://www.update.microsoft.com/microsoftupdate/v6/default.aspx?ln=en-us>
- **Novell:** <http://support.novell.com/patches.html>
- **Cisco:** <http://tools.cisco.com/security/center/home.x>

Summary

Twenty pages to tell you what you already know: There is more involved in being the security officer than just clicking Install when new updates are available. You need to not only be savvy about how to update a system, but also when, why, where you need to go for answers, and how you balance functionality with allowing your employees to be productive. There is a fine line.

The main challenge in managing security updates is determining which of the many available updates are appropriate to the needs and vulnerabilities of your enterprise systems and business requirements.

Some updates are critical and require immediate action to protect your environment. For example, the updates that address risks from newly discovered exploitations, viruses, and worms are considered critical updates.

Some updates can be useful, can increase performance or stability, or can make the end-user experience better, but they might not be considered critical to the safety of your enterprise. Other updates might not be necessary to your enterprise and can be ignored.

Some updates could create problems (for example, break other line-of-business applications) for your enterprise if you use them.

To keep your enterprise secure, you must establish processes for the following:

- Receiving information about the latest software updates and vulnerabilities
- Auditing your enterprise for applicable software updates
- Assessing and authorizing available software updates
- Deploying authorized software updates within your enterprise in a timely, accurate, and efficient manner
- Tracking update deployment across your enterprise

To learn how to determine which updates are critical, useful, irrelevant, or harmful to your enterprise, and to create a software update management process for your enterprise, you need to be familiar with the current state of the resources in your enterprise. This includes knowing the following:

- The computers in your enterprise
- Operating systems and versions functioning on the computers
- Software updates in use on your computers (service pack versions, software updates, and other modifications)
- The function each computer performs in your enterprise
- The applications and programs running on each computer
- Ownership and contact information
- The assets present in your environment and their relative value, to determine which areas need the most protection
- Known vulnerabilities and the processes your enterprise has for identifying new vulnerabilities or changes in vulnerability level
- Countermeasures that have been deployed to secure your environment

This information should be updated regularly and should be readily available to those involved in your update management process.

Chapter Review and Questions

This is typically where you'd find the chapter review and questions to spark some cognitive recognition behind your tired eyes. I am not including chapter review questions here. Instead, I am including a list of links that might be beneficial to you, the security officer. Listed are some checklists, best practice links, security websites, and so on that I hope you find useful.

Cisco Security Intelligence Operations: <http://tools.cisco.com/security/center/home.x>

Novell Patches and Security: <http://support.novell.com/patches.html>

How to Use Microsoft Update:

www.update.microsoft.com/microsoftupdate/v6/default.aspx?ln=en-us

Microsoft TechNet Security Archive: <http://technet.microsoft.com/en-us/library/dd365874.aspx>

Microsoft Security Best Practices: <http://technet.microsoft.com/en-us/library/dd366071.aspx>

Microsoft Security Checklists: <http://technet.microsoft.com/en-us/library/dd366061.aspx>

Symantec Virus Updates and Security Threats:

www.symantec.com/business/security_response/definitions.jsp

VMware Security Advisories: www.vmware.com/security/advisories/

Federal Agencies Security Practices: <http://csrc.nist.gov/groups/SMA/fasp/index.html>

NIST CSRC: <http://csrc.nist.gov/index.html>

CVE International: <http://cve.mitre.org/index.html>

Network Security Standards and Guidelines

By the end of this chapter, you should know and be able to explain the following:

- Identify resources for use within your organization
- Harden a Cisco IOS or operating system
- Harden an Apple OS X operating system
- Harden a Microsoft Server and desktop environment

Let's have an '80s movie flashback with Chevy Chase to set the stage for this chapter. One of his best movies was and still is *Fletch*. In this scene, you see Fletch looking for clues in an aircraft hangar dressed as a mechanic....

Willy: What do you need ball bearings for?

Fletch: Awww, come on guys, it's so simple. Maybe you need a refresher course. [*leans arm on hot aircraft engine*]

Fletch: Hey! It's all ball bearings nowadays. Now you prepare that Fetzer valve with some 3-in-1 oil and some gauze pads. And I'm gonna need 'bout ten quarts of antifreeze, preferably Prestone. No, no make that Quaker State.

Although ball bearings might have worked given the right circumstance, they weren't the right tools for the job. That is what this chapter is about: using the right security tool for the right job. In many cases, it's also knowing where that tool is and how to correctly apply it; fortunately, much of the work is already done for you so that you can protect that Fetzer valve from any ball bearings Fletch might want to test it with!

This chapter discusses and provides an overview of common design guidelines and provides some example of how they should be considered and adapted in a production environment.

You learn some of the best practices and guidelines from major industry players such as Cisco, Microsoft, and Apple, along with delving into the some best practices set forth from the National Security Agency (NSA) and the Computer Security Division Computer

Resource Center of NIST (National Institute of Standards and Technology). Then this chapter talks about how to apply these best practices within your organization.

We'll start with a few Cisco resources/tools: Cisco SAFE and Cisco Validated Designs, and follow up with best practices from Cisco on hardening the IOS, and configuring a firewall/ASA and an Intrusion Prevention System (IPS).

Cisco SAFE 2.0

Cisco SAFE is a tool that I wish I had access to when I was first designing systems back in the late 1990s. It is a consistent framework for all phases of network design for network security: design, planning, and implementation.

Overview

Cisco SAFE is a resource offered by Cisco to walk you step by step through designing, planning, and implementing a consistent security policy across all aspects of your environment, taking into consideration both Cisco security devices (that is, firewall/ASAs and ASA) and Cisco network devices (that is, routers and switches). SAFE uses a security control framework that employs various technical design components and implementation guides, all designed to increase the visibility of the components in your environment, help you identify the shortfalls, and assist you in avoiding “stove-pipe” security solutions that focus solely on one aspect of your environment, such as the data center or the edge devices.

Purpose

Using SAFE enables consistent security policy deployment across your environment, combining network and security devices into a seamless security platform for the campus/data center, Internet edge, and remote (branch office, virtual office, and clients) aspects of your environment. This also assists you in the design and planning phases of your security deployment to create a truly unified security strategy.

SAFE has several benefits: collaboration of devices, modular, and consistent implementation guides. The security control framework that SAFE uses enables you to witness how a security implementation affects Cisco network devices and Cisco security devices across the entire implementation, from data centers outward to branch offices through multiple Cisco devices, ensuring consistent Layers 2 and 3 design. Furthermore, it has a modular approach. There are different design guides for each segment on the network: campus, Internet edge, branch office, and so on. These first two benefits give birth to a document that is invaluable: a customized, yet consistent design implementation guide for your environment that enables you to view your security shortfalls and avoid the stove-pipe effect of securing one aspect of your environment while neglecting the others.

We encourage you to look into this golden egg of a tool at www.cisco.com/go/safe.

Cisco Validated Design Program

The Cisco Validated Design (CVD) program is another resource unique to Cisco that can assist you and your organization with faster, more reliable, and more predictable deployment of systems and solutions, and you'll have confidence that it will be done securely and accurately the first time. These designs incorporate a wide range of technologies and product solutions. CVDs are organized by solution area; in other words, you can research based on your needs: architecture (that is, borderless networks, data centers, and collaboration), technology (that is, security, mobility, and unified communications) or industry (that is, education, government, and healthcare).

This chapter provides a basic overview of just a few of the design guides, along with links to the website where you can download these resources, or you can do more research based on your needs.

Branch/WAN Design Zone Guides

Overview: These guides assist the network designer, or network engineer, in deploying high-value network services on a secure branch network connected to a central site, despite the variety of WAN technologies.

The website

(www.cisco.com/en/US/netsol/ns816/networking_solutions_program_home.html) breaks down the design guides by the branch size (small, medium, or large), multicast implementation, and overall design guides for both branch and WAN.

Campus Design Zone Guides

Overview: This section consists of four subsections: Overall Campus Design, High Availability, Network Virtualization, and Unified Access.

The Campus Design guide addresses enterprise campus architectures using the latest advanced services technologies from Cisco and is based on best-practice design principles. It introduces the key architectural components and services necessary to deploy a highly available, secure, and service-rich campus network.

The High Availability design guide gives you an understanding of what enterprise campuses require to maintain a highly available, secure, intelligent network infrastructure to support business solutions (that is voice, video, and wireless). It can help you plan for, and understand, how the system recovers from component outages (planned and failures) and what the expected behavior is during such an outage. Knowing this is a critical step in designing, upgrading, and operating a highly available, secure campus network, and it is provided to you by Cisco.

The Network Virtualization guides provide multiple solutions to business problems and drivers that range from simple to complex. This guide provides solutions for an end-to-end network virtualization solution separated into the following three functional areas: access control, path isolation, and services edge. If you need to provide something as simple as

Internet access to visitors or something more complex, such as providing Internet access as a line-of-business and a revenue stream to various clients, you need to peruse these guides.

And finally the Unified Access guide combines switching, wireless, location, identity, policy, and management in a system design to simplify IT network-access management.

Unified Access systems also provide for a consistent experience when users access the network. Check out the website at

www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html.

Data Center Design Zone Guides

The Data Center Design Center has a lot to offer, including the following:

- **Application Networking:** Cisco provides strategies for implementing various application platforms, such as IBM, Oracle, SAP, Citrix, Seibel, Microsoft, and more.
- **The Business Continuity Design Guide:** Provides instruction and guidance for implementing high availability clusters to distributed data centers.
- **The Cisco Unified Computing Design Guide:** Shares how to deliver a scalable, cost-effective architecture for your current and future needs for virtualized Microsoft applications, using computing and storage resources with ease using Oracle applications and Cisco UCS, or improving computing and storage resource allocation within the data center.
- **Cloud Computing:** This is a hot ticket item. This design guide goes through the steps required to create a flexible data center, enabling you to efficiently share your resources and become better prepared for rapid change management.
- **Data Center Networking:** Cisco Data Center Business Advantage provides architecture to unify the virtualization capabilities of individual devices to create a fully virtualized data center. This includes such items as Wide Area Application Services (WAAS) designs or creating a virtual data center infrastructure and integrating those WAAS with it.
- **Security (Data Center):** Comply with regulations and protect your data center from attack. But it goes beyond that. Sure you can get solutions for securing your data center, but you can also learn about security and virtualization in the data center. And with the big push right now to open everything up to the cloud, you had better be sure that you can protect your back-end information. Even if you think you don't need to read this, you should. Take some time in your afternoon and become better educated on security and virtualization.
- **Server Networking:** This section tells you how to efficiently and securely deploy your various server networks. This includes data center blade server integration and server farm security.
- **Service Provider:** Controlling operational and capital costs while maximizing return on server, storage, and network infrastructure. The Service Provider CVD provides

validated architectural guidance for building a baseline Service Delivery Center network infrastructure.

- **Storage Networking:** The Storage Networking CVD shows how to efficiently and cost-effectively implement a storage area network (SAN) solution using the Cisco MDS 9500 series multilayer director or intermingling the IBM FICON with the Cisco MDS 9000. This is an invaluable tool for resource allocation within your data center.
- **Virtualization:** Bring network, computer/storage, and virtualization platforms closer together to provide unparalleled flexibility, visibility, and policy. This includes three other subtopics: network, server, and desktop virtualization.

You can find more information about the Data Center Design Center at www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html.

Security Design Zone Guides

The design zone for security has several sections; Cisco SAFE is one section we explained. But where SAFE differs from the rest of the sections is that SAFE is a tool you can use. The other sections are design guides for network foundation protection, enterprise campus security, and threat control:

- **Cisco SAFE:** For a detailed description, refer to the previous section, “Cisco SAFE 2.0.”
- **Enterprise Campus Security:** Leverages network virtualization for security. This CVD represents a chapter of the overarching SAFE Design Guide.
- **Network Foundation Protection:** Details security architects for the enterprise network. This too represents a separate chapter in the SAFE Design Guide. This portion goes over best practices for securing the network infrastructure by doing things such as setting security policies for infrastructure device access, routing infrastructure, network telemetry, and network policy enforcement.
- **Secure Campus:** Integration of wireless and security into the campus networks. The overarching purpose of this design guide is to describe the integration and collaboration of network security technology and the Cisco Unified Wireless Network.
- **Secure Technology Partners:** Solutions to security issues for data security systems to implementing a security information and event management system from Cisco and its partners.
- **Security in Branch:** Security solutions for branch locations. Branch offices can offer their own challenges, such as high availability, infrastructure protection, secure and persistent connectivity to the home office, and threat defense. This CVD offers guidelines and best practices for addressing these issues.
- **Security in WAN:** This CVD speaks specifically on the implementation of an IPsec virtual private network (VPN) WAN design. It defines the components required to build a site-to-site VPN system in relation to your WAN connectivity.

- **Threat Control:** This is another chapter in the SAFE Design Guide. It specifically covers the threat detection and mitigation capabilities available on Cisco Firewall/ASAs, Cisco IPS, Cisco Security Agents (CSA), Cisco Network Admission Control (NAC), and web/email security appliances.

You can find more information about the Security Design Zone Guides at www.cisco.com/en/US/netsol/ns744/networking_solutions_program_home.html.

Cisco Best Practice Overview and Guidelines

Trying to apply best practices and guidelines is difficult enough, but trying to do it when you need to worry about how putting in a certain access control list (ACL) affects the components of the configuration of your network or security components is something entirely different. You not only need to be aware of how they operate but also how they interoperate. Cisco provides a handy location where all the whitepapers, discussion groups, and so on are easily located. You can drill down to a specific IOS, and then a specific type configuration within that IOS, say *Security and VPN*, and from there you can narrow that beam a little more to find what it is specifically you're looking for, say, *Access Control Lists (ACL)*, for instance.

Following are a few models and some links and tools for configuring your Cisco IOS device, firewall/ASA, or an intrusion prevention system (IPS).

Basic Cisco IOS Best Practices

Two access modes are available for Cisco IOS-based devices: basic and privileged. Both modes should be password protected. When you log in to the device and successfully enter the initial login sequence, the system enters basic mode. You can then enter privileged mode by typing the **enable** command followed by the password.

Secure Your Passwords

The passwords for basic mode and privileged mode should be different. Another useful tool in protecting your passwords is to use the **enable secret** command when setting your **enable** password. It uses an improved encryption algorithm over the **enable password** command. The **enable secret** command also provides more security for your configuration files should they be stored remotely on a TFTP server. Best practice is to always use the **enable secret** version of the **enable** password because the older version is easily cracked with free online tools. And finally, passwords should never be seen in cleartext when you view any configuration file. The capability of IOS to automatically hide the real passwords when a configuration is displayed is accomplished using the global command **service password-encryption**:

```
Switch(config)# service ?
password-encryption  Encrypt system passwords
<<<output omitted for brevity>>>
Switch(config)#
```

The password authentication can take on one of three modes: password, secret, or TACACS+.

```
Switch(config)# enable ?
  last-resort  Define enable action if no TACACS servers respond
  password     Assign the privileged level password
  secret       Assign the privileged level secret
  use-tacacs   Use TACACS to check enable passwords
```

The **password** and **secret** commands enable you to set an encrypted password that users must use to enter into privileged mode. Here's the trick. The difference between **enable password** versus **enable secret** is in the encryption algorithm used to encrypt the password. Using the **enable password** command uses a reversible algorithm, which is necessary to support certain authentication protocols, notably CHAP, which sends the passwords in cleartext. The **enable secret** command encrypts the passwords using the MD5 algorithm. MD5 is not reversible and is more secure.

Following is an example of the configuration options for the **enable secret** command:

```
Switch(config)# enable secret ?
  0      Specifies an UNENCRYPTED password will follow
  5      Specifies an ENCRYPTED secret will follow
  LINE   The UNENCRYPTED (cleartext) 'enable' secret
  level  Set exec level password
```

Limit Administrative Access

Many times you might want to assign particular members of your staff a subset of the privileged **enable** commands. Cisco has provided for this eventuality by enabling 16 various privilege levels (0–15). Level 1 is basic mode, and level 15 is the current privileged mode.

Limit Line Access Controls

At a minimum, users should be authenticated before gaining device access. Use consistent authentication mechanisms if possible to simplify keeping track of passwords.

Various line access options are available, such as **console**, **aux**, and **vtty**. The console port is useful in initial configuration and in cases where the network is down. Use the **aux**, or auxiliary, ports for modem support and asynchronous connections. The **vtty** are usually reserved for remote console access where administrators can Telnet into a device to access and do their daily jobs as if they were physically connected via the console port.

The line access should be controlled by authentication with a username and an encrypted password. For authentication, you can use a local login account or a TACACS+ server. To set the login authentication type for the device, see the following example.

Note TACACS+ is a Cisco proprietary protocol that provides access control for routers, network access servers, and other networked devices. RADIUS is an alternative solution used in many organizations. However, unlike RADIUS, TACACS+ provides separate authentication, authorization, and accounting services for each networked device, whereas RADIUS combines authentication and authorization in a user profile. Another major difference between the two authentication services is TACACS+ uses TCP, whereas RADIUS uses UDP.

```
Switch(config-line)# login ?
  local   Local password checking
  tacacs  Use tacacs server for password checking
  <cr>
```

The *local* keyword denotes using the local database for authentication. We do not recommend using a local database on each router or switch in your environment. This is an administrative challenge if you have more than two devices, especially if you are using a local database on your router for VPN authentication for your external clients. It is best to use one centrally managed database server such as a TACACS+.

For the aux and vty ports, it is our recommendation that you limit them to a local login or disable them entirely. Following is a sample line access configuration in which the console access is secured using a simple password but access to the vty and aux ports is secured using a local database of users:

```
!
line con 0
  password 7 047E050200335C465817
line aux 0
  login local
line vty 0 4
  login local
```

Limit Access to Inbound and Outbound Telnet (aka vty Port)

You can limit access to inbound and outbound Telnet connections on vty ports by putting in access lists that permit or deny access from, or to, only specific networks or host devices. A more detailed description of access lists can be found in Chapter 7, “Firewalls.” Following is an example of an access list applied to a vty port:

```
access-list 5 permit 130.109.6.0 0.0.0.255
!
line con 0
  password fastrouter
line vty 0 15
  access-class 5 in
  login local
!
end
```

Another way to handle interactive access is to completely prevent it by using the configuration command **no exec** on any asynchronous line. This command enables only an outgoing connection for a line. When an outside user attempts to connect via Telnet to a line with the **no exec** command configured, the user gets no response. You can also establish which protocols can be used to connect to a specific line by using the **transport input** line configuration command; for example, SSH and Telnet or just one:

```
!
line con 0
  password fastrouter
line vty 0 15
  login local
  no exec
transport input ssh
!
end
```

Establish Session Timeouts

The default timeout period for unattended console or vty connections is 10 minutes. This can be modified—and should be modified—with the **exec-timeout** command (as shown in the following, where I’ve established a timeout of 2 minutes and 30 seconds for the console port and vty lines).

```
!
line con 0
  exec-timeout 2 30
line vty 0 15
  access-class 5 in
  exec-timeout 2 30
  login local
  transport input ssh
!
end
```

Make Room Redundancy

If you have critical network segments, the Cisco devices supporting these segments should be configured with Hot Standby Router Protocol (HSRP). What HSRP does is provide high network availability (HA) by routing IP traffic from hosts without relying on the availability of a single router.

Basically, you take two routers and configure them such that there is a virtual MAC address and an IP address that is shared among the routers running HSRP. One is selected as the primary, or active, router. The active router receives and routes packets destined for

the group's MAC address. HSRP detects when the active router fails, at which point the secondary, or backup, router assumes control of the group's MAC address and IP address. The backup becomes active, and if another router is in the group, it becomes the standby.

HSRP is configured in the interface configuration mode. Go to www.cisco.com/en/US/docs/internetworking/case/studies/cs009.html to see a good whitepaper on configuring HSRP for fault-tolerant IP routing.

Protect Yourself from Common Attacks

Whenever possible, put into place filters to ensure only valid network addresses are permitted past the routers. All corporate infrastructure routers should have filters in place to disallow any obviously malicious traffic. For example, any edge router should deny traffic whose source address is one of the RFC reserved addresses listed in Table 4-1.

Table 4-1 *RFC Reserved Address Space*

Network IP Address	Subnet Mask
10.0.0.0–10.255.255.255	255.0.0.0
172.16.0.0–172.31.255.255	255.240.0.0
192.168.0.0–192.168.255.255	255.255.0.0

You can do some forensics on your routers to find where the attacks originate by using the logging feature of the `access-list` command. It is enabled by adding `log-input` to an access list entry, for example.

```
Router(config)# access-list 100 permit ip any any log-input
```

The result is an output that looks similar to the following example:

```
%SEC-6-IPACCESSLOGP: list 100 permitted udp 130.109.35.3(53)
(Gigabit0/1) -> 130.109.69.45 (5775), 1 packet
%SEC-6-IPACCESSLOGP: list 100 permitted icmp 130.109.35.3(53)
(Gigabit0/1) -> 130.109.69.45 (0/0), 1 packet
```

Another common attack is a TCP/SYN attack or a flooding attack. A TCP/SYN attack occurs when a large number of TCP/SYN packets are sent to a server. These packets have a source destination that is spoofed and not in use. When the server receives these packets, it responds to them using a SYN-ACK. However, because the source IP does not exist, the TCP handshake is never completed. This forces the server into a wait mode because it sits and waits for the ACK from the nonexistent sending device. To sit in this wait mode, the server allocates resources to record the information it received in the SYN packet and sent out in the SYN-ACK. This waiting results in half-open TCP requests. The TCP connection queue determines how many half-open requests your server can manage. After the queue is filled, the server can no longer accept any other packets, and legitimate users are denied the services offered by the server.

So, how can you protect your servers from this maliciousness? Use the **ip tcp intercept** command. This command keeps track of the following:

- Number of session requests in the last minute
- Number of incomplete sessions
- Time until final acknowledgment

A great guide for configuring TCP intercept and preventing denial-of-service attacks can be found at www.cisco.com/en/US/docs/ios/sec_data_plane/configuration/guide/sec_cfg_tcp_intercpt.html.

Firewall/ASAs

Firewall/ASAs are a key component when you talk about securing your corporate network infrastructure. This section on firewall/ASAs reviews basic firewall/ASA best practices for the core components of your network security architecture: identity, integrity, confidentiality, availability, and audit.

The Cisco ASA enables you to establish stateful firewall/ASA protection and secure VPN access with a single device. PIX Firewall/ASA provides a scalable security solution with failover support available for selected models to provide maximum reliability. ASA uses a specialized operating system that is more secure and easier to maintain than software ASAs found on the higher-end routers that use a general-purpose IOS, which are subject to frequent threats and attacks.

This section covers some basic industry best practices, including examples. At the conclusion of the section, we've provided links to configuration guides from the Cisco website.

Encrypt Your Privileged User Account

When you enter the password into the configuration, it is encrypted using an MD5 algorithm. The following example shows how to enter the Cisco ASA enable password encrypted:

```
enable password getsmarter
show enable password
enable password fecGHTsjguFGH encrypted
```

If you use the **encrypted** command during configuration, you are telling the ASA that the word preceding **encrypted** is already encrypted, as demonstrated in the following example:

```
enable password getsmarter encrypted
show enable password
enable password getsmarter encrypted
```

Do you see the difference? And understand why the former is the recommendation of these authors?

Second, ensure you change your default passwords before any network infrastructure device is put into place.

Consider one final thing concerning basic authentication on your Cisco ASA. Like many Cisco devices, you can access the Cisco ASA via the web-based configuration tool: the Cisco ASDM. Secure who can access your firewall/ASA via HTTP by using the following commands on your ASA:

```
http ip_address [netmask]
passwd password
```

This command enables you to specify a host address users can use to connect to the Cisco ASA web browser after you enable the Cisco ASDM.

Limit Access Control

Similar to other Cisco devices, there are various ways to interact with the Cisco ASA, such as a console or Telnet. You want to ensure that if you are allowing these connections, you do it correctly. By using AAA authentication, you can ensure you are using either TACACS+ or a RADIUS server.

The command follows:

```
aaa authentication [any/telnet] console tacacs+|radius
```

By using the **any** keyword, you are telling the ASA that any connection must be authenticated (console or telnet). If you use the **telnet** keyword, only Telnet connections to the ASA will be authenticated.

Note Access to the console is available only from the inside (facing toward your organization's internal network) interface.

The **telnet** command enables you to decide who can access the ASA via the Telnet protocol. You can have up to five simultaneous connections to the ASA via Telnet. To establish a password for Telnet access, you must configure the **passwd** command as demonstrated earlier.

For more information on using AAA on your ASA, look no further than www.cisco.com/en/US/docs/security/pix/pix63/configuration/guide/mngacl.html#wp1090040.

Make Room for Redundant Systems

Typically the firewall/ASA is a major component to your organization's security and well-being. You do not need this device to be a single point of failure. To ensure this doesn't happen, Cisco suggests that you have a redundant firewall/ASA and use the **failover** command to ensure fast, dynamic recovery if there is an outage of the primary firewall/ASA.

The following link provides you with a good failover configuration example for configuring failover on your ASA device. There is even an option for interface high availability: www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_configuration_example09186a00807dac5f.shtml.

Note The failover functionality is supported only between identical ASA models running the same software version and having identical hardware.

General Best Practices

The following list is a set of best practices, in no particular order, that you should consider to ensure that your ASA is configured for optimal performance and effectiveness:

- Deny all traffic by default, and enable only needed services.
- Disable or uninstall any unnecessary inspections and features on the ASA that are not specifically required.
- Limit the number of applications that run on the firewall/ASA to let the firewall/ASA do what it's best at doing. Consider running antivirus, content filtering, VPN, DHCP, and authentication software on other dedicated systems behind the firewall/ASA unless you get a dedicated expansion module to offload those tasks to.

These listed best practices are not the end-all, be-all in configuring your firewall/ASA or ASA device.

Configuration Guides

You can find a good resource for correctly configuring your ASA at www.cisco.com/cisco/web/solutions/small_business/products/security/ASA_5500_series/index.html.

Intrusion Prevention System (IPS) for IOS

First-time users should read through the “Getting Started Guide,” which you can find at www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6634/prod_white_paper0900aecd805c4ea8.html.

This is a must-read because there can be many unintended consequences if an IPS is misconfigured. Do not try to configure the IPS subsystem on your IOS router or ASA without first reading the overview on how it all fits together.

After you finish reading through the getting started guide, make sure you heed the following Cisco best practices.

1. Always remember to retire all signatures first:

```
router(config)# ip ips signature-category
router(config-ips-category)# category all
router(config-ips-category-action)# retired true
router(config-ips-category-action)# end
```

Do you want to accept these changes? [confirm]

2. Never unretire the “all” signature category.
3. For routers with 128 MB memory, start with the IOS IPS basic category.
4. For routers with 256 MB or more memory, start with the IOS IPS advanced category.
5. Use CCP/CSM to customize the signature set by unretiring/retiring a few signatures at a time according to your network needs.
6. Pay attention to the free memory every time after you unretire/retire signatures. When router free memory drops below 10 percent of the total installed memory, stop unretiring signatures. Adding more memory does not necessarily significantly increase the number of signatures that can be loaded.
7. You must unretire and enable a signature to have it loaded and take configured actions when triggered. Enabling it does not load a signature if using IOS IPS in a network with a lot of out-of-order packets.

Note To gain the IPS feature set, you must use 12.4(9)T2 or 12.4(11)T or later T-Train releases. You cannot use Mainline image. If you use a firewall/ASA, you must use an IOS that supports the Classic IOS Firewall/ASA configuration; using a Zone-Based Firewall/ASA will not work with out-of-order packets.

This is just a brief overview of best practices for implementing an IOS-based IPS.

NSA Security Configuration Guides

The National Security Agency (NSA) has graciously provided a wealth of resources to help in securing everything from third-party applications such as Adobe or Oracle to the Department of Defense (DoD) Bluetooth peripheral device security requirements. Because NSA is nonbiased in its goal to produce a secure platform for its customers, you can find a resource for just about any device you have installed and operating within your walls.

You can find a few of the Security Configuration Guides on the NSA site. NSA partnered with Microsoft, Defense Information Systems Agency (DISA), National Institute of Standards and Technology (NIST), U.S. Air Force, U.S. Navy, U.S. Marine Corps, U.S. Army, Department of Homeland Security, and the Office of Management and Budget for actual security setting decisions.

Cisco Systems

For Cisco devices you can find resources for securing and managing your routers, Layer 2 and Layer 3 switches, and Voice over IP (VoIP) Call Managers.

The Router Security Configuration Guide, which you can find at www.nsa.gov/ia/guidance/security_configuration_guides/cisco_router_guides.shtml, provides technical guidance for network administrators and security officers. It contains principles and guidance specific to ensuring you have a secure configuration for your IP routers. You can use the presented information to control access, help resist attacks, shield other network components, and help protect the integrity and confidentiality of network traffic.

Switches Configuration Guide

The NSA's Cisco IOS Switch Security Configuration Guide, which you can find at www.nsa.gov/ia/guidance/security_configuration_guides/switches.shtml, provides technical guidance for secure configuration of switches, with detailed instructions. You can use the information presented to control access, help resist attacks, shield other network components, and help protect the integrity and confidentiality of network traffic. This applies to both Layer 2 and Layer 3 switches.

VoIP/IP Telephony Security Configuration Guides

With the proliferation of Voice over Internet Protocol (VoIP) and IP Telephony (IPT) being installed and maintained throughout the U.S. government and private industry, it is more important than ever to make sure your Call Managers and Unified Communication environment is secure, stable, and functional. The Information Assurance Directorate, Systems and Network Analysis Center (SNAC) of the NSA has provided general guidance to make the implementation of Cisco Unified Communications Manager Express (CUCME) 7.0 and Cisco Call Manager. Then it follows up by giving advice on general security guidance of IP telephony systems and recommended IP telephony architectures for your organization in the following document:

www.nsa.gov/ia/guidance/security_configuration_guides/voip_and_ip_telephony.shtml.

Microsoft Windows

You've all heard the complaints. From television, neighbors, co-workers, forums, chat rooms, strangers on planes, conferences, the co-author who is an Apple fan boy...the list goes on and on. Microsoft is unsecure and a pain to manage. Well...not necessarily. The earlier versions of Microsoft Windows were unsecure and a pain to manage. However, the programmers have learned from their mistakes, listened to the lessons learned, and taken the best parts of what other companies have done to their operating systems and incorporated them. Today, Microsoft can give you a platform that is as stable and more inherently secure right out of the box than anything that rivals it on the market. That being said, the

NSA has taken the liberty of providing the IT-savvy person in charge of your IT infrastructure another tool. You can find some of the papers and security configuration guides the NSA offers at www.nsa.gov/ia/guidance/security_configuration_guides/operating_systems.shtml#microsoft.

Microsoft Windows Applications

Microsoft includes several pieces of software in its default installation of its operating system. Things such as Wordpad, notepad, calculator, games (Solitaire, Minesweeper, and so on). Other than taking up valuable space on your hard drive, these programs are relatively unintrusive. However, there may be a reason to disallow certain pieces of software, and that is where a software restriction policy (SRP) comes in handy. It can be configured as a local computer policy or as a domain policy using Group Policy with Windows Server 2003 domains and later. The SRP enables administrators to control which applications are allowed to run on Microsoft Windows. By using this guide, administrators can configure SRP to prevent all applications in their domain from running except applications they explicitly allow. Using SRP as an application white-listing technique significantly increases the security posture of the domain by preventing some malicious programs from executing.

Microsoft Windows 7/Vista/Server 2008

The NSA has taken a different approach to the latest Microsoft operating systems. It considers the Special Security – Limited Functionality (SSLF) settings in Microsoft's Windows 7 Security Guide to track closely with the security level represented in its own guidelines. It is the NSA's belief that the guide it produces establishes the latest best practices for securing the product and recommends that traditional customers use the Microsoft Security Compliance Manager when securing Windows 7.

The NSA's website provides several papers on the subject of Windows 7: Security Highlights, Center for Internet Security (CIS) Windows 7 Benchmark, and the Microsoft Security Compliance Manager.

Check out Microsoft Security Compliance Manager at <http://technet.microsoft.com/en-us/library/cc677002.aspx>.

Note You can find the Windows 7 Version 1.1.0 benchmark at <https://benchmarks.cisecurity.org/en-us/?route=downloads.form.windows7.110>.

Microsoft Windows XP/Server 2003

Microsoft also offers a security configuration guide on Server 2003 and Windows XP. This chapter does not cover these because these will be phased out of your environment soon; however, you should be aware that they are there, and if needed you can download them and implement ASAP. It would be a good start to make sure your current security posture on those two operating systems is worthy.

Apple

The recommendations in Apple's Mac OS X Security Configuration for Version 10.5 Leopard and 10.6 Snow Leopard track closely with the security level historically represented in NSA guidelines. It is the NSA's belief that the guide produced by the manufacturer establishes the best practices for securing the product and recommends that traditional customers of its security recommendations use the Apple guide when securing either version of the Mac OS X systems.

Its website (www.nsa.gov/ia/guidance/security_configuration_guides/operating_systems.shtml#AppleMac) does provide some spectacular links to some resources for hardening your Mac OS X 10.5 and 10.6. We suggest you go and check them out.

Microsoft Security

When you talk about Microsoft security, it is not an easy thing to pin down. With the many operating systems, back office products, database servers, and so on, you need to focus this beam a little. Then the problem becomes how can you narrow this beam and make sure it is accurate and pertinent. This is a difficult task. Following are broad categories and an overview, which include links to the more specific items as they pertain to Windows XP or Microsoft Windows 7.

Anything older than Microsoft Windows XP and Microsoft Windows Vista is not included. Microsoft Windows XP is still a stable operating system, somewhat secure (if updates are applied regularly), and viable in the workplace. Microsoft Windows 7 is becoming the new de facto system in office environments; it is more stable, secure, and up to date than Microsoft Windows XP and requires less maintenance to maintain its security posture.

Security Policies

Are we talking about account policies, local policies, or audit policies? And then which operating system? Microsoft Windows 2003 Server or Microsoft Windows XP Professional? To be fair, the major components of three operating systems are covered: Microsoft Windows XP Professional, Microsoft Windows Server 2003, and Microsoft Windows 7 Professional.

Microsoft Windows XP Professional

For Microsoft Windows XP Professional, there is a security guide detailing various aspects of authorization and access control. The Microsoft Windows XP Professional operating system includes a number of features that you can use to protect selected files, applications, and other resources from unauthorized use. These features include access control lists (ACL), security groups, and local and group policies. These tools provide a powerful access control infrastructure for your organization's network infrastructure. If you have Microsoft Windows XP Professional systems deployed as your working desktop environment, we highly recommend reading through the security guide (<http://technet.microsoft.com/en-us/library/bb457115.aspx>).

Microsoft Windows Server 2003

For Microsoft Windows Server 2003, you can manage and secure several features through Group Policies—things such as account settings, event audit settings, software environment, wireless, and so on. When you edit security settings in a local GPO, only the security policy settings on that computer are affected. When you edit security policy settings in a GPO in Active Directory directory service, the policy settings affect sites, domains, and organizational units (OU) to which the GPO is linked. Some settings, such as password policy settings, are operative only at the domain level. You can find some guidelines at the following sites:

- **Security and Protection Overview:** [http://technet.microsoft.com/en-us/library/dd582586\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd582586(WS.10).aspx)
- **Security Policy Settings:** [http://technet.microsoft.com/en-us/library/cc739328\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc739328(WS.10).aspx)
- **Security Policy Planning and Architecture Best Practices:** [http://technet.microsoft.com/en-us/library/cc739214\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc739214(WS.10).aspx)

Microsoft Windows 7

There have been many changes and improvements to Windows 7 over Windows XP and Windows Vista. These changes include new features such as BitLocker, AppLocker, multiple active firewall profiles, user account control, Internet Explorer security features, auditing enhancements, safe unlinking in the kernel pool, Windows Biometric Framework, smartcard support, and service accounts. Table 4-2 gives a brief overview.

Table 4-2 *Security Improvements in Windows 7*

Improvement	Description
BitLocker	Encrypts entire volumes, including system volume, nonsystem volumes, and removable drives.
AppLocker	Provides flexible control over which applications users run.
Multiple active firewall/ASA profiles	Provides different firewall/ASA profiles for the physical network adapter and virtual network adapters used by VPNs.
User Account Control (UAC)	Gives standard users the opportunity to provide administrative credentials when the operating system requires them. For administrators, it runs processes with standard privileges by default and prompts the administrator to confirm before granting administrative privileges to a process.
Internet Explorer security features	Reduced risk of phishing and malware attacks when browsing the Internet.
Auditing enhancements	Provides more control over which events are monitored.
Safe unlinking in the kernel pool	Reduces the risk of overrun attacks.
Windows Biometric Framework	Provides a uniform interface for fingerprint scanners.
Smart cards	Provides support for a standard smart card driver interface.
Service accounts	Enables administrators to create accounts for services without needing to manage service account passwords.

The collection of Windows 7 security and protection guidelines, which you can find at [http://technet.microsoft.com/en-us/library/dd571075\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd571075(WS.10).aspx), provides detailed information about security features listed in Table 4-2. These features enable the IT professional to design, deploy, and maintain Windows 7 desktop environment in a secure fashion.

However, Windows 7 is not only deployed at the office, but also many individuals use Windows 7 at home. If you need help and how-to security information for using Windows at home, see Windows Help and How-to at <http://go.microsoft.com/fwlink/?LinkId=168437>.

Windows Server 2008

Microsoft has replaced the Windows Server 2008 security guide with the more intuitive Windows 2008 Security Compliance Management Toolkit, as part of its Security Management Compliance Toolkit series. You can find the toolkit at www.microsoft.com/download/en/details.aspx?DisplayLang=en&id=17606.

Many tools Windows Server 2008 deploys can assist the administrator in planning, deploying, and administering a Windows 2008 server securely. Following is a link that

provides guidance for everything from managing user accounts to diagnosing overall system security: Windows 2008 Security Tools: [http://technet.microsoft.com/en-us/library/cc722416\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc722416(WS.10).aspx).

Microsoft Security Compliance Manager

Similar to the Cisco Validated Design (CVD) program mentioned earlier in the chapter, Microsoft has the Security Compliance Manager. This product provides centralized security baseline management features, a baseline portfolio, customization capabilities, and security baseline export flexibility. This enables you as an administrator or the chief information security officer (CISO) to increase your organization's capability to efficiently manage the security and compliance process for the most widely used Microsoft technologies.

The Security Compliance Manager is the next evolution in security from Microsoft. It incorporated its previous guidance and documentation into this tool. This tool enables you to access and automate all your organization's security baselines in one central place.

This tool enables you to access the complete database of Microsoft recommended security settings, customize your baselines, and then export them in various forms, including XLS, Group Policy Objects (GPO), Desired Configuration Management (DCM) packs, or Security Content Automation Protocol (SCAP). Table 4-3 outlines the key benefits and features of the Microsoft Security Compliance Manager.

Table 4-3 *Microsoft Security Compliance Manager Key Features and Benefits*

Key Feature and Benefit	Description
Centralized Management and Baseline Portfolio	The centralized management console of the Security Compliance Manager provides a unified tool to plan, customize, and export security baselines. The tool gives you full access to a complete portfolio of recommended baselines for Windows client and server operating systems and Microsoft applications. The Security Compliance Manager also enables you to quickly update the latest Microsoft baseline releases and take advantage of baseline version control.
Security Baseline Customization	Enables you to duplicate any of the recommended baselines from Microsoft—for Windows client and server operating systems and Microsoft applications—and quickly modify security settings to meet the standards of your organization's environment.
Multiple Export Capabilities	Export baselines in formats such as Excel .XLS files, group policy objects (GPO), desired configuration management (DCM) packs, or security content automation protocol (SCAP) to enable automation of deployment and monitoring baseline compliance.
Standard Security Baselines	Security baselines and security guides for Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, Hyper-V, Windows 7, Windows Vista, Windows XP, BitLocker Drive Encryption, Windows Internet Explorer 8, Microsoft Office 2010, and Microsoft Office 2007 SP2.

Following is a link to the Security Compliance Manager:
www.microsoft.com/download/en/details.aspx?displaylang=en&id=16776, which you should download, install, and use.

Chapter Summary

This chapter explained what you should consider to secure your networking infrastructure, starting with security tools from Cisco and working through the various other industry security guides from NIST and the NSA. Finally, you learned some standard settings of Microsoft and Apple from their own security guides.

You should take away from this section the importance of positive control of all your devices to ensure that no one can tamper with the network by reconfiguring the devices, attacking a router or firewall/ASA, and ultimately getting into your network infrastructure to wreak havoc.

General concepts and specific features used in Cisco devices and Microsoft operating systems were explained, and you were shown how to incorporate additional elements of a security architecture, including integrity, confidentiality, availability, and audit.

Chapter Link Toolbox Summary

As with the previous chapter, there are no questions here; instead, the following list provides a central repository for all the links provided in the chapter:

Cisco SAFE: www.cisco.com/go/safe

Cisco Validated Design (CVD) Branch and WAN Guide www.cisco.com/en/US/netsol/ns816/networking_solutions_program_home.html

Cisco Validated Design (CVD) Campus Design Guide:
www.cisco.com/en/US/netsol/ns815/networking_solutions_program_home.html

Cisco Validated Design (CVD) Data Center Design Guide:
www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html

Cisco Validated Design (CVD) Security Design Guide:
www.cisco.com/en/US/netsol/ns744/networking_solutions_program_home.html

Cisco HSRP Configuration Guide: www.cisco.com/en/US/docs/internetworking/case/studies/cs009.html

Cisco Guide for configuring TCP Intercept and Preventing Denial-of-Service Attacks:
www.cisco.com/en/US/docs/ios/sec_data_plane/configuration/guide/sec_cfg_tcp_intercpt.html

Configuring AAA on your Cisco PIX/ASA Guide: www.cisco.com/en/US/docs/security/pix/pix63/configuration/guide/mngacl.html#wp1090040

Configuring Failover on your Cisco PIX/ASA: www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_configuration_example09186a00807dac5f.shtml

Cisco PIX/VPN Configuration Guide: www.cisco.com/en/US/docs/security/pix/pix63/configuration/guide/config.html

Cisco ASA Configuration Guide: www.cisco.com/en/US/docs/security/asa/asa71/configuration/guide/basic.html

Configuring your IPS: www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6634/prod_white_paper0900aecd805c4ea8.html

NSA Security Guide for Cisco Routers: www.nsa.gov/ia/guidance/security_configuration_guides/cisco_router_guides.shtml

NSA Security Guide for Cisco Switches: www.nsa.gov/ia/guidance/security_configuration_guides/switches.shtml

NSA Security Guide for VoIP/IP Telephony: www.nsa.gov/ia/guidance/security_configuration_guides/voip_and_ip_telephony.shtml

NSA Security Guide Microsoft Windows: www.nsa.gov/ia/guidance/security_configuration_guides/operating_systems.shtml#microsoft

NSA Security Guide for Apple: www.nsa.gov/ia/guidance/security_configuration_guides/operating_systems.shtml#AppleMac

Microsoft Security Guidelines for Windows XP Professional: <http://technet.microsoft.com/en-us/library/bb457115.aspx>

Microsoft Security Guidelines for Windows Server 2003 Security and Protection Overview: [http://technet.microsoft.com/en-us/library/dd582586\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd582586(WS.10).aspx)

Microsoft Security Policy Settings for Server 2003: [http://technet.microsoft.com/en-us/library/cc739328\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc739328(WS.10).aspx)

Microsoft Security Policy Planning and Architecture Best Practices for Windows Server 2003: [http://technet.microsoft.com/en-us/library/cc739214\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc739214(WS.10).aspx)

Microsoft Windows 7 new security features: [http://technet.microsoft.com/en-us/library/dd571075\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd571075(WS.10).aspx)

Microsoft Server 2008 Security Guide: www.microsoft.com/download/en/details.aspx?DisplayLang=en&id=17606

Microsoft Windows Server 2008 Security Tools: [http://technet.microsoft.com/en-us/library/cc722416\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc722416(WS.10).aspx)

Microsoft Security Compliance Manager: www.microsoft.com/download/en/details.aspx?displaylang=en&id=16776

Overview of Security Technologies

“We can’t help everyone, but everyone can help someone.” —Ronald Reagan

This chapter discusses the use of technologies that have evolved to support and enhance network security. Many of these technologies are used today without the user understanding when or where they operate. After reading this chapter, you will understand the benefits of these technologies, where they operate, and some of the operational risks associated with them. By the end of this chapter, you should know and be able to explain the following:

- How you can employ packet filtering to reduce threats to a network
- Understand precisely what stateful packet inspection is, and why its important for firewalls to use this technique
- The role and placement of a proxy technology within a secure network
- Network Address Translation (NAT) and how you can use it to allow the Internet to continue to grow in IPv4
- How Public Key Infrastructure (PKI) has the potential to protect the flow of information in a global manner

Answering these key questions and understand the concepts behind them will enable you to understand the overall characteristics and importance of the security technologies covered in this chapter. By the time you finish this book, you will have a solid appreciation for network security, its issues, how it works, and why it is important.

So far, this book has painted in broad strokes the steps an attacker could possibly take to gain access to sensitive resources. The first step in protecting these assets is the global security policy created by combining the many aspects discussed in Chapter 2, “Security Policies.” This chapter introduces some of the more broadly used security technologies. Each of these technologies contains a concept or specific role that increases the security of your network when designed and implemented in a layered design.

Security First Design Concepts

Network security can be a hydra (many-headed beast) with regard to potential attacks and threats against the network. The resources and opinions on this subject are incredible, and opinions vary greatly depending on whom you ask. For example, in 2004 when I wrote the first edition of this book, a simple Google search on “designing a secure network” returned almost half a million results. In 2012, that same search string returns more than five and a quarter million hits. It is no wonder that conflicting security concepts bombard people, causing a great deal of confusion. To be honest, if you were to look up network security books, any bookstore also reveals almost as many!

The point is that experts in each area of network design have written so much on designing secure network architecture that to try to do the subject justice here is beyond the scope of this book. Books and websites deal with every aspect of network security, server security, application security, and so forth. We endeavor to provide you with a strong foundation upon which to build the security knowledge required for your role or network.

This book illustrates good network security design principles to build the strongest possible foundation. However, it covers some important design concepts of which you must be aware:

- **Layered security:** A network that implements layered security truly understands that a single point of defense is doomed to eventual failure. Thus, as Figure 5-1 demonstrates, consistently implementing security throughout a network at as many points as possible is considered good design. This concept of layering a network’s security is the single-most important design concept in this chapter and is often referred to as *Defense in Depth*.
- **Controlling access:** The network is ultimately your responsibility and, as a result, you determine what is allowed into and on your network. One highly recommended practice is to make access decisions with the mindset of “block everything, and allow only what is needed to conduct business.” This has also been referred to as the Policy of Least Privilege (POLP). This is the default action of Cisco firewalls and access control lists (ACL).
- **Role-specific security (Role Based Access Control [RBAC]):** When deciding upon access and privilege (that is, trust), one of the most useful templates to use is based on a user’s role within the organization. For example, a web developer would clearly need access to the organization’s website, whereas an administrative assistant would not.
- **User awareness:** Stories abound about users writing down passwords, changing them five times in a row, and then using their original password again. It is not that users are intentionally bypassing security; they do not understand the purpose of the security and may have become complacent. Okay; let’s be honest; some users definitely try to bypass security, but more on that later! Thus, user awareness through training and visibility is essential to get users to understand the importance of security. One great idea for getting users to attend training and learn why it is important is to serve ice cream with all the trimmings. This method appeals to a basic human love of

sweets, but it is also effective and fun; you will become a popular person! It is crucial to have your user truly aware of security and supportive of security policies; making security training a pleasant experience can help make that happen.

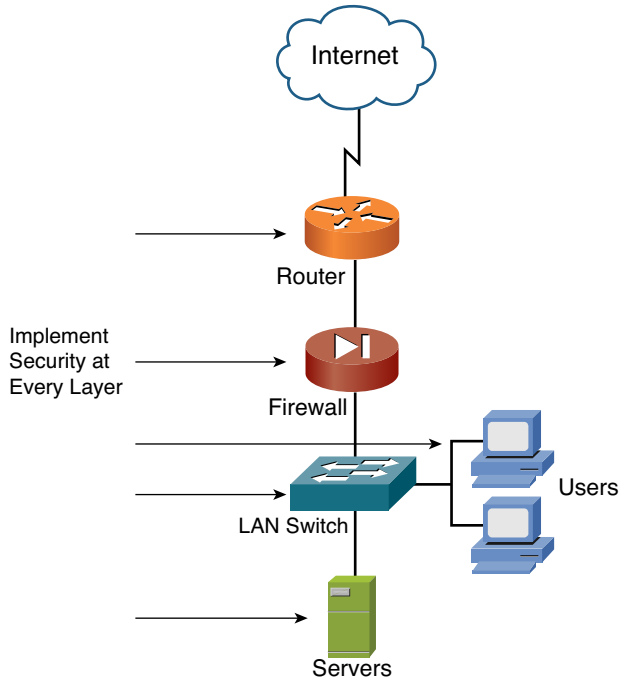


Figure 5-1 *Layered Security Points (Defense in Depth)*

- **Monitoring:** Perhaps one of the most forgotten aspects of security is monitoring. Many organizations believe that it is enough simply to have security. They forget that monitoring their systems to ensure that they remain secure and are not subject to attack is also crucial. The truth is, security devices report every little thing, and it's hard to do an effective job if you're not listening and monitoring what they are saying. One of the ways to achieve this is to "tune" the device; another is to have every device on the network report to a central device that you tune and monitor. It is much easier to monitor one device than ten. Cisco has an effective product for this, referred to as *Cisco Security Manager*. More information on the Cisco Security Manager is available at www.cisco.com/en/US/products/ps6498/index.html. Chapter 11, "Intrusion Detection and Honeypots," discusses the methods used to monitor for attacks: intrusion detection systems (IDS). A strongly recommended practice is to include provisions for IDS when designing a network's security solution in wired or wireless networks.
- **Keep systems patched:** Patching or updating systems is a fundamental task that is often forgotten by system administrators with their busy schedules. Fortunately,

many newer operating systems can remind you when new updates are available. For example, I use an Apple Mac Book Pro running OS X (aka Snow Leopard); within this operating system is a built-in functionality that automatically checks for updates, as shown in Figure 5-2.

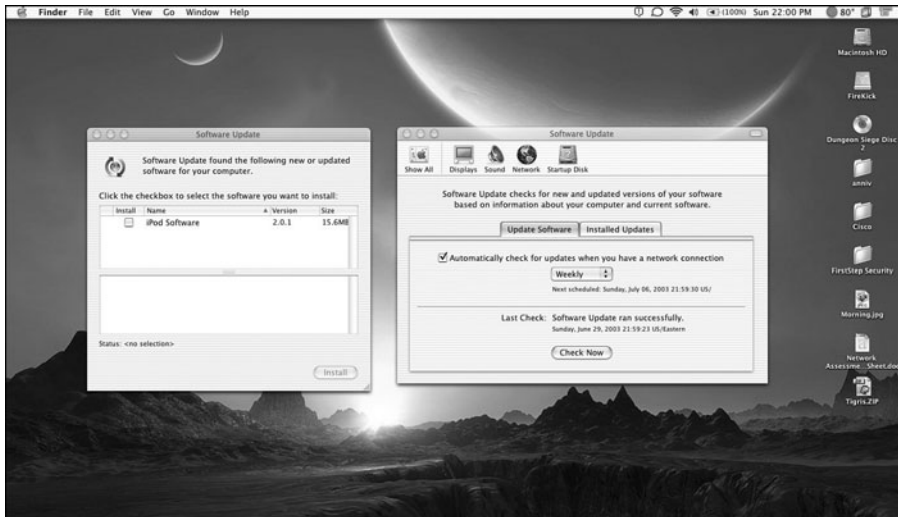


Figure 5-2 *MAC OS X Automatic Update Functionality*

The only downside in this example is that I do not yet have an Apple iPod, which would require this update. Regardless, you can understand the point: Always make the time to check for patches for your systems because hackers are always pushing to find and exploit. For Windows users, Microsoft has also included this automatic update functionality in newer versions of its operating systems. The trouble is that Microsoft set the auto updates to occur at 3 a.m. by default. This is great if you leave your computer on 7x24, but if you're like most people, you shut it off when done using it. The moral to this story is that auto updates are good, but be proactive to ensure that they are happening, and at the right time! Patches must also be tested before inserted into production networks, and not all systems are as patch friendly as others. You need to understand where your patches are coming from (in some cases they are hashed) so that you can be sure they are not malicious code masquerading as a patch to one or more of your critical systems. Apple handles updates in a more elegant manner, as shown in Figure 5-2.

- **Incident Response Teams:** Security concerns will inevitably be brought to you in some form or another. Perhaps your systems have become the target of an attack or you have detected that the compromise and damage has already been done. This aspect of design deals with how an organization responds to an attack and deals with whatever situation it experiences. It is best to include and consider incident response teams and the process of responding in practice rather than when you are under pressure and the situation is extreme. So, design it now; the benefits come later. Practice

makes perfect, and dry runs can help point out a plan's flaws that do not seem evident at the time the plan and policy is written.

These first-step security design considerations will enable you to understand how to begin securing any network. The next section begins to discuss the specifics of how you can use security technologies and their roles in protecting a network.

Packet Filtering via ACLs

As you probably already know, all information that flows across the Internet uses TCP/IP and, in turn, this information is sent in small pieces known as *packets*. In the early days of the Internet, filtering based on packets was common and, in many cases, routers in many networks still use packet filtering. Packet filters are often used as a first defense in combination with other firewall technologies. Today, their most common implementation is seen in the ACLs of routers at the perimeters of networks.

Packet filtering is one of the oldest and most common types of packet inspection technologies available. It begins by inspecting a packet's contents and applying rules to determine whether a packet should be dropped or allowed. Although many characteristics are possible within a TCP/IP packet's header (that is, protocol, port, and so on), this discussion refers to filtering based on the source or destination IP address, as shown in Figure 5-3.

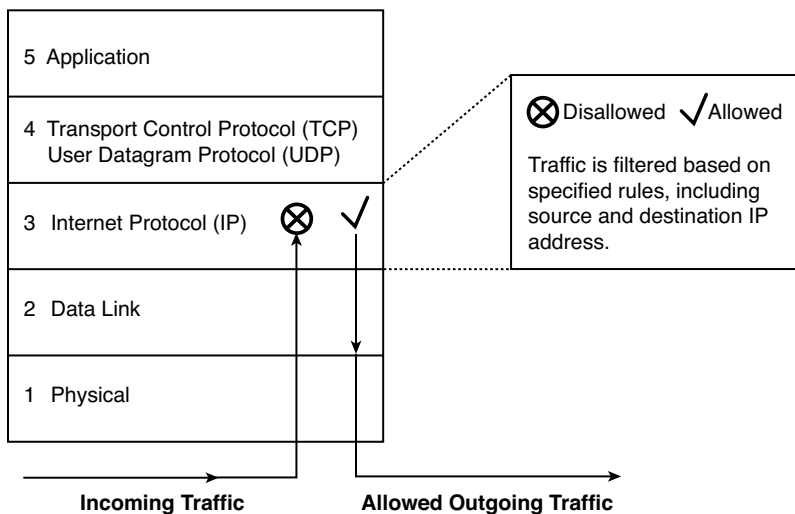


Figure 5-3 *Packet Filtering at Layer 3 of the TCP/IP Model*

The two main types of ACLs are standard ACLs, which filter based on IP address, and extended ACLs, which look further into a packet header, if so configured.

Note Standard ACLs are source address–based and extended ACLs are source-based and destination-based and have more capabilities, such as specifying port or protocol. The following ACL styles for IP are supported:

- **Standard IP ACLs:** Use source addresses for matching operations
- **Extended IP ACLs (control plane only):** Use source and destination addresses for matching operations and optional protocol type and port numbers for finer granularity of control
- **Named ACLs:** Use source addresses for matching operations

Refer to the following URL for more information about configuring ACLs and Cisco devices (Cisco.com account required): www.cisco.com/en/US/partner/products/sw/secursw/ps1018/products_tech_note09186a00800a5b9a.shtml#types

Packet filters inspect each packet individually, examining source and destination IP address and ports as defined in the filter. Only the beginning of each packet header is examined; for this reason, they can quickly decide packet flow because the packet is read only enough to determine whether it is a match. The characteristics of each of these inspection points determine whether the given packet should be allowed or denied. The use of ACLs is how packet filtering is conducted on Cisco devices; they are one of the focal points of this section.

Because every packet of every connection is checked against the access control rules, larger, complex packet-filtering rule bases could decrease performance of the device upon which they are applied. In addition, because packet filters can check only low-level attributes, they are not secure against malicious code hiding in the other layers.

The use of ACLs is one of the most confusing topics to many. As you see in the following section, a good understanding of ACLs can be less confusing when superimposed over a good analogy that relates to real life.

Grocery List Analogy

This analogy based on going grocery shopping is just one way to introduce and explain the concepts behind packet filtering via ACLs. You must consider certain key principles while considering this grocery list analogy. Table 5-1 begins the analogy by comparing packet filtering via ACLs with creating a grocery list.

In planning a turkey dinner, my wife and I discovered that we needed some things to finish cooking; we decided to make a list. This way, I would not forget what we needed when I went to the store. We knew that we had the following things, so they are not going on the grocery list:

- Turkey
- Stuffing

- Bread
- Cheese

Table 5-1 *Access List/Grocery List Analogy Overview*

ACL Characteristics	Grocery List Analogies
ACLs are effective	Following a grocery list is efficient and saves money.
Top-down processing	The order of the items on the list is important.
Place denials first	There are items not on the list, so do not buy them.
Always have a permit	A list must always include things that are permitted.
Implicit deny all	You can buy only what is on the list.

In other words, I cannot buy these ingredients because my wife says that we do not need them. When I make a list of the things I am allowed to buy, my list is rather broad. I am happy with the list; it will do the job, so I am ready to head to the grocery store to get the following items:

- Milk
- Pie
- Potatoes
- Gravy
- ****Buy nothing else****

This list is broad because there are many types of milk and *many* types of pies and because of how the list is written, I can buy any sort of pie I want because they are all allowed. She just might be in trouble because I happen to enjoy mincemeat pies and she does not! Because we need these ingredients, I can buy them. This broad grocery list analogy can relate directly to a standard ACL when expressed as follows:

[standard acl] Regular Grocery List

[deny] Turkey

[deny] Stuffing

[deny] Bread

[deny] Cheese

[permit] Milk

[permit] Pie

```
[permit] Potatoes
[permit] Gravy
[implicit deny all else] **Buy nothing else - end**
```

Notice the last line; my wife imposes this restriction on me because I have a great deal of affection for chocolate ice cream and on-sale items. Now, she does not need to actually *say the words to me* because I implicitly understand that I am not allowed to buy anything else.

I decide to show my list to my wife to make sure I did not miss anything. She reviews the list and decides I need more specific instructions because it is important to buy the right “kind” of groceries. She begins writing on my list:

```
[extended acl] Extended Grocery List (that is, wife's version)
[deny] Turkey
[deny] Stuffing
[deny] Bread
[deny] Cheese
[permit] Milk – 2% White
[permit] Pie – Mrs. Smith's Pumpkin
[deny] Potatoes – Red because a guest is allergic to this type
[permit] Potatoes – Any potatoes other than red is okay
[permit] Gravy - White Country
[implicit deny all] **Buy nothing else - end**
```

This type of list allows for a more granular level of filtering or, in my case, a more rewarding return home with the ingredients I was permitted to buy. Did you notice the difference between the two lists? The first list was rather broad and not specific at all, whereas the second list was extremely specific and told me not only exactly what not to buy, but more specifically what I was permitted to buy. Ultimately, the implicit understanding is that everything else is denied. You probably relate to the challenges of shopping when you are married and are also wondering how this relates to ACLs and packet filtering.

Packets have identifiable characteristics that access lists use to classify them and take an action—either permit or deny. Consider Example 5-1, which shows what a standard access list based on my analogy might look like.

Example 5-1 *Analogy as a Standard Access List*

```
access-list 10 deny any turkey
access-list 10 deny any stuffing
access-list 10 deny any bread
access-list 10 deny any cheese
```

```

access-list 10 permit any milk
access-list 10 permit any pie
access-list 10 permit any potatoes
access-list 10 permit any gravy

```

The standard access list in a Cisco device is primarily used to filter packets based on IP addresses. In addition, numbering them identifies a standard access list; specifically, they use 1–99 and 1300–1399 as identification numbers. If you were to take this example a technical level deeper and use IP addresses and subnets, it would look like Example 5-2 in a Cisco device’s configuration.

Example 5-2 Standard Access List Filtering Packets

```

access-list 10 permit any 192.168.10.0
access-list 10 permit any 192.168.20.0
access-list 10 permit any 192.168.30.0
access-list 10 permit any 192.168.40.0

```

You are probably wondering what happened to the **deny** statements. With Cisco ACLs, there is that implicit **deny** everything else at the end, which you do not “see” in the configuration. Thus, you do not have to enter the **deny** statements. You could take the standard ACL and expand it to be even more specific by using an extended ACL; this is what my wife did when she gave me more specific instructions.

Because they are designed to identify packets, ACLs fulfill many roles in the world of networking. After a packet is identified, it can be acted upon in some manner. This action might include sending it after a more important packet, or perhaps filtering the packet. Figure 5-4 shows the placement of an ACL to filter packets.

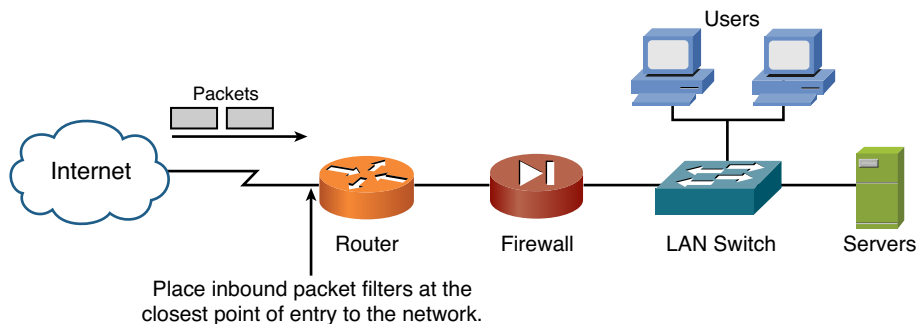


Figure 5-4 Placement of Packet Filters

If you consider the analogy of the entrance to my local grocery store to where the packets are entering the router, you can understand that nothing is getting in without permission!

A secure router at the edge or perimeter of your network might be your first step/layer in a strong defense-in-depth methodology.

Limitations of Packet Filtering

It is time to talk about the drawbacks of using packet filtering. Certainly, you can stop many things with their use. Consider that you have a web server in a DMZ; all web/HTTP traffic must be able to reach this server. This server happens to run Microsoft's IIS web server software, and an attacker decides to directly attack the web server using web/HTTP traffic. Because the attack targets vulnerabilities in IIS, the packets are allowed. So, although packet filtering is not enough security (on its own), it most certainly is another technique that will increase the depth of your networks security by creating another layer of protection.

Note You can find additional ACL information and techniques at the following Cisco.com URL (Cisco.com account required). The article is titled "Protecting Your Core: Infrastructure Protection Access Control Lists": www.cisco.com/en/US/partner/tech/tk648/tk361/technologies_white_paper09186a00801a1a55.shtml

The next section takes packet filtering a step further by discussing stateful packet inspection.

Stateful Packet Inspection

This section discusses the more advanced technique of packet inspection: *Stateful Packet Inspection (SPI)*. To understand how SPI operates, you must briefly review the TCP/IP model.

Note Many people are confused about the relationship between the OSI reference model and the TCP/IP model—simply put, the use of OSI is a reference for developers whereas, in education, functionally TCP/IP is used. Therefore, you must use the TCP/IP model when inspecting packets.

Figure 5-5 shows the five layers of the TCP/IP model. The stateful inspection component is concerned with how TCP (Layer 4—transport) makes connections. Tracking the state of the TCP connection is done via Layer 4 of the TCP/IP model.

In most cases, SPI occurs in a firewall, which sits behind the secure router that connects your network to the Internet. If you have implemented packet filtering with ACLs on the router as your first line of defense (and you should), the next line/layer of defense will be SPI at the firewall, as shown in Figure 5-6.

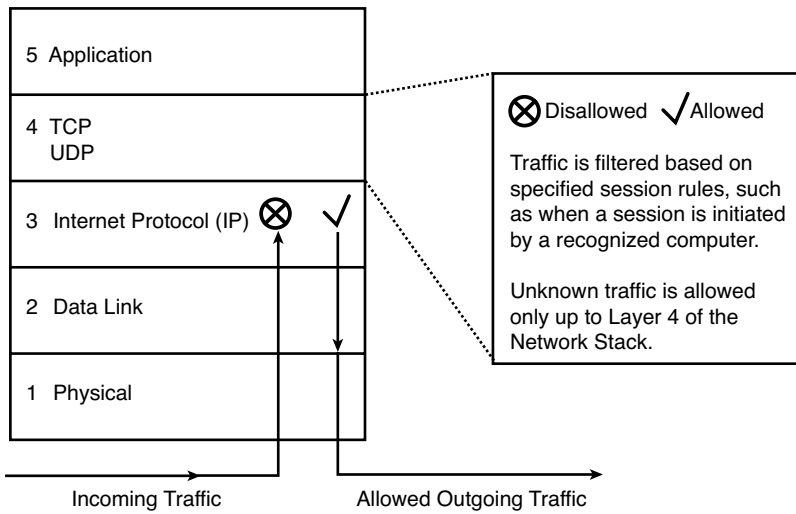


Figure 5-5 TCP/IP Model

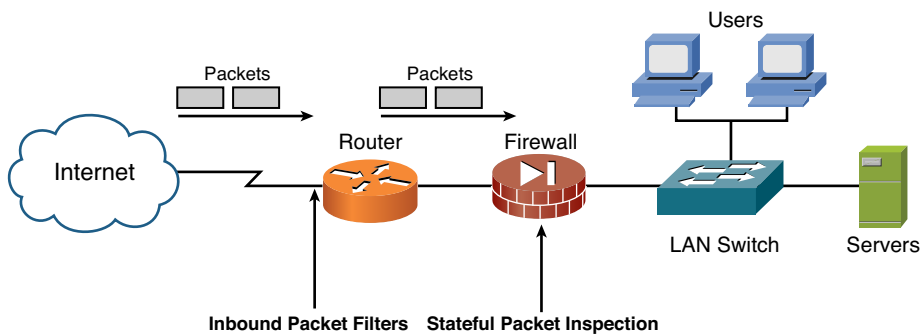


Figure 5-6 Placement of Stateful Packet Inspection

Note There is an Internet standard known as RFC 2827, which can guide you through the process of creating your first line of defense. This RFC is titled “Network Ingress Filtering: Defeating Denial of Service Attacks,” which employs IP Source Address Spoofing.

This placement and added security enables the defense in depth to be layered at yet another level, with the goal of completely securing the network via multiple layers of protection.

SPI is usually implemented in a firewall, so the TCP/IP connections can be inspected more closely. Thus, this technology is considered *connection-aware* in that SPI monitors and understands that a connection between two computers usually consists of many packets that flow back and forth between the computers. This connection-aware functionality happens because the firewall is tracking every connection that comes into it and out of it

to track the state of the connection. Yes, this connection was opened by one of my internal users (permit) or no, it was not opened (deny).

Stateful inspection of packets occurs during the first packets used to create this connection. As the connection is inspected, an entry is created in a table. Then, as future packets are received, they are verified against entries in this table to see whether they belong to an existing and recorded connection. If the packets pass this verification phase, they are allowed to pass. At a high level, that is how SPI occurs. The following section examines this process in more detail.

Detailed Packet Flow Using SPI

Because this book strives to always present best practices regarding network security and the associated technologies, this more detailed discussion is based on the assumption that the external router is in place and that it is configured to prescreen connection attempts into the network by using packet filtering. Therefore, picking up the packet as it passes through the router and its packet filtering, the next step is the packet arriving at the firewall:

1. When a packet arrives at the firewall, a decision must be made to determine whether the packet should be allowed (forwarded) to the internal network.
2. The device performing the stateful packet inspection takes each arriving packet and inspects its headers to determine whether they match the set of rules that control what kind of packets are allowed.
3. When inspecting the packet's headers, the inspection includes the packet's source and destination addresses, its protocol type (TCP, UDP, ICMP, and so forth), its source and destination ports, flags set on the packet (SYN, ACK, FIN, RST, and so on), or other such basic header information. Incoming packets are inspected until enough information has been gathered from the packets received (using information such as TCP sequence numbers) to determine the connection's "state."
4. This inspection data is compared against the rule set that has determined what should be allowed and what should be denied. For example, all HTTP traffic *only* might be allowed to a web server, whereas other traffic should be denied trying to access the web server. This is a common rule wherein only a certain type of traffic should be allowed to only a certain server.
5. Depending on the connection status, this inspection information is then compared to a stateful table that would have entries for each TCP/IP connection the device has enabled. For example, most devices enable everyone from inside the network to access anything they want outside the network, and that connection would have formed an entry in the state table. Rather than enabling all packets that meet the rule set's requirements to pass, only those packets that are part of a valid, established connection are permitted.

6. Ultimately, packets are either permitted or denied depending on these inspection steps. Because these rules/tables are consulted only once, complex inspection rules do not greatly impact performance.
7. All permitted and denied access should be logged to a secure syslog server that has accurate NTP sync. These logs can be fed into a security information management system for further analysis and reporting. In Cisco Security Manager (CSM), rules can be audited and hit counts analyzed to make sure that rule usage is being monitored, templates are followed, and there is no rule overlap or mistakes in existing rules.

SPI rules are not as easy to create as packet-filtering rules because of the added level of complexity. However, they are certainly worth the money and effort because they add an additional level of security to your network. They are also fast and can handle large amounts of network traffic. If the metrics recorded for the connection do not match the entry in the connection database, the connection is dropped.

Note Usually, firewalls are the devices of choice for performing stateful packet inspection; however, routers can also be used in this role. However, this is not advised because mixing network devices' roles alters the functions they were designed to perform. Some might argue that you can successfully combine roles and devices; perhaps this might be appropriate in the distant future—for today and for the networks I am responsible for securing, I advise against it.

Limitations of Stateful Packet Inspection

Although SPI devices have improved scalability and benefits over packet filtering, they are not the ultimate point of protection for your network; again they are but a layer in a layered defense. Consider the following two major disadvantages of stateful packet inspection:

- **No application-level inspection:** SPI cannot look at a packet any higher than Layer 4 of the OSI reference model. In practice, this is how attacks can succeed against servers that are accessible in some manner and protected by firewalls performing stateful packet inspection. Keep in mind that many attacks today are focused on Layer 4 and higher.
- **No connection state for every TCP/IP protocol:** Certain protocols within TCP/IP have no method of tracking the state of their connection between computers. Specifically, ICMP and UDP have no connection state; thus, in the layered defense model, these protocols should be subjected to packet filtering because they have no connection state to track.

This section discussed the capability of security devices, such as firewalls, to track the state and thereby the validity of a connection to determine whether it should be allowed into the protected area of your network. The next section focuses on the various means of

further ensuring the validity of packets entering your network by using additional security to inspect them at Layer 5 (application) of the TCP/IP model or Layer 7 of the OSI model to provide a map of the layers in the models.

Network Address Translation (NAT)

The Internet has grown larger than anyone ever imagined. Although its exact size is unknown, the current estimate is that there are approximately 100 million hosts and more than 350 million users actively on the Internet. This is more than the entire population of the United States. The Internet is effectively doubling in size each year.

When IPv4 addressing first appeared, everyone thought there were plenty of addresses to cover any need. Theoretically, you could have 4,294,967,296 unique public addresses (2³²). The actual number of available public addresses is smaller (somewhere between 3.2 and 3.3 billion) because of the way the addresses have been separated by the Internet Engineering Task Force (IETF) into classes (A, B, C) and the need to set aside some of the addresses for multicasting, testing, or other specific uses (Class D).

Note In addition to arranging groups of IPv4 addresses into classes, you might be wondering what happened to the millions of public IPv4 addresses that I said were no longer available. To ensure that every network in need of private IP addresses can have them, the Internet Engineering Task Force (IETF) has set aside a large range of addresses for internal network routing by means of Network Address Translation (NAT). Many of these addresses are referred to as private IP addresses; these addresses are not accessible on the public Internet, thus the word *private*. Private addresses are to be used within any organization that needs them and never used (routed) on the Internet. The addresses used (routed) on the Internet are referred to as public IP addresses.

With the explosion of the Internet and the ever-increasing need for IP addresses in home networks and business networks, the number of available IPv4 addresses is simply insufficient. The obvious solution is to redesign the IP addressing scheme to allow for more possible addresses. This is being developed in a solution known as IPv6, but it will take many years to implement because it requires the modification of the Internet's entire infrastructure. As a result, the process of converting from IPv4 to IPv6 has been slow and will likely continue slowly as NAT further extends the life of IPv4. Because of the massive number of addresses that IPv6 provides, NAT was not built into IPv6 initially. RFC 6052 is the most recent update for IPv6.

NAT enables organizations to resolve the problem of IP address depletion when they have existing networks and need to access the Internet. Sites that do not yet possess NIC-registered IP addresses must acquire them from the Internet Assigned Numbers Authority (IANA) and American Registry for Internet Numbers (ARIN), who delight in causing bureaucratic delay. Many sites do not pass ARIN's bureaucratic detailed examination or justification process and are denied public IP addresses; therefore, NAT is the solution for most organizations.

Note The IANA has reserved the following three blocks of the IP address space for private networks:

- 10.0.0.0–10.255.255.255 (10/8 prefix)
- 172.16.0.0–172.31.255.255 (172.16/12 prefix)
- 192.168.0.0–192.168.255.255 (192.168/16 prefix)

NAT enables companies to use *public IP addresses on the outside of the network* (that is, on those devices that connect directly to the public Internet). However, as discussed, there probably will not be enough public IP addresses for every network printer, PC, server, switch, router, wireless device, and so forth to be assigned a public IP address.

These devices need an IP address to connect with TCP/IP, so we use *private IP addresses on the internal network*. The use of private IP addresses inside our network provides for all devices to now communicate using TCP/IP, which was the goal. However, you must activate NAT because the private IP addresses are not allowed out onto the Internet.

NAT is deployed and implemented on a device (firewall, router, or computer) that sits between an internal network using private IP addresses and the Internet, which uses public IP addresses. The device performing the address translation from private to public is usually a firewall and, to a lesser extent, a router. The device performing NAT usually sits with one part connected to the internal network and another part connected to the Internet (or some external network). Figure 5-7 shows the placement of NAT as part of a layered defense-in-depth architecture.

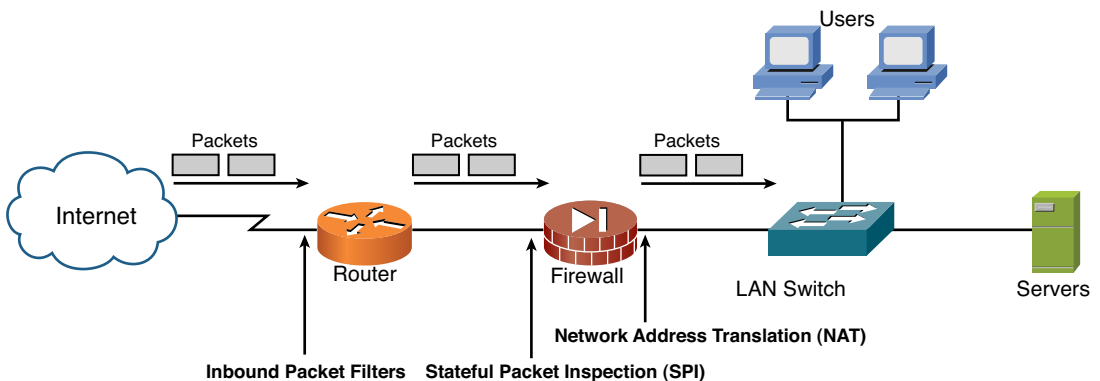


Figure 5-7 Placement of NAT in a Network

Discussion of how NAT also provides an additional level of security to your network is discussed later in the section “Increasing Network Security.” NAT has many forms and can work in several ways:

- **Static NAT:** Provides for mapping a private IP address to a public IP address on a one-to-one basis. This is particularly useful when a device needs to be accessible from outside the network; for example, if your web server has an internal IP address of (10.0.0.1) and it needs to be accessible from the Internet—it is your web server, after all! NAT must be *statically* configured to enable users who have only a single public IP address for it always to be translated to 10.0.0.1. The use of static NAT is quite common for devices such as web servers, which must always be accessible from the Internet.
- **Dynamic NAT:** Provides for mapping a private IP address to a public IP address from a group of registered IP addresses. In this type of NAT, there is a one-to-one relationship in the mapping from private to public. For example, if your PC is assigned an internal IP address of 10.0.0.2 and your co-worker is assigned 10.0.0.3, each of you would be assigned a public IP address at the firewall via NAT as your traffic went to the Internet. Dynamic NAT is helpful, but it might not be the right solution in many cases. For example, what if your other co-worker wanted to access the Internet and the firewall was out of available public IP addresses? He would be denied. This could introduce a serious problem; therefore, NAT overloading was developed.
- **NAT Overloading (aka PAT):** A form of dynamic NAT that provides for the dynamic translation of multiple private IP addresses to a single public IP address by using different TCP ports. This is also known as Port Address Translation (PAT) or single address NAT. Its many names are not important, but how it functions is crucial. Because, with 65,535 TCP ports possible per single IP address, NAT enables an effective means of providing Internet access to many users who have been assigned private IP addresses. This type of NAT is the most commonly used because it serves large numbers of users at once.

Increasing Network Security

Solving the IPv4 address depletion and waste problems was the leading reason for the development of NAT, which also provides for yet another layer of security to protect your network. In general, using NAT makes it slightly more difficult for an attacker to do the following:

- Discover and map the target’s network topology and determine connectivity
- Identify the number of systems running on a network
- Identify the type of machines and the operating systems they run
- Implement denial-of-service (DoS) attacks such as SYN (synchronize/start) flooding, port and service scans, packet injection, enumeration, and escalation of privilege on your network

NAT's Limitations

It is clear that the introduction of NAT to the realm of networking and the Internet has solved or at least extended the IP address depletion problem. Many people have asked whether networks will ever evolve to IPv6 now that NAT works so well. The question is not actually if, but when will this conversion take place. For example, the Asia/Pacific region of the world is leading the implementation of IPv6 with many networks already using it.

As connectivity and convergence increase, the need for additional IP addresses will grow and expand. We will therefore make the change to IPv6 eventually; NAT has simply delayed the inevitable. NAT is useful and has brought advantages; however, it does have some limitations:

- **Issues with UDP:** NAT tracks and controls connections based on state and, as discussed earlier in this chapter, UDP has no inherent mechanism to determine state because it is connectionless as a protocol. Thus, NAT has no way of knowing whether a packet is part of an ongoing conversation or an isolated transmission. NAT devices then need to guess at how long a conversation involving UDP should remain open after the last packet; this is known as the *idle time*. Cisco firewalls provide the functionality to set idle time on UDP sessions to limit such cases.
- **Sensitive protocols:** Some protocols hide, alter, or otherwise obscure aspects of the packets that NAT requires to properly perform the translation. For example, IPsec VPN, Kerberos, X-Window, remote shell, and Session Initiation Protocol (SIP) can have trouble operating through a NAT device. This trouble is caused by applications that have embedded IP addresses in the packets where this issue occurs. Cisco firewalls have special “inspect” for different protocols, such as Skinny for telephony, that enable these applications to work when the inspect command is activated.
- **Interferes with encryption and authentication systems:** Many data encryption systems attempt to ensure the integrity of packets by ensuring that packets were not tampered with in transit. By its design, NAT tampers with packets, thus causing encryption and authentication technologies to not work well with NAT (by default). This is commonly seen with IPsec VPNs when a VPN device expects unaltered packets but the user is behind a firewall performing NAT. This means my VPN packets leave my computer and get NAT'd to be sent off onto the Internet, and <boom> the VPN breaks.
- **Complicated logging:** When devices log through a device, the correlation of the logs requires users to understand the translations being performed by NAT. Correlation of system logs with the NAT system can thus become highly complicated and tedious to understand which internal systems were actually involved.
- **One size fits all:** If your organization is using PAT, and one person in the company authenticates to a protected resource outside your company, it's possible that the rest of your organization now has access to that resource as well. Remember that if you use PAT, you're using only one IP address that has been multiplexed using port num-

bers. The protected resource that requires authentication sees all conversations from your company as coming from the same IP address.

As NAT has matured, there have been ways of addressing many of these limitations allowing them to work seamlessly; the VPNs requirement for special exemption from the packet-checking process of IPsec is one example. The final point to reinforce is that NAT is useful in many regards, from enabling an entire company to access the Internet to providing an additional layer of security. If you go back to the network referenced in figures throughout this chapter, you can see that including NAT adds another layer of protection (refer to Figure 5-7).

The following section looks at how security can be further deepened through tools and technologies that look deeper into a TCP packet.

Proxies and Application-Level Protection

Stateful packet inspection firewalls are enhanced versions of basic firewalls that just do packet filtering. The devices discussed here provide additional enhancements by analyzing the packets at the application layer. As you can see, we started with simple packet filters, added more advanced stateful packet inspection, and now we look even deeper into the packet at the application data contained within the packet.

You can use several types or technologies to provide application layer protection, and they are known by many different names. Although each technology operates slightly differently, their goal is the same: to increase the security of your network.

Application-level firewalls provide the most secure type of data connections because they can examine every layer in the TCP/IP model of the communication process. To achieve this level of protection, these firewalls—also known as *proxies*—actually mediate and control connections by intercepting and inspecting every connection. If the proxy determines that the connection is allowed, it opens a second connection to the server from itself on behalf of the original host, as shown in Figure 5-8. This sort of functionality is commonly seen when users surf the Internet; their computers talk to a proxy that, in turn, talks to servers on the Internet on their behalf.




Figure 5-8 Placement and Packet Flow of a Proxy

The data portion of each packet must be stripped off, examined, rebuilt, and sent again on the second connection. As shown in the list and in the following sections, different types of firewalls can be used to accomplish this:

- **Standard proxy firewalls:** A proxy firewall does not route packets; it simply forwards them, and it operates at the application layer of the TCP/IP model. Functionally, a proxy firewall receives packets from one interface, inspects the packets according to the defined rule set (perhaps access to porn is blocked), and passes the packets out to the firewall if the request is permitted (checking the weather). A connection is never made from the outside to the inside by PCs; as far as the PCs inside the firewall know, all their information comes from the proxy firewall.
- **Dynamic proxy firewalls:** Originally developed from the concepts described for standard proxy firewalls, a dynamic proxy firewall was designed to take the benefits of standard proxies and add the benefits of packet filtering. A dynamic proxy firewall performs a complete inspection of the packet; when a connection is first made and, after it is approved, the faster and weaker packet filtering mechanism handles all additional packets. To summarize, connections are first inspected at the application layer and then at the network layer.

Because these proxy firewalls have full visibility into the application layer of the TCP/IP model, they can look for more specific pieces of data than any other type of technology discussed thus far. For example, they can tell the difference between an email and Java data contained within a packet, as shown in Figure 5-9.

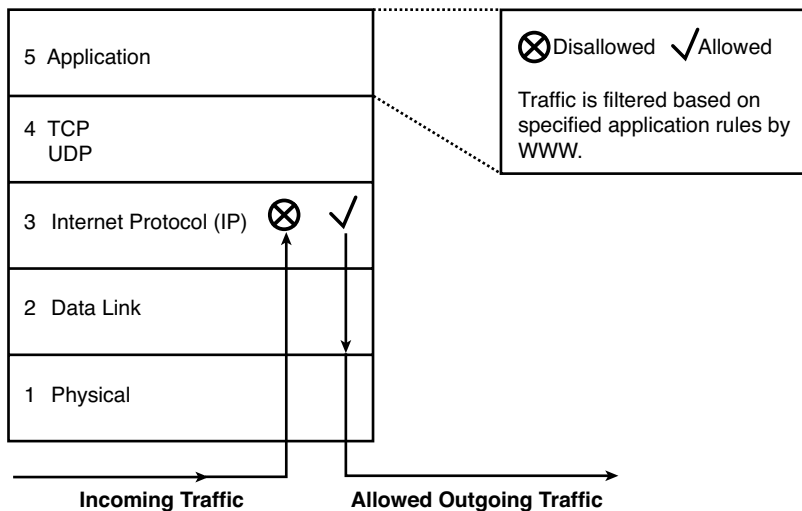


Figure 5-9 Proxy Packet Inspection

As the packet is inspected upon being received by the proxy server in Figure 5-9, all aspects of the TCP/IP header information is removed from the *actual data* and just data is

inspected. The information gathered by this inspection would then be compared against the proxy server rules, and the packet would then either be denied or permitted based on this comparison. If the packet were deemed as something that should be permitted, the proxy firewall stores the connection information from the headers, rewrites the headers, and retransmits the packet accordingly. If the packet were denied, it would be thrown in the bit bucket. Often, the proxy gives users a web page stating why the website they were trying to go to is not allowed; for example, a reference to the Acceptable Usage Policy (AUP).

Note Have you ever heard the phrase *bit bucket*? It is a lighthearted way of saying trash or garbage can. When saying that a packet is thrown in the bit bucket, this actually means that the router, firewall, or proxy has chosen to discard the packet; because all data is ultimately only bits (1s and 0s), this is proof that nerds have a sense of humor.

Limitations of Proxies

Hopefully by now you have realized that implementing any technology and especially security has limitations or drawbacks that you must consider. The folks that sell and market these devices would be thrilled if you believe that their new security gizmo is perfect for solving all your problems. Reality is frequently not the rosy picture they would like you to believe, and proxy firewalls are no different. Following are some of the limitations of proxy firewalls:

- **Reduced performance:** This thorough examination and handling of packets means that proxy firewalls are secure and generally slower than normal processing. Reduced performance could result because of the inspection of essentially every part of every packet being subjected to this level of security.
- **Not always current:** As new protocols and applications are developed, proxy servers must be expanded to recognize what is acceptable. This expansion means that, to stay current, new software must be developed and tested; this takes time and results in a security device that might not always be current.

From a security standpoint, the most secure firewall is a standard proxy firewall that inspects all traffic on an application layer. However, that is not always the most practical solution in many of today's networks. Careful planning and understanding of the required network security and the traffic therein is important for developing a strong security solution. For example, a landscaping company has different security needs than a company that builds electronic components for the military. You should be aware of your application-level traffic through baselines and apply only the necessary security controls applicable to your baselined traffic until things evolve.

Of the two types of firewalls discussed—stateful and proxy—it is crucial that you use at least one of them as part of your layered approach to network security and defense in depth. Add to them the presence of packet filtering on your edge router and a firewall device that also uses NAT, and you will have developed the beginning of a layered

defense. The following section examines how you can also use content filters to protect your network and its users.

Content Filters

Content filtering is a subject so vast that its implications and possible solutions have spawned entire businesses dedicated to providing the right solution for you, regardless of whether you are a home user or a large business. Everyone seems to be faced with the need to filter some sort of content at every aspect of how they connect. Consider some of the challenges that have recently emerged in politics and the media:

- **Public libraries and pornography:** For some reason, there is a group of people who think people have the right to surf pornography on computers that tax dollars pay for. Making this issue worse is that they do this in the middle of libraries—the same place where children go to read. Content filters could be used in libraries to disallow access to this type of content. Businesses are also using content filters to filter out user attempts at going to sites on the Internet.

Unfortunately, the problem is not only about pornographic websites—there are also those sites dedicated to drug use, criminal activity, terrorism, violence, threats to the safety of children, and hate speech.

- **Spam:** If you have email, you have spam—of that there can be no doubt. All types of businesses are fighting back against spam, and it has always been a fight to detect and stop spam. Every time a solution is discovered, spammers get more creative and do something different. For example, many people spell out their email addresses now—*tom dot thomas at netcerts dot com*—in hopes of fooling the programs that search for email addresses. It might for a little while, but it will not last long. In the arena of spam prevention, content filters can identify those annoying ads for low mortgage rates. They are so silly; who would want to get a mortgage with a company that had to spam to get your business? Trust me, you have not lost money in Nigeria either that was found by some mysterious individual who is emailing you; if any of these things were true, you would not be contacted via email. But you knew this; if only the gullible people who didn't would buy this book!
- **Viruses and Trojan horses (malicious code):** Many of the ways viruses are spread follow the growth patterns of the Internet. Virtually everyone who connects to the Internet has email—thus sending a malicious attachment in an email has become commonplace. Content filters would examine the content of such attachments and filter them before any damage was done.
- **Malicious web pages:** Attackers can now code into web pages ways to learn more about you when you visit those pages, and they can do this in many ways. Content filters would examine the actual HTML code that makes the website and filter it as needed. This happens more frequently than you imagine; users didn't do anything or go someplace they shouldn't. A normal website can be hacked with bogus content in place with the end results that every visitor gets infected.

- **Increased organization success:** You might wonder how content filtering can increase a corporation's overall success. Companies and government agencies can face significant risk because of their employees' behavior. Consider the implications to any organization if an employee were to access offensive or illegal material via that organization's network. For example, employees visiting websites with offensive content can create a hostile work environment and negatively affect morale or productivity, which might lead to potentially costly legal fees with the resulting negative bad press. Do you recall the concept of downstream liability discussed in Chapter 1, "There Be Hackers Here." If an employee were to access child pornography, the organization could be held liable, have assets seized (network), and suffer additional negative publicity.

Internet access has become critical to businesses, and the rewards to many organizations can be high. However, issues arise where employees have unmanaged access to the Internet, as just discussed. None of the technologies discussed thus far address the potential security risks just listed. You might be correctly thinking that not all these risks are applicable to your organization, and that might be true. The goal of this chapter is to discuss the technology surrounding content filtering, which could clearly be applied to many different problems, depending on your need. Benefits of content filtering include the following:

- Reduce the legal liability by not letting your organization's resources be used in a compromising manner or through the inadvertent disclosure of confidential information.
- Optimize employee productivity; who wants to pay people's salaries while they are surfing the Internet for pleasure?
- Improve reporting on employee Internet usage. This is critical because you might feel protected or safe. There is no way to know for sure unless you also watch what happens on your network.
- Enforce company Internet access policies that would be documented in the Acceptable Use Security Policy, as discussed in Chapter 2:
 - Disallow the accessing of illegal or offensive material.
 - Prevent the downloading of unauthorized software.
 - Sorry, no holiday shopping during work hours.

You can filter the content of packets in a variety of ways as they flow through your network. Entire companies and many products provide any type of filtering service for you from spam to content. To do them justice by explaining them all is beyond the scope of this chapter. There are some common fundamental similarities, regardless of the product selected.

Note Your organization's Acceptable Use Policy should inform employees about what is expected from them as users of corporate resources, and the content monitoring or filtering monitors and reports on compliance.

The key to content filtering solutions is the ability to monitor and filter content from the Internet, chat rooms, instant messaging, email, email attachments, Word, PowerPoint, PDFs, and from web browsers. There are several ways to filter traffic, which can be classified into two main categories:

- **Client-based filtering:** This filtering solution involves loading software onto individual PCs that check content and filter it according to a defined set of rules. In the case of home users, this is the most common type of solution and usually comes in the form of a subscription to a server that contains updates.
- **Server-based filtering:** In this filtering solution, individual client PCs do not require specialized software to be loaded because everything is loaded and controlled by a server that the client PCs in turn access. This type of filtering is commonly used for email spam and virus detection; all email comes into a central server, which is the most logical place to filter it.

For content filtering, a device such as a proxy server, content engine, or WAN optimization device forces all web traffic through it so that the user requests to view web pages. Users can be inspected to determine whether the request should be permitted or denied. Content filtering is accomplished using a library or database of terminology, words, and phrases as the set of rules defining what is not allowed.

In many cases, requests are regarded as the replies; for example, some attempts to access a website might be classified via the database or library when the client makes a request (such as www.showmeporno.com), whereas other requests might require the filtering device to analyze the content of the web page before making a filtering decision.

These same examples of browsing the Internet using content filtering is extremely similar to how spam and virus filtering is accomplished. Ultimately, a database contains ways of identifying what should be filtered and what should not. As traffic enters the network, it is verified against this database. For example, many products and tools can be used at the server level to identify and stop spam. Although nothing is ever 100 percent accurate; so many email clients also have some sort of built-in way of allowing users to further identify spam email.

Limitations of Content Filtering

Content filtering can play a large role in protecting your network and ensuring the proper use of network resources. However, it does have some disadvantages that, if you are aware of them, allow for the filtering to operate better:

- An estimated 3 to 5 million websites are introduced to the Internet as new or re-named every week. This makes the tracking of good or bad sites extremely difficult to do and requires dedicated service to ensure that your filters are always up to date.
- Content is always changing; in addition to new websites, new ways to spam, new viruses, and other threats make it difficult to keep on top of the changes.
- Nothing is perfect, so you can expect to see false positives to a certain degree. Therefore, retaining some sort of control of the system is important, and blind reliance on outside classifications is probably not a good idea—for example, www.msexchange.com being seen by content or URL filters as “m sexchange” rather than “ms exchange.”
- In the higher education environment, a balance between security and freedom of academia is often a balance that must be struck. RIAA also comes into play here from a compliance-related perspective on downloads and sharing protected music through open programs riddled with security threats.

Content filtering is probably in use in your network in some form or another. The extent of its implementation varies widely depending on the size and sensitivity of your business. The following section looks at ways to completely secure your network: PKI.

Public Key Infrastructure

Have you ever bought anything online or otherwise engaged in some sort of electronic commerce on the Internet? Most likely, you saw the little lock in the corner of your browser window that told you that this was a secure transaction. With what you have learned so far in this book, do you honestly believe that?

The little key or lock in your browser means that you are on a website (server) that uses a *Secure Socket Layer (SSL)* certificate, so you can rest assured that they are who they say they are. Go ahead—buy and enter your credit card number!

Note The little lock means that an SSL connection has been engaged. Anyone can cause a secure connection to take place, so be careful even when you see a little lock.

Have you ever noticed that, while you are conducting e-commerce, the `http://...` changes to `https://...`? The presence of the “s” means that you are using HTTP over SSL to communicate back and forth.

Ultimately, what is actually occurring is that your web browser is taking in the SSL certificate, contacting whoever certified it to ensure its validity, and then proceeding to communicate in a secure mode with the server so that you can complete your transaction in complete security. Do you still believe that this is a good system?

Did I mention that this SSL certificate session is 40 bits in length? Certain aspects of the certificate that reside on the server are 1024 bits. Compare this 40-bit length to an IP address, which is 32 bits in length or 3DES encryption at 128 bits. You should never feel 100 percent secure when conducting e-commerce at this stage in the Internet's evolution because the security is not there yet. As the use of e-commerce continues to rise, the level of fraud is increasing even more. This includes forging certificates that may use valid certificates from the "lock" perspective that encourages man-in-the-middle attacks. This trend is taking a toll on the growth and confidence in e-commerce and online transactions of all kinds. Of course, none of this is ever talked about in polite sales and marketing circles. Not to fret—an advance in securing e-commerce is coming in the form of PKI.

Public Key Infrastructure (PKI) is an evolving technology that will eventually become standard. The goal of PKI is to provide a foundation for a system that supports a variety of security services, such as data integrity, data confidentiality, and nonrepudiation; thus preventing destruction, alteration, and disclosure. PKI can provide this through a combination of hardware, software, procedures, and policies so that users can communicate and exchange information securely, regardless of location.

This system involves the verification and authentication of each side of a transaction over a network. Consider for a moment the impact that online credit-card fraud has on people and businesses. At this time, everyone is losing when fraud occurs—the people because they had their credit card or identity stolen, and the businesses because they are trying to provide a service while remaining profitable.

PKI provides for authentication through the use of advanced digital certificates and certification authorities and subordinate certification authorities to verify and authenticate the validity of each side of a transaction. This transaction could be something as sensitive as an online Internet purchase or as straightforward as exchanging sensitive information via email. PKI is going to be the next step in the evolution and enablement of secure communication and e-commerce.

You can find additional PKI resources online at the following locations:

www.pki-page.org/

www.pkiforum.org/

PKI's Limitations

In researching PKI, I began to think this was a great next step in security—even more so when my identity was stolen—see, no one is safe or perfect! Of course, I did the right thing and called the police; I was amazed at the lack of concern shown by our law-enforcement agencies. The ease with which people dismissed the crime was amazing, not to mention that businesses felt it was just a risk whose loss they had to absorb. Trust me,

preventing loss is where you should spend your time! Certainly then, PKI would be a good step; however, there are some serious challenges in its future:

- E-commerce is working and flourishing on the Internet, regardless of the occasional risks involved.
- Serious laws in states like Utah and Washington are on the books, saying that if someone were to crack your key or illegally use it, you are still responsible for the debt they created. Having seen the bills created by the theft of my wife's identity, this is extremely worrisome to me if I am ever forced to use PKI!
- Security is today, and it is likely to continue to be under PKI, the responsibility of the certificate holder. Thus, you must trust that they have taken all the necessary precautions without exposing new vulnerabilities. PKI is coming; however, there are still some questions in my mind about it.
- PKI does not support a single login infrastructure (single sign on), so users will need to log in and authenticate multiple times to access different resources; this is a recipe for disaster. Users will find ways to "simplify" (that is, defeat) the security PKI provides, and mistakes will happen.

So, is a technology such as PKI good or bad? That is difficult to say because PKI is not mature enough to be fully vetted. However, PKI does provide for increased security that could help in many areas. The verdict on PKI is still up in the air and is subject to the whims of the PKI vendors and how they listen and evolve their products. Of course, organizations then have to choose to spend money on PKI to correctly implement it; PKI's adoption will take some time. The following section looks at some methods currently available for authenticating access to the network.

Reputation-Based Security

Internet users are under attack, and an increasingly common characteristic of malware is the presence of a URL that a user must visit as part of the attack. Organized criminals methodically and invisibly exploit vulnerabilities in websites and browsers to infect computers, stealing valuable information (login credentials, credit card numbers, and intellectual property), and turning both corporate and consumer networks into unwilling participants in propagating spam and malware. Simply allowing a user to visit their favorite website, or clicking a link from their top ten search results, is all it takes for the malware infection process to unknowingly begin.

For most malware creators, recognition for creating a clever piece of malware is no longer the point. With a thriving, maturing malware economy in place, it's more valuable to create malicious code that generates revenues for online criminal networks—for example, through click-fraud, massive spam campaigns, or identity and data theft.

To be successful, the malware must be both easy to distribute to as many victims as possible and difficult to detect. That makes the Internet an attractive malware delivery mechanism. Originally, malware was delivered directly through email, but the visibility of large

attachments and the store-and-forward nature of email made it relatively simple to stop. The near-real-time nature of Internet websites, with threats hidden directly in the content, makes malware exponentially more difficult to stop.

The growing significance of the Web as a threat delivery mechanism is shown by the fact that more than 80 percent of spam messages include URLs, which can direct a user to a web server where malware is located. That percentage is even higher for malicious emails, such as phishing campaigns. These URLs are intended to lure readers to websites that engage them in questionable transactions or download malware onto their computers.

Typically, both the spam messages and the malicious websites the messages refer to use a combination of social engineering and software vulnerabilities to compromise users. Malicious websites, specifically created to distribute malware, are not the only sites compromising users. Hackers are now frequently distributing malware through legitimate websites that have been compromised, taking advantage of security flaws in web applications.

More often, malware writers are targeting legitimate, trusted websites as the starting point for malware distribution. Both BusinessWeek.com and MSNBCsports.com had portions of their websites used for distributing malware. Although no threat is present on these websites today, users became infected simply by visiting trusted sites. Knowing these websites are trusted by millions of users makes them easy targets for malware writers.

As Figure 5-10 shows, the attacker's traffic mixes with that of trusted visitors. If the attackers gain control of the site, they often insert attacks to those trusted users.




Figure 5-10 *Criminals Compromise Legitimate Websites to Infect Unsuspecting Users*

Reactive Filtering Can't Keep Up

Traditional methods of protection are usually not fast, accurate, or comprehensive enough to assess and protect users from these new, dynamic web-based threats, which are growing in record numbers.

IP blacklists and URL-filtering solutions typically cover only a small percentage of all URLs and IP addresses—and only the known bad ones. They are also normally binary, offering only “block/malicious” or “allow/safe” options for the URLs and IP addresses they do cover, instead of providing detailed, granular information about any possibly suspicious URL, IP address, or object—even those that haven't been known offenders before.

Even with security categories enabled, these URL-filtering solutions can't help when a legitimate, normally trustworthy website has been turned into a redirection hub for malware distribution. The website's URL is trusted and not on any blacklist. Consequently, acceptable-use policies designed to protect a network by preventing access to certain sites can't prevent users from getting infected on acceptable websites. Because traditional URL-filtering technologies are concerned only with the initial domain request, they don't examine the additional objects needed to load the web page correctly or their origins, and thus don't observe the malicious redirection. When a web page has an average of 150 objects, traditional URL-filtering technologies simply can't keep up. This was the case on September 13, 2009, for visitors to NYTimes.com; a trusted source often categorized by URL filtering lists as “news.” A seemingly legitimate advertisement (inserted via a single object on the site—when there are so many objects linked to each web page) began presenting a pop-up, alerting visitors that a virus had infected their system. Victims were then redirected to a malware site that offered legitimate-looking antivirus software, which was actually a malicious Trojan.

The sophistication, innovation, rapid pace, and dynamic nature of these attacks often render traditional defenses useless. URL filtering and IP blacklisting are reactive and cannot adequately assess new or previously uncompromised sites in a timely fashion, whereas signature-based scanning solutions have trouble keeping up with the constant mutation of malware. Protecting users from today's web-based threats requires a layered, holistic, and integrated approach that uses multiple advanced methodologies to assess each threat and type of network traffic. The solution to this new threat asks a simple but powerful question:

“What is the reputation of this URL?”

When assessing the trustworthiness of a URL, a great deal can be determined by analyzing data that is hard to forge, such as how long the domain has been registered, in what country the website is hosted, whether the domain is owned by a Fortune 500 company, whether the web server is using a dynamic IP address, and more.

For example phishing site creators can spoof the content of their websites to perfectly replicate legitimate banking and e-commerce sites. Phishing sites cannot, however, spoof the URL on which they are located. The reputation of the URL assigns a reliability score to the vast majority of URLs and can therefore protect users. Analyzing data, even the

most difficult to manipulate elements, can reveal much about the trustworthiness of a URL. Data analysis can determine how long a domain has been registered, whether it was registered by machine or manually, who owns it, whether it is associated with an IP address that has previously been associated with a web-based threat, whether the IP address is dynamic or static, what country the website is hosted in, and more. By gathering this information and assigning a score to each category when a user attempts to access a URL, this score is calculated, and access is either permitted or denied. As shown in Figure 5-11, sophisticated algorithms analyze and correlate threats with more than 200 different web traffic- and network-related parameters (*Cisco products only*) to accurately evaluate a web object's malware risk. Using this data, a dynamic score ranging from +10 to -10 is generated for web reputation. The same technology in senderbase.org, now sensorbase, has been adapted to intrusion prevention system (IPS) technologies, but the scoring assumes the attack is malicious and is -1 to -10 as an additional anomaly detection over and above traditional methods.

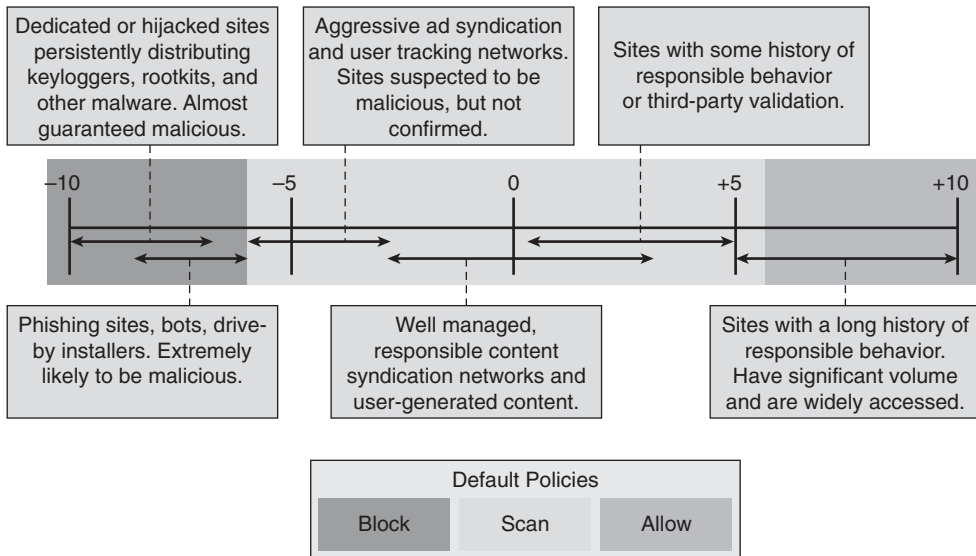


Figure 5-11 URL Reputation Examples

Cisco Web Reputation Solution

Cisco Web Reputation Filters are the world's premier reputation system, in part because of the Cisco acquisition of Ironport. Powered by the Cisco Security Intelligence Operations and the Sensor Base network, Cisco Web Reputation Filters have visibility into more than 100,000 global networks—including Cisco IPS, with more than 30 percent of the world's email and real-time traffic insights from customer participation. Cisco built the Sensor Base reputation database from more than 800,000 sensors deployed by customers globally. Each sensor now has the capability to anonymously contribute what it is detecting directly to Cisco Sensor Base.

Unlike traditional URL-filtering solutions, Cisco Web Reputation Filters examine every request made by the browser. Instead of just looking at the initial HTML request, they also analyze all subsequent data requests, considering each element on a web page and its origins—including live data (such as JavaScript, ads, and widgets), which might be fed from different domains. This enables Cisco Web Reputation Filters to give users a much more precise and accurate assessment and block web content in a far more fine-grained way than URL-filtering and IP-blacklisting solutions.

AAA Technologies

Today, we live in a world in which almost everything must be protected from misuse and nothing is free. It does not matter whether you are a system administrator, manager, student, or a network engineer. If you access services via a network, you always need three things:

- Authentication
- Authorization
- Accounting

These components are collectively known as AAA (Commonly referred to as *Triple A*). As discussed in the following sections, each of these components plays an important role.

Authentication

Authentication ensures that the network's users are who they claim to be. This is important because you do not want these people accessing the network if they are not supposed to. Usually a shared secret or a trusted third-party software application provides authentication.

Authentication enables the network administrators to identify who can connect to a network device or Internet by including the user's username and password. Normally, when a user connects to a router remotely via Telnet, the user must supply only a password to gain access to the router. This is functional but not secure because, if the router is connected to the Internet, an attacker could try and try to connect, and you might never know that this was occurring. All the attacker would need to do is guess a single password to access your router. How hard could that be when he has all the time in the world?

When someone logs on to one of your network devices and makes a change, how do you know who the person is and what she has done? With AAA authentication, whenever a user logs on, the user must enter a username and password pair (which the network administrator assigned). The following code snippet shows an example of a remote user accessing a Cisco router with AAA configured to request a username:

```
User Access Verification
Username: tom_thomas
Password: xxxxxxxx
MyNetworkDevice>
```

As shown in the preceding example, the user must enter a valid username and password to gain access to the router. Typically, a database that contains the valid usernames resides locally on the device or on a remote security server such as Cisco Access Control Server (ACS).

Authorization

After the user is authenticated, there must be a way to ensure that the user is authorized to do the things he requests. For example, if you are a normal user, you do not have the permissions to access all the files in a file system.

Authorization enables administrators to control the level of access users have after they successfully gain access to the router. Cisco IOS Software enables certain access levels (called *privilege levels*) that control which IOS commands the user can issue. For example, a user with a 0 privilege level cannot issue any IOS commands. A user with a privilege level of 15 can perform all valid IOS commands. The local or remote security server can grant access levels.

You can display your privileged level on a Cisco router with the **show privilege** command, as shown in the following command line:

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

Authorization can also dictate the types of protocol activity in which the user can engage, such as allowing a user to invoke only FTP, Telnet, SSH, or HTTP traffic. The higher the privilege, the more capabilities a user has with the IOS command set.

Accounting

Accounting occurs after the authentication and authorization steps have been completed. Accounting enables administrators to collect information about users and the actions that they take when connected to network devices. The information gathered through accounting can provide network forensic evidence of tampering or hacking because you have a road map of the user's times/dates and activities. Specifically, administrators can track which user logged in to which router or switch, which IOS commands a user issued, and how many bytes were transferred during a user's session. For example, accounting enables administrators to monitor the routers that have had their configurations changed. A router or a remote security server can collect accounting information.

Note If you use wireless in an airport, for example, to access the Internet, you use a form of AAA when you authenticate and receive authorization into the service provider's network. Accounting is the process in which the network service provider collects network usage information for billing relating to how long you were connected, capacity planning,

and other purposes. This is important for the service provider—there is no such thing as a free lunch.

After AAA is configured, you can use external security servers to run external security protocols—such as RADIUS or TACACS—that will stop unauthorized access to your network. Both RADIUS and TACACS can be implemented on Cisco network devices and are reviewed in the upcoming sections.

Note You must use AAA if you intend to use RADIUS or TACACS security server protocols. As AAA collects the information, it sends it to the security servers to determine each of the characteristics associated with AAA.

Remote Authentication Dial-In User Service (RADIUS)

RADIUS is a client/server-based system that secures a Cisco network against intruders. RADIUS is a protocol implemented in Cisco IOS Software that sends authentication requests to a RADIUS server. A RADIUS server is a device that has the RADIUS daemon or application installed. RADIUS must be used with AAA to enable the authentication, authorization, and accounting of remote users. When a RADIUS server authenticates a user, the following events occur:

1. The remote user is prompted for a username and password.
2. The username and password are encrypted and sent across the data network.
3. The RADIUS server accepts or rejects a username and password pair. In some instances, a user might be asked to enter more information. (This is called a *challenge response*.) For example, if a user's password has expired, a RADIUS server prompts the user for a new password.

Note Traffic between the Network Access Server (NAS) and RADIUS is not encrypted—as opposed to TACACS, which does encrypt authentication message traffic.

Note A RADIUS server is usually software that runs on various platforms, including Microsoft NT servers or a UNIX host. RADIUS can authenticate router users, authenticate vendors, and even validate IP routes.

The following steps are required to enable RADIUS on a Cisco router:

- Step 1.** Use the `aaa new-model` command. AAA must be used with RADIUS.
- Step 2.** Specify the RADIUS server with the `radius-server host` command, as shown in Example 5-3.

Step 3. Specify the password used between the router and the RADIUS server.

Note Of course, you must also ensure that you have entered users and passwords into the RADIUS server before activating RADIUS.

Example 5-3 displays the required configuration for a Cisco router to authenticate users from the RADIUS server with the host address 10.99.34.50.

Example 5-3 RADIUS Configuration

```
radius-server host 10.99.34.50

radius-server key <password>
```

Let's move on to TACACS, which is an alternative protocol to RADIUS that also works with AAA.

Terminal Access Controller Access Control System (TACACS)

Cisco IOS supports three versions of TACACS: TACACS, extended TACACS, and TACACS+. All three methods authenticate users and deny access to users who do not have a valid username and password.

The first version of TACACS provides simple password verification and authentication. Accounting is limited in that only requests and denials are listed. Next, extended TACACS replaced the first version of TACACS. TACACS+, also referred to as *TACACS plus*, provides detailed accounting and must be used with AAA (in other words, the **aaa new-model** command must be enabled). TACACS+ (yes, the plus sign is important) supersedes the earlier releases of TACACS. In general, TACACS provides a centralized security system that validates users from any remote location. Typically, TACACS runs on a Windows Server or UNIX operating system. When a TACACS server authenticates a user, the following events occur:

1. The remote user is prompted for a username and password.
2. The username and password are sent across the data network and is authenticated.
3. The TACACS server accepts or rejects the username and password pair. The user might be asked to enter additional information (called a challenge response).

For example, a challenge response might appear when an error occurs during authentication. TACACS+ requires AAA, but TACACS and extended TACACS do not use AAA.

The configuration tasks required to enable TACACS+ on a Cisco router are as follows:

- Step 1.** Use the **aaa new-model** command. AAA must be used with TACACS+.
- Step 2.** Specify the TACACS+ server with the **tacacs-server host** command.

- Step 3.** Specify the authentication key used between the router and the TACACS+ server.
- Step 4.** Because TACACS+ must be used with AAA, you must specify TACACS+ authentication, authorization, and accounting.

Example 5-4 displays the required configuration for a Cisco router to authenticate users from the TACACS+ server with the host address 10.99.34.50.

Example 5-4 TACACS Configuration

```

aaa new-model

aaa authentication enable default tacacs+

! Sets router to use the tacacs server to authenticate enable
! password

aaa authorization exec tacacs+

! Sets tacacs+ plus to authorize exec commands on local router

aaa accounting exec start-stop tacacs+

! Accounting information is gathered for exec commands

tacacs-server host 10.99.34.50

tacacs-server key <password>

```

Example 5-4 is a basic TACACS + configuration; you can set other configuration options to enable complex AAA commands.

Caution If you enable AAA on a router, you could get locked out if you are not careful. If you fat finger any commands and exit out of your configuration, you might not be able to re-enter; make sure you are certain of your work before disconnecting.

TACACS+ Versus RADIUS

Comparing the two server protocols, RADIUS and TACACS+, shows that both require AAA to be enabled on a Cisco router (unless you use the older versions of TACACS+, namely TACACS and extended TACACS). RADIUS and TACACS+ both require a username and password pair to obtain access. The difference between the two protocols is in the protocol itself and the fact that TACACS+ is a centralized validation service, whereas RADIUS is based on client/server technologies.

Two-Factor Authentication/Multifactor Authentication

As we have shown when reviewing RADIUS and TACACS, they are a means to securely authenticate to access a secure device. Those authentication methods relied on the user knowing a secret codeword or password, typically things the user knows. Although they are efficient, they are single layers in the defense of your network; additional layers, for defense in depth, are provided by using two-factor authentication. Two-factor authentication is when there are two independent and separate methods of authentication that the user must pass to gain access. Usually one of the two methods is something the user has that must be applied as part of the authentication process. You might think that this is a new concept, but you have likely have been engaging in two-factor authentication for quite awhile without realizing it. Accessing an ATM or paying with a bank card is two-factor authentication, the first authentication is “having the card,” and the second authentication is “knowing the PIN.”

Perhaps you have access with a token card that is synced to generate a unique code that when applied together enables you to gain access to your company’s network via a VPN. Two-factor authentication is becoming more and more common these days; your kids are likely using it if they play the online game *World of Warcraft*. As shown in Figure 5-12, you gain access to the game by “knowing” your username/password and “having” your token to generate an authentication code.




Figure 5-12 *Warcraft Authenticator Code Is Two-Factor Authentication*

IEEE 802.1x: Network Access Control (NAC)

Organizations continue to embrace mobility for their users by expanding wireless LANs (WLAN) for PCs and a whole range of mobile devices. Wireless networks are attractive because they are much easier to deploy and use than wired networks and cheaper, making them a better solution in many offices. However, security is a big concern because the open nature of wireless LANs brings a whole slew of concerns about user and corporate data being pulled from the air. Also increasing is the concern for the security of wireless networks with all sorts of new threats emerging, as discussed, and the bleed to wired networks. Organizations require security mechanisms that ensure that when credentials are transmitted, they remain secure and enable organizations to ensure that users trying to connect are whom they claim to be.

Enter 802.1X, a standard for port-based network access control, developed by the Institute of Electrical and Electronics Engineers (IEEE). 802.1X was originally designed for use in wired networks but was adapted to address WLAN security concerns because of its robust, extensible security framework and powerful authentication and data privacy capabilities. The 802.1X standard also defines the encapsulation methodologies for the transport of EAP over PPP or Ethernet. The 802.1X standard delivers powerful authentication, and security enables you to enforce port-based network access control when devices attempt to access the network. The 802.1X standard has three main components:

- **Supplicant:** Software that resides on the user's machine or device and is used to request access to a wired or wireless network
- **Authenticator:** Devices, such as switches or wireless access points, that sit between the supplicant and the authentication server
- **Authentication server:** A server that receives authentication messages which in turn takes the request and validates against a back-end data store such as Active Directory, eDirectory, or LDAP

802.1x has become popular in wireless and wired networks in large part because its operation is secure and straightforward. When attempting to access an 802.1X-enabled network, instead of the user or device simply being granted Layer 3 access, it is challenged for its identity. If the user's device is not configured for use in an 802.1X-based network, that is, it does not have a running supplicant, it will be denied network access. If the user's device is configured with an operational supplicant, it will respond to the challenge for its identity and start the 802.1X authentication process. The supplicant passes network credentials (user and/or device identification information) to the authenticator, which verifies the connection to the network and passes the identification information on to the authentication server to determine access. Figure 5-13 demonstrates how 802.1X would work in a wireless network.

In a wired network, switches can use IEEE 802.1X to perform user authentication, rather than the types of device authentication performed by many of the other features described in this section. User authentication requires the user to supply a username and password, verified by a RADIUS server, before the switch can enable the switch port for

normal user traffic. Requiring a username and password prevents the attacker from simply using someone else's PC to attack the network without first breaking the 802.1X authentication username and password.




Figure 5-13 802.1x Authentication Process Flow

Cisco has a comprehensive identity management solution based on 802.1X called TrustSec. TrustSec is an integrated solution that uses Cisco products that offer authentication, access control, and user policies to secure network connectivity and resources. These products include the following:

- Cisco Catalyst family of switches
- Wireless LAN access points and controllers
- Cisco Secure ACS
- Cisco Secure Services Client

Additional and optional components include X.509 public key infrastructure (PKI) certificate architecture. You can find detailed TrustSec information including configuration and deployment guidelines at www.cisco.com/go/trustsec.

Information on how the TrustSec and 802.1x solution is integrated into Cisco NAC is covered in the following section.

Network Admission Control

Network Admission Control (NAC) is a multipart solution that validates the security posture of an endpoint system before entering the network. With NAC, you can also define what resources the endpoint has access to, based on the results of its security posture. NAC is a key part of the Cisco Self-Defending Network Initiative (SDNI). The SDNI mission is to dramatically improve the capability of the network to identify, prevent, and adapt to threats.

NAC Appliance or Cisco Clean Access (CCA) enables an organization to enforce security policies by blocking, quarantining, and performing remediation of noncompliant systems. Remediation occurs at the discretion of the administrator. The policies and requirements enforced by the Cisco NAC Appliance include checks for latest antivirus software, operating system (OS) patches, and security patches. The Cisco NAC Appliance can also perform vulnerability scanning on the end-user machine in addition to role-based authentication on users attempting to connect to the network. The NAC Appliance solution can

restrict what resources these users can access, based on their role. All these policies and configurations are done in the Clean Access Manager (CAM). The Cisco NAC Appliance has three major components:

- **Clean Access Server (CAS):** Acts as a network control device
- **Clean Access Manager (CAM):** Manages one or more servers
- **Clean Access Agent (optional):** Serves as an endpoint lightweight client for device-based registry scans in unmanaged environments

Cisco TrustSec

The traditional network and physical perimeter is no longer the only border where information must be defended. Collaboration, IT consumerization, mobility, and new computing technologies are increasing productivity while presenting new security requirements. There is greater pressure on IT to meet the demands of a dynamic workforce—both in terms of service delivery and security challenges. New solutions are needed to protect borderless networks and to help further improve business efficiencies in the mean time. Cisco TrustSec is such a solution.

Solution Overview

Cisco TrustSec enables organizations to secure their networks and services through identity-based access control to anyone, anywhere, anytime. The solution also offers data integrity and confidentiality services, policy-based governance, and centralized monitoring, troubleshooting, and reporting services. TrustSec can be combined with personalized, professional service offerings to simplify solution deployment and management, and is a foundational security component to Cisco Borderless Networks.

The Cisco TrustSec solution offers the following benefits to customers:

- **Compliance support:** Expands real-time access visibility and audit trails across an increasingly complex network to address mandated monitoring, auditing, and reporting requirements
- **Strengthened security:** Extends security across the borderless network by enforcing consistent security policy, ensuring endpoint health, and delivering a secure network fabric
- **Increased efficiency:** Reduces IT overhead through centralized identity services, integrated policy enforcement, a consistent user experience, and dynamic assignment of user and device access

The core Cisco TrustSec functional areas follow:

- **Identity-aware user and device access:** Dynamically provides role-based access. Noncompliant devices can be quarantined, remediated, or denied access.

- **Guest user access and lifecycle management:** Sponsored guests receive restricted access to specific resources (Internet, printers, and so on) through a customized web portal. Internal network access is blocked and activity is tracked and reported.
- **Nonuser device discovery:** Nonuser devices (printers, cameras, phones, and so on) are centrally discovered. Access is provided based on policy, and device behavior is monitored and audited to prevent spoofing.
- **Data integrity and confidentiality:** Data paths can be encrypted via MACsec, from the endpoint client to the network core, while allowing critical tools (firewalls, IPSs, content inspection, QoS, and so on) to retain visibility into data streams.
- **Monitoring, management, and troubleshooting:** Centralized, policy-based corporate governance and compliance includes centralized monitoring and tracking of users and devices to maintain policy compliance. Provides sophisticated troubleshooting, detailed auditing, and historical and real-time reporting.
- **Professional services:** TrustSec services provide policy review, analysis, and design expertise to prepare a network to deploy a TrustSec solution.

Figure 5-14 illustrates the mechanics of how Cisco TrustSec works.

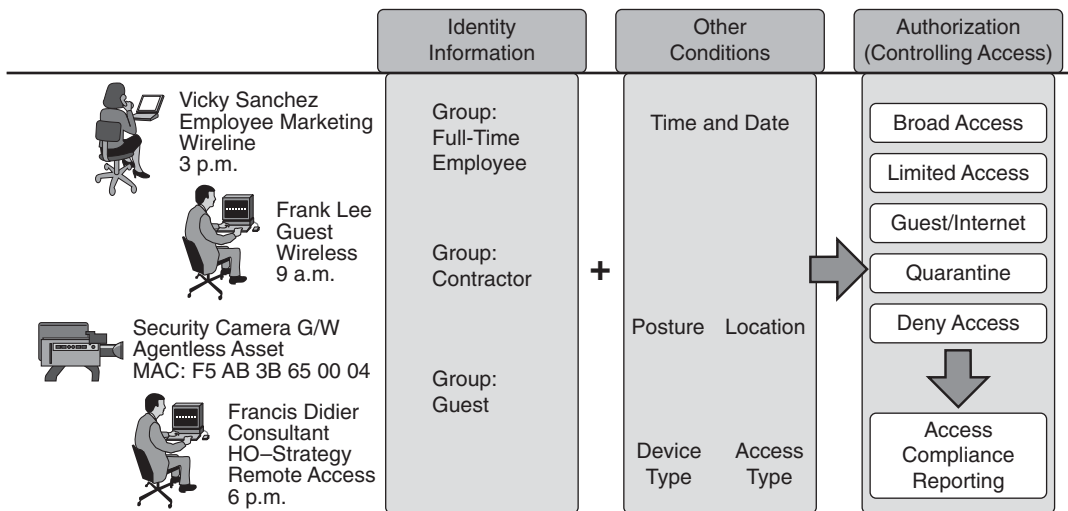


Figure 5-14 How Cisco TrustSec Works

Network users are authenticated with flexible authentication mechanisms to support different device types, operating systems, and access methods.

Cisco Identity Services Engine

Traditional corporate network boundaries and siloed services are a thing of the past. Today's networks must accommodate an ever-growing array of consumer IT devices while providing user-centric policy and enabling global collaboration. The Cisco TrustSec architecture addresses this shift by using identity-based access policies to tell you who and what is connecting to your network, allowing IT to enable appropriate services without sacrificing control.

The first release of ISE focuses on the pervasive service enablement of TrustSec for Borderless Networks. Cisco Identity Services Engine (ISE) delivers all the necessary services required by enterprise networks (AAA, profiling, posture, and guest management) in a single appliance platform. In the future, the same ISE platform can be used to propagate consistent service policies throughout the borderless network, from any endpoint to the video delivery optimization, branch service personalization, and data center server and service agility.

As part of the Cisco TrustSec solution and the Cisco SecureX architecture for Borderless Networks, the Cisco ISE provides a centralized policy engine for business-relevant policy definition and enforcement. ISE complements global contextual information offered by Cisco Security Intelligence Operations (SIO) with localized context awareness for effective access policy enforcement.

- **Security:** Secures your network by providing real-time visibility into and control over all users and devices on your network.
- **Compliance:** Enables effective corporate governance by creating consistent policy across an infrastructure.
- **Efficiency:** Helps increase IT and network staff productivity by automating traditionally labor-intensive tasks and streamlining service delivery.
- **Business-relevant policies:** Enables centralized, coordinated policy creation and consistent policy enforcement across the entire corporate infrastructure, from head office to branch office.
- **Systemwide operational visibility:** Discovers, assesses, and monitors users and endpoints and employs advanced troubleshooting capabilities to give IT teams complete visibility into who and what is on the corporate network.
- **Context-aware enforcement:** Gathers information from users, devices, infrastructure, and network services to enable organizations to enforce contextual-based business policies across the network. Cisco ISE acts as the “single source of truth” for contextually rich identity attributes, including connection status, user and device identity, location, time, and endpoint health.
- **Flexible services architecture:** Combines AAA, posture, profiling, and guest management capabilities into a single appliance platform. Cisco ISE can be deployed across the enterprise infrastructure, applying the appropriate services supporting 802.1x wired, wireless, and VPN networks. Figure 5-15 demonstrates these aspects of ISE.

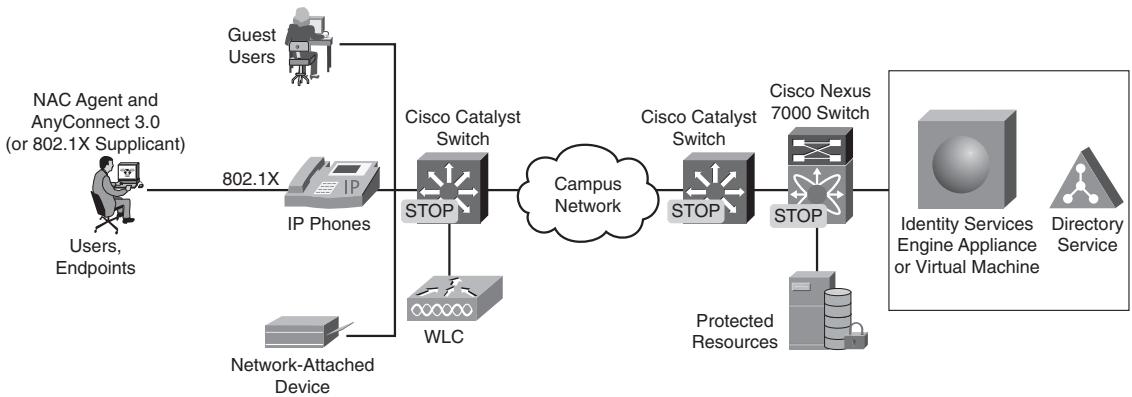


Figure 5-15 ISE-Based TrustSec LAN Deployment

The Cisco ISE is part of an infrastructure-based Cisco TrustSec deployment using Cisco network devices to extend access enforcement throughout a network. Additional deployment components include Cisco NAC Agent and Cisco AnyConnect (or a 802.1x supplicant) on the endpoint; Cisco Catalyst switches and Cisco wireless LAN controllers acting as policy enforcement points for the LAN; and Cisco Adaptive Security Appliances for secure remote access. Cisco ISE also integrates with directory services such as Microsoft Active Directory and Sun ONE Directory Server as policy information points.

Putting Cisco TrustSec and ISE together is a layered solution, as shown in Figure 5-16.

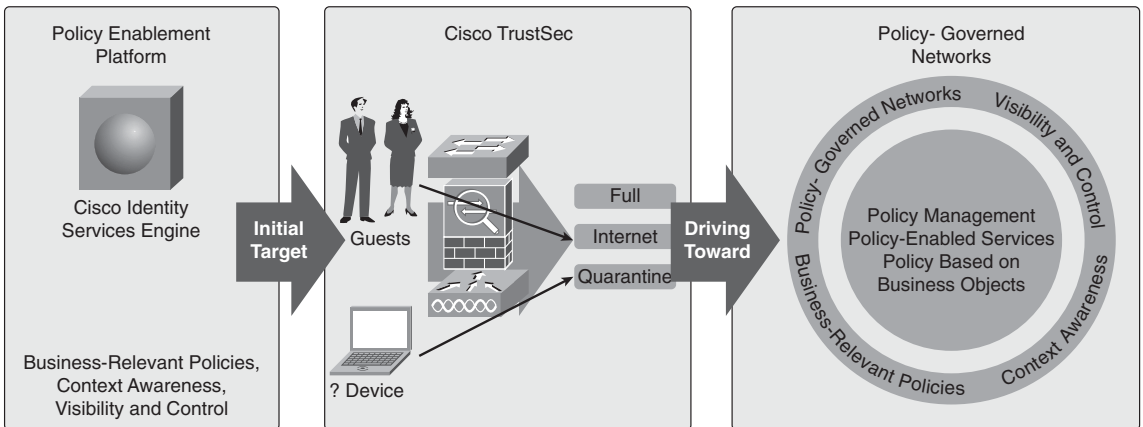


Figure 5-16 ISE and TrustSec

Chapter Summary

This chapter began with a discussion of the importance of a layered network security design. This layering of security provides a deeper level of protection for your network. You must avoid what I call “the orange syndrome,” as in the fruit, in which only a single layer of protection exists before you get to the good stuff. You do not want attackers to defeat a single security layer and get to the good stuff in your network.

This chapter looked at many technologies that you can use to provide a layered approach to security:

- Packet filtering via ACLs
- Stateful packet inspection
- Network Address Translation
- Proxies and application level protection
- Content filters
- Public key infrastructure
- AAA technologies

Separately, each of these technologies is just a single layer of protection, but combined, they provide you with several layers of protection and keep the good stuff safe.

Chapter Review Questions

The following questions assist in reinforcing the concepts covered in this chapter.

1. What are the six security design concepts you should consider when looking at the security technologies for securing your network?
2. What rule is always implicitly present at the end of every packet filter?
3. When a device performs stateful packet inspection, what characteristics in a packet’s header are inspected, and why are they important?
4. What are some limitations of stateful packet inspection?
5. Define the differences between public and private IP addresses.
6. Compare and contrast the three different versions of NAT, and identify which of them is the most commonly used.
7. What are the two types of proxy firewalls?
8. Why is content filtering so important to networking?
9. What is the potential value of PKI to securing a network and e-commerce?
10. AAA provides security for what aspect of a network?
11. Search the Internet and find three potential vendors that can offer an effective RADIUS solution. Describe what features about each are beneficial.

Security Protocols

“...The wisest mind [always] has something yet to learn.”—Author Unknown

By the end of this chapter, you should know and be able to explain the following:

- The difference between DES and 3DES encryption, including their limitations
- AES encryption and its strengths
- The function and role the MD5 hash plays in securing connections
- What a message digest is and how an SHA hash functions
- The differences between PPTP and L2TP
- The breadth and scope of SSH and how it is more secure than Telnet

Answering these key questions will enable you to better understand the overall characteristics and importance of network security. By the time you finish this book, you should have a solid appreciation for network security, its issues, how it works, and why it is important.

Some of you might be wondering why this chapter is called “Security Protocols” because in the IT realm, the term *protocol* is usually reserved for routing, or routed, protocols of some sort. The best routing protocol is Open Shortest Path First (OSPF), and you should learn more about it when you can. At this time, however, the discussion focuses on security. According to Newton’s *Telecom Dictionary*, a protocol is defined as “a set of rules governing the format of messages that are exchanged between computers and people.” I have also seen it defined like this: “A sequence of operations that ensure protection of data. Used with a communications protocol, it provides secure delivery of data between two parties.”

In the realm of security, a *security protocol* is defined as a secure procedure for regulating data transmission between computers. This chapter concerns the methods of securely encrypting data for transmission over a network. Chapter 9, “IPsec Virtual Private Networks (VPNs),” covers the means of transporting data securely.

This chapter enables to develop an understanding of how you can secure data. In many ways, being able to protect data through encryption is yet another layer of a network's security.

Consider that each day, information is being disclosed to people whom you do not want to have it; more often than not, this is sensitive information. In many cases, this is not intentional, nor is it related to criminal activity or attackers in any way. Do you find this difficult to believe? You should not. Think about the following points:

- Sensitive data is placed on servers connected to your LAN for other people to access.
- Sensitive data is copied to USB flash drives, CDs, and DVDs, or printed and then handed to the (in many cases unauthorized) recipient.
- Sensitive data is emailed across the network, or perhaps the Internet, often unencrypted.
- Sensitive data is transmitted in some other manner.
- Sensitive data is placed on a web server and then often removed or altered.

Certainly, these common examples of “business as usual” and “how we do business” are easily recognizable scenarios to many people. We have all done this at some time or another. The danger here is that the sensitive data is being sent *in the clear*; this means that anyone can read the data if they intercept it intentionally or accidentally, or even unintentionally. (Have you ever sent an email to the wrong person?) You might ask yourself what possible kind of data could be used in a negative manner. Consider the following types of data:

- **Personally identifiable information:** Have you ever entered your full name, address, phone number, date of birth, driver's license number, vehicle registration plate number...dare I say Social Security number into a web page or an email?
- **Financial data:** Do you use Quicken or other money-management software on your computer? Is that computer ever connected to a network? What about checking bank account information online, tracking stocks you own, or entering credit-card data online?
- **Customer data:** Does your company enter customer information into a database or take orders online?
- **Medical data:** When was the last time you walked into a hospital or doctor's office and did not see a computer, no less in a common area? Last time I was there, the doctor had a Palm Pilot with all my data loaded onto it. What would happen if he lost it or it was stolen?

These are the most commonly known types of data, but what about movies, music, new product plans, future projections, source code, and so on?

Most of the time, there is no danger of any sort; however, this is not always the case. When there is a mistake, it can be extremely serious. The point here is that everyone and

every company has important data that they would not want shared. This chapter discusses ways to protect this data.

Note When discussing encryption, the password is often referred to as the key; these two terms can be and are used interchangeably.

Triple DES Encryption

The predecessor to Triple DES was DES, which was a fantastic answer to a problem in the 1970s; however, what the developers did not expect, or anticipate, was how much the world would change in less than 30 years. They did not understand that they were on the leading edge of the IT revolution. Ultimately, however, technology has made the protection level of DES such that it left businesses needing another solution.

The DES algorithm became obsolete after it was cracked. To fill the gap, Triple DES (written as *3DES*) was developed from the original DES algorithm. The development of 3DES happened quickly because it was based on the existing DES algorithm.

Looking at the names of the two different algorithms, you might be inclined to believe that 3DES makes your encryption three times more difficult to break. 3DES actually makes your encryption five billion, trillion, trillion times harder to break—that is, 5×10^{33} .

The 3DES algorithm uses three separate keys when running its encryption algorithm and associated computations. Through the use of three 64-bit keys, the key length has effectively been increased from 8 to 24 characters, thereby resulting in 192 bits worth of encryption strength. Mathematically, this means that the number of possible key combinations can be expressed as

$$2^{168} = 3.7 \times 10^{50} \text{ (370 trillion trillion trillion trillion) different combinations}$$

Earlier in this chapter, I mentioned what would happen if you could crack keys at the rate of 1 million per minute. I have no idea how long it would take using 3DES, but I will be long gone from this earth by the time you finish. This is why 3DES is considered strong. You can read more about cracking 3DES in financial ATM applications in the article, “Extracting a 3DES Key from an IBM 4758,” which you can find online at <http://public.planetmirror.com/pub/descrack/>.

Encryption Strength

3DES is an extension of DES that takes three keys and encrypts the data, as shown in Figure 6-1.

The overall procedure to encrypt data is the same in 3DES and DES; however, in 3DES, the encryption process is repeated three times. The plain text data, such as an MS Word document, is encrypted with the first key. The result is then encrypted with the second key, and that result is then encrypted with the third key—hence the name 3DES.

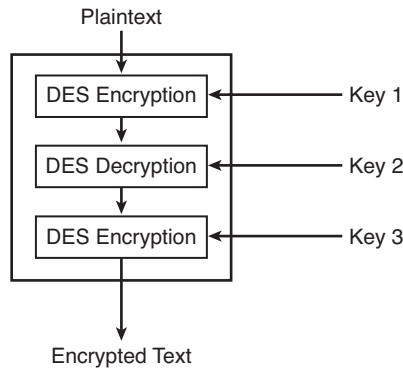


Figure 6-1 *Triple DES Encryption Steps*

Note DES, the block cipher from which 3DES is derived, is now considered to be insecure for many applications. This is primarily because the key size is inadequate; it is only a 65-bit key size. Furthermore, DES has been withdrawn as a standard by NIST, the National Institute of Standards and Technology.

Limitations of 3DES

The resulting actions of having to encrypt every piece of plain text data three times means that 3DES runs slower than normal DES. If used properly with three different keys, 3DES is several magnitudes stronger than DES.

You want to avoid having the same key for each of the three encryption steps. If any of the keys are the same, the end result is that you are using a slower version of DES. As discussed in this section, 3DES is a stronger method of encryption than DES and is used today in many places.

Advanced Encryption Standard (AES)

Federal Information Processing Standards Publication 197 (FIPS PUBS), dated November 26, 2001, announced the advanced encryption standard (AES) to the world. AES specifies a FIPS-approved cryptographic algorithm used to protect electronic data. The publication defines it as, "...a symmetric block cipher that can encrypt (encipher) and decrypt (decipher) information...." The U.S. government adopted this standard, and in June 2003, the U.S. government (NSA) announced that AES was secure enough to protect classified information up to the TOP SECRET level.

Different Encryption Strengths

The AES standard uses one of three block ciphers, AES-128, AES-192, and AES-256, that were adopted from a larger collection originally published as Rijndael.

Each encryption key size causes the algorithm to behave slightly differently, so the increasing key sizes not only offer a larger number of bits with which you can scramble the data but also increase the complexity of the cipher algorithm, forcing the number of rounds to increase from 10, 12, and 14, respectively, required to open the virtual vault you have encompassing your data.

Limitations of AES

Limitations of AES? That is hard to say; when this standard was first introduced it was declared completely unbreakable. What we can say is that AES is used to encrypt everything from the U.S. government's most secret documents to financial transactions from banks and e-commerce sites around the globe. A tear in the AES fabric would open up valuable personal and business information to hackers and foreign governments alike. It's only a matter of time before someone looks for the missing scale on this dragon.

That is just what happened in the spring 2009, when Biryukov, Khovratovich, and Nokolic found a key recovery attack on AES-256 with a time complexity of 2^{131} . This enterprise was completely impractical, but it marked the first time anyone had published an attack on the full AES cipher. Shortly after that, the time was reduced to 2^{119} and the first attack on AES-192 was attempted and succeeded. As a result, AES is no longer considered theoretically secure.

Is AES broken? No. The latest attack techniques on AES-192 and AES-256 are impractical outside a lab setting, but they do nonetheless provide theoretical proof that versions of AES are susceptible to attack. Think about all the newer practical uses for chaining gaming systems together and cryptooffloading to video acceleration cards.

Message Digest 5 Algorithm

With the development of the Internet and the evolution of the world to become oriented in data and connectivity, we have also learned that “there be hackers” out there. This means that you must be concerned with issues such as security, authenticity, and integrity of data.

These issues are important for almost everyone, from the military/government to health-care/personal records to financial data. All organizations require secret or private data to be kept from those who should not have access to it. Security in the form of authenticity and integrity of data is driven as follows:

- *Authenticity* is responsible for ensuring that the group or person sending the data is who he says he is. A digital signature is an example of the importance of authenticity.

- *Integrity* is responsible for ensuring that the data is not altered during transmission and that exactly what was sent was received. Have you ever downloaded a software application or operating system patch? It is important that the downloaded file has not lost any of its integrity; this is the importance of integrity.

Message Digest 5 (MD5) is one of the better available methods of ensuring that these security needs are met. A message-digest algorithm is designed to accept data and generate fixed-length output; this output is called a *hash value*, *fingerprint*, or *message digest* and is the key to the security that MD5 provides.

Note The term *hash* comes by way of analogy with its standard meaning in the physical world: to chop and mix. When teaching, I often run across technologies that hash. I find that the best way to explain a hash and make it memorable to students is through an analogy. A hash is basically a grinder that takes something recognizable—such as beef or pork—hashes it, and the result is something based on the original but is unique. In this case, it is hamburger or sausage, of course! Try and put that back together.

Developed in 1994 by Rivest, MD5 is a one-way hash algorithm that takes any length of data and produces a 128-bit nonreversible fingerprint known as a hash. (RFC 1321 officially describes MD5.) This output hash/fingerprint cannot be reverse engineered to determine the data that was used to produce it. Functionally, this means that it is impossible to derive the original file contents from the MD5; this is why they call it one way.

Note A one-way hash is the result of an algorithm that turns data of any type into a string of digits, thus creating a digital signature. These digital signatures are then used to verify the authenticity and integrity of data. Like a written signature, the purpose of a digital signature is to guarantee that the individual sending the message actually is who she claims to be.

MD5 does not actually *encrypt* or alter any data; instead, it creates a hash from which the data's authenticity and integrity can be determined. Because MD5 does not encrypt data, it is not restricted by any exportation rules. You can freely use and distribute this MD5 anywhere in the world.

Note Authentication is the process of identifying an individual or device based on the correct username and password combination. Authentication does not determine what an individual is allowed to access, but merely that he is who he claims to be. Authorization defines what an individual is allowed to access—assuming that he has been authenticated, of course!

The following section looks at MD5 in action and where you might have unknowingly encountered it. The actual mathematics of how MD5 creates these hashes is beyond the

scope of this book. Readers wanting to learn more about MD5 are encouraged to read RFC 1321, “The MD5 Message-Digest Algorithm” (<http://tools.ietf.org/html/rfc1321>).

MD5 Hash in Action

If you own a computer, you have most likely experienced MD5 without even knowing it. MD5 plays a large role in networking, and it can help you in a variety of ways:

- When downloading files from the Internet, you can use MD5 to ensure that the downloaded file has been unaltered after being made available on a server. The MD5 hash is calculated after a file is downloaded and compared.
- Ensure that the integrity of system files is maintained—various tools, such as tripwire (covered later), use MD5 to monitor and consistently verify that operating system files have not been altered. This protects crucial systems and alerts administrators if something has changed because the hashes no longer match.

When using a one-way hash operation such as MD5, you can compare a calculated message digest against the received message digest to verify that the message has not been tampered with. This comparison is called a *hash check*.

MD5 checksums are widely used in software development to provide assurance that a downloaded file or patch is unaltered. By verifying a published MD5 checksum and comparing MD5 checksum on record with the software provider with a downloaded file's checksum, a user can be sure that the file is the same as that offered by the developers if a match occurs.

This comparison procedure protects everyone by providing a measure of protection when downloading software and by ensuring that no Trojan horses or computer viruses exist. As previously discussed, this is the definition of a digital signature. Digital signatures are especially important for electronic commerce and are a key component of most authentication schemes. To be effective, digital signatures must be unbreakable, which is an idealistic goal. As a viable compromise, the signature must be independently verifiable, difficult to break, and have a design that enables its strength to increase and evolve. As demonstrated in the discussion of DES, the growth of technology can quickly overtake security if you do not take the proper precautions or follow up on updated security needs.

Secure Hash Algorithm (SHA Hash)

The Secure Hash Algorithm, or SHA Hash, is published by the National Institute of Standards and Technology (NIST) as a U.S. Federal Information Processing Standard - FIPS PUB 180-3, which specifies three flavors of the SHA Algorithm:

- **SHA-0:** No longer used.
- **SHA-1:** The most widely used version
- **SHA-2:** Comes in four different variants: SHA-224, SHA-256, SHA-384, and SHA-512

When a message of any length less than 2^{64} bits (SHA-1, SHA-224, and SHA-256) or less than 2^{128} bits (SHA-384 & SHA-512) is input to a hash algorithm, the result is an output called a *message digest*. The message digests range in length from 160 to 512 bits, depending on the algorithm.

The five hash algorithms specified in this standard are called secure because, for a given algorithm, it is computationally infeasible to find a message that corresponds to a given message digest, or to find two different messages that produce the same message digest. Any change to a message will, with a high probability, result in a different message digest. This will result in a verification failure when the secure hash algorithm is used with a digital signature algorithm or a keyed-hash message authentication algorithm.

Types of SHA

Of the three flavors, I'm going to concentrate on the variants of SHA: SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512. Following is an informative overview—without getting too deep into the weeds. All these are a cryptographic hash function designed by the National Security Agency (NSA) and published by NIST.

SHA-1

The original specification of the algorithm was published in 1993 in FIPS PUB 180-1. This is the most widely used of the existing SHA hash functions and is employed in several widely used security applications and protocols, such as transport layer security (TLS), secure socket layer (SSL), pretty good privacy (PGP), Secure Shell (SSH), Secure/Multipurpose Internet Mail Extensions (S/MIME), and Internet Protocol Security (IPSEC). SHA-1 hashing is also used in distributed revision control systems such as Arch, Mercurial, Monotone, and BitKeeper to identify revisions and detect data corruption or tampering. And, yes, even when you're at home enjoying some guilty pleasure of killing a complete stranger over the Internet through your Nintendo or trying to stay fit using your Wii, the SHA-1 hash is being used for signature verification during your boot process.

SHA-2

In August 2001, NIST published FIPS PUB 180-2, introducing SHA-2 to the general populace. SHA-2 includes a significant number of changes from its predecessor, SHA-1. SHA-2 is a family of four similar hash functions with differing digest lengths, known as SHA-224, SHA-384, SHA-256, and SHA-512. These algorithms are collectively known as SHA-2. The same vulnerabilities found in SHA-1 in 2005, these same attacks have not been extended to SHA-2 or its variants.

Like its predecessor, the SHA-2 hash function has been implemented in TLS and SSL, PGP, SSH, S/MIME, and IPsec. However, SHA-2 implementation is not as widely used as SHA-1, despite its better security. Reasons vary: lack of support on Microsoft systems older than Windows XP SP2, a lack of urgency, or perhaps even waiting for SHA-3 to come around (see the note). Currently, SHA-256 is used for authentication on certain Linux

packages; SHA-512 is also a part of an authentication system for archival video from the International Criminal Tribunal of the Rwandan genocide. UNIX and Linux vendors are pushing for use of the SHA-256 and SHA-512 for secure password hashing.

Note SHA3: Now I know you are saying to yourself, he didn't mention it beforehand, and you're right. SHA-3 is a new hash standard currently under development. There is an ongoing NIST hash function competition that is scheduled to select a winning function in 2012. The new SHA-3 algorithm will not be derived from SHA-2.

Point-to-Point Tunneling Protocol (PPTP)

This section discusses the Point-to-Point Tunneling Protocol (PPTP), which was developed by Ascend Communications, Microsoft Corporation, 3Com/Primary Access, ECI Telematics, and U.S. Robotics.

PPTP operates at Layer 2 of the OSI reference model and is based on the Point-to-Point Protocol (PPP) standard for dial-up networking that enables any user with a PPP client to use an Internet service provider (ISP) to connect to the Internet. PPTP builds on the functionality of PPP, which is used for broadband access to commercial networks, by enabling users to securely connect via a VPN (covered in Chapter 8, "Router Security") to secure networks such as that of their employers or business partners.

PPTP is a protocol (or a set of communication rules) that enables corporations to extend their own corporate network through private "tunnels" over the public Internet. Effectively, a corporation uses a WAN as a single LAN. A company no longer needs to lease its own lines for wide-area communication but can securely use the public networks. This kind of interconnection is known as a *virtual private network (VPN)*.

PPTP Functionality

PPTP packages data within PPP packets and then encapsulates the PPP packets within IP packets (datagrams) for transmission through an Internet-based VPN tunnel. PPTP supports data encryption and compression of these packets. PPTP also uses a form of Generic Routing Encapsulation (GRE) to move data to and from its final destination.

PPTP-based Internet remote access VPNs are by far the most common form of PPTP VPN. However, PPTP VPNs are not the most common VPNs in use. IPsec is far more secure and popular today. Cisco IOS Software does enable the use of PPTP VPNs; however, you should consider the shortcomings, which are explained in the section, "Limitations of PPTP." When PPTP tunnels are established with a two-step creation process

1. The user wanting to connect using a PPTP client connects to his ISP using PPP dial-up networking (in most cases, traditional modem or ISDN)—or perhaps the user is permanently connected via cable modem, for example.
2. The PPTP client is launched, and it creates a control connection via TCP (port 1723) between the client and the server, thereby establishing the tunnel.

After the PPTP tunnel is established, two types of rather obvious packets of information flow through this tunnel: *control messages*, which manage the PPTP tunnel, and data packets.

PPTP relies on the inherent functionality of PPP to maintain the connection, encapsulate packets, and authenticate users. PPTP uses the Challenge Handshake Authentication Protocol (CHAP) or the Password Authentication Protocol (PAP).

PPTP directly handles maintaining the VPN tunnel and transmits data through the tunnel. PPTP also supports some additional security features for VPN data beyond what PPP provides.

PPTP remains a popular choice for VPNs, thanks to Microsoft. PPTP clients are freely available in all popular versions of Microsoft Windows. Windows servers and certain Cisco devices can function as PPTP-based VPN servers to terminate PPTP client connections.

Limitations of PPTP

As just discussed, the extensive use of Microsoft products is driving PPTP in many ways. Although other PPTP providers are available because many organizations have Windows servers, it is natural to want to use what you have. There are some limitations and drawbacks to using PPTP that revolve around its use in general.

The PPTP standard does not define how authentication and data encryption tasks are to be handled. This means that two different vendors might produce a PPTP-capable device or client, and yet they might not be able to work together, which introduces compatibility issues within an organization using different PPTP implementations. For example, if Vendor A implements PAP and Vendor B implements CHAP, they will not interoperate.

Concerns also persist around the security involved in the use of PPTP connections when compared with other available solutions—specifically, surrounding the implementation of Microsoft’s PPTP solution, which is the leader today.

A company called Counterpane Internet Security is a managed security services provider founded by Bruce Schneier, who is also the current CTO. Mr. Schneier is also an author who helped develop several encryption technologies—specifically, Blowfish and Twofish. Mr. Schneier also wrote *Applied Cryptography*. These folks have written some excellent and detailed papers on a variety of security-related subjects. Specifically, Bruce Schneier of Counterpane and Mudge of L0pht Heavy Industries conducted detailed analysis on Microsoft’s implementation of PPTP. In their own words, they summarize their findings:

“The Point-to-Point Tunneling Protocol (PPTP) was designed to solve the problem of creating and maintaining a VPN over a public TCP/IP network using the common Point-to-Point Protocol (PPP). Although the protocol leaves room for every type of encryption and authentication imaginable, most commercial products use the Microsoft Windows NT version of the protocol. This is the implementation that we crypto-analyze in this paper.

We have found Microsoft’s authentication protocol to be weak and easily susceptible to a dictionary attack; most passwords can be recovered within hours. We have found the encryption (both 40-bit and 128-bit) to be equally weak, and have discovered a series of bad design decisions that make other attacks against this encryption possible. We can open connections through a firewall by abusing the PPTP negotiations, and can mount several serious denial-of-service attacks on anyone who uses Microsoft PPTP.”

If you would like to read the entire paper or refer others to it, you can find it at

www.counterpane.com/pptp-paper.html

This additional paper from Counterpane Internet Security covers its findings and analysis of Microsoft’s implementation of CHAP—including both of the versions, MS-CHAP and MS-CHAPv2, respectively:

www.counterpane.com/pptpv2-paper.html

For a little humor regarding the findings, refer to the FAQs at

www.counterpane.com/pptp-faq.html

If you have taken the time to read through some of these papers on the Counterpane website, you are probably wondering how you can implement a VPN securely if PPTP is not advised. The primary alternative to PPTP is an IPsec-based VPN. IPsec is an open standard that has been developed under the direction of the Internet Engineering Task Force (IETF) in its normal public process and is not owned by any one company. This is an important distinction because the manner in which Microsoft has implemented PPTP has made it proprietary to Microsoft. The following section examines the Layer 2 Tunneling Protocol (L2TP), which is also an extension of PPP. Chapter 9 discusses IPsec.

Layer 2 Tunneling Protocol (L2TP)

Layer Two Tunneling Protocol (L2TP) is an extension of the PPTP that is documented and defined in RFC 2661. L2TP is used to enable the operation of a VPN over the Internet. RFC 3193 defines using L2TP over a secure IPsec transport. In this approach, L2TP packets are exchanged over User Datagram Protocol (UDP) port 1701. IPsec Encapsulating Security Payload (ESP) protects UDP payload to ensure secure communication. Cisco and Microsoft agreed to merge their respective L2TP, thereby adopting the best features of two other tunneling protocols: PPTP from Microsoft and Layer 2 Forwarding (L2F) from Cisco.

The two main components that make up L2TP are the L2TP Access Concentrator (LAC), which is the device that physically terminates a call, and the L2TP Network Server (LNS), which is the device that terminates and possibly authenticates the PPP stream.

L2TP is similar to PPTP in its use of PPP and in both function and design. In this blending of two of the largest IT-related companies, some areas definitely benefited—specifically, the area of securing sensitive data.

L2TP Versus PPTP

L2TP and PPTP have a variety of features and benefits in common that reflect their original design and function within networking. These similarities are as follows:

- Both provide a logical transport mechanism for sending PPP payloads.
- Both provide tunneling and encapsulation so that PPP payloads based on any protocol can be sent across an IP network.
- Both rely on the PPP connection process to perform user authentication and protocol configuration.

Although L2TP and PPTP share some similarities, they are different in the following ways:

- With PPTP, data encryption begins after the PPP connection process (and therefore PPP authentication) completes. With L2TP/IPsec, data encryption begins before the PPP connection process.
- L2TP/IPsec connections use either DES or 3DES—again, we strongly prefer 3DES.
- PPTP requires only user-level authentication, and L2TP requires the same user-level authentication, as well as computer-level authentication through a computer certificate.

The following section discusses some of L2TP's important benefits and how it can be used more securely than its predecessor, PPTP.

Benefits of L2TP

ISPs have been able to build VPN solutions using L2TP (because of its Internet standard status) as the method in which customers gain the benefits of VPNs within a carrier's network. Some of the more specific benefits of L2TP include the following:

- Because it is standards-based, interoperability of L2TP-capable devices between vendors is greatly increased.
- L2TP VPNs have become products for service providers.
- In Cisco-powered networks, end-point-to-end-point quality of service (QoS) can be provided through the use of QoS technologies such as DiffServ to categorize, tag, and prioritize traffic accordingly.
- IPsec is responsible for the encryption, which is also standard-based (that is, defined in RFC 4308 most recently from 2005). IPsec provides per-packet data origin authentication (proof that the authorized user sent the data), data integrity (proof that the data was not modified in transit), replay protection (prevention from resending a stream of captured packets), and data confidentiality (prevention from interpreting captured packets without the encryption key). By contrast, PPTP provides only per-packet data confidentiality.

- Support for multiprotocol environments because, by design, L2TP can transport any routed protocols, including IP, IPX, and AppleTalk. L2TP also supports any WAN transmission technology, including Frame Relay, ATM, X.25, or SONET. It also supports LAN media such as Ethernet, Fast Ethernet, Token Ring, and FDDI.

In many ways, L2TP is the best of both vendors (Cisco and Microsoft); personally, I think Microsoft was the big winner because its tinkering with PPTP left a lot to be desired. The following section examines how L2TP functions.

L2TP Operation

As discussed previously, L2TP enables the support of legacy protocols and over the tunnel through the use of GRE. This permits an architecture to be created that enables L2TP tunnels to connect rather easily over the public Internet or dial-up.

Note Traditional dialup networking services support only registered IP addresses, thereby limiting the types of applications implemented over VPNs.

Figure 6-2 shows a common architecture used when an L2TP network is implemented. In this figure, note that the equipment shown is what an ISP or carrier would use when implementing a complete L2TP solution with all the aspects and benefits that we have described. It is commonplace for companies to use a subset of this design on which to build based on current and future requirements.




Figure 6-2 L2TP Network Architectures

L2TP uses the Internet and its network connections to make it possible for its endpoints to be in different geographic locations. In Figure 6-2, the user's PC creates a dial-up connection (Layer 2) to the L2TP Access Concentrator (LAC), which then authenticates them using the AAA server and forwards the connection, which is encrypted, to the L2TP network server.

L2TP's greatest security strength is its use of standards-based IPsec, which provides connections with confidentiality, per packet authentication, and antireplay protection for control and data packets. In contrast, the Microsoft Point-to-Point Encryption (MPPE) used by PPTP encrypts only data and does not prevent forgery or replay, like IPsec does.

The following list describes the actual call sequence steps as home users used to dial in to their ISP to create an L2TP connection to their corporate office:

1. The remote user uses the analog telephone system or broadband to initiate a PPP connection from her home to an ISP.
2. The ISP network LAC accepts the connection at its point of presence (POP), and the PPP link is established.
3. After the end user and LNS negotiate LCP, the LAC partially authenticates the end user with CHAP or PAP. The username, domain name, or DNIS is used to determine whether the user is a VPDN client. This is how ISPs can offer these services because each company and user is unique. The AAA server connected to the LAC defines each user.
4. If the user is not a VPN client (using L2TP), authentication continues and the client accesses the Internet as a normal user. If the user is a VPN client, her connection names a specific endpoint (the *L2TP network server [LNS]*) where the user's VPN terminates. The user's information is sent to the AAA server, which is connected to the LNS, for further authentication.
5. The tunnel endpoints—the LAC and the LNS—authenticate each other before any data is transmitted from the user into the tunnel.
6. After the VPN tunnel (using L2TP) is created, an L2TP session is created for the end user to the corporate network.

The end result is that the exchange process appears to be between the dial-up client and the remote LNS exclusively, as if no intermediary device (that is, the LAC) is involved. Figure 6-3 offers a visual representation of the L2TP incoming call sequence with its own corresponding sequence numbers. Note that the sequence numbers in Figure 6-3 are not related to the sequence numbers described previously.

The following section examines one of my favorite protocols and tools for IT professionals today: Secure Shell (SSH). It is also a robust security protocol; no good book could be written without mentioning it, so I have to include it!

Secure Shell (SSH)

SSH is used to log in to a remote computer system using port 22, much in the same way that Telnet (port 23) has been used in the past for the same purpose. The big difference between Telnet and SSH, however, is that SSH provides significantly enhanced security for your connection. SSH is a program/client that provides an encrypted communications

path between two hosts over an untrusted, potentially insecure network such as the Internet. Therefore, it prevents users' passwords and other sensitive data from being transmitted across the network in clear-text form. SSH helps solve one of the most important security problem on the Internet: hackers stealing or cracking passwords.

Used since 1995, SSH1 was designed to replace the nonsecure UNIX commands (**rlogin**, **rsh**, and **rcp**). These protocols provided UNIX users with a variety of useful tools; however, they were fraught with security concerns. The IETF released SSH2 in 1997 and improved the security and functionality of SSH1. SSH1 is slowly being phased out in favor of SSH2.

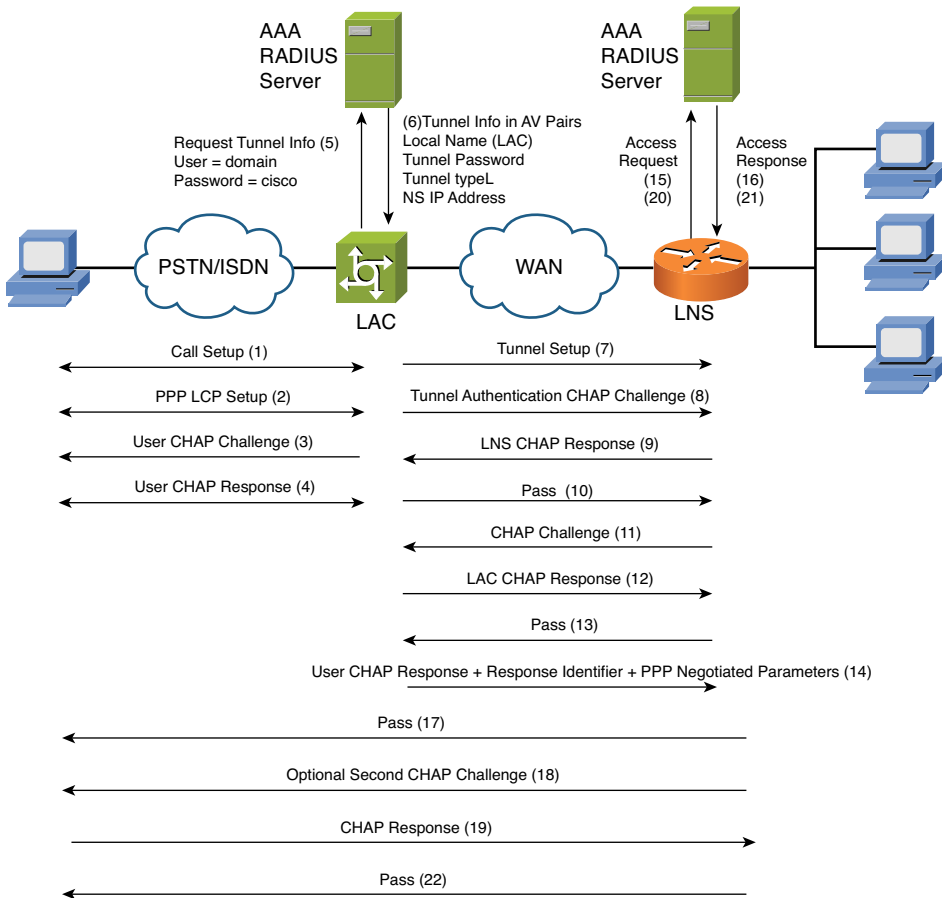


Figure 6-3 L2TP Creation Steps

Note You might be wondering what the difference is between SSH1 and SSH2 and whether they are compatible. In a nutshell, they are not compatible, and SSH2 is a complete rewrite of SSH1 resulting in a completely different protocol implementation. SSH2 encrypts packets more securely and references only host keys because it exchanges a hash.

The most common use of SSH is for creating a secure command shell (remote login) like the more common protocol, Telnet. However, SSH takes the basic functionality and vulnerabilities of Telnet and solves them in a manner that has made SSH the de facto connection standard for secure network remote device access.

SSH extends Telnet capabilities both in features and functionality. Today, SSH is available as a client on virtually all computer platforms: Macintosh, Microsoft Windows, UNIX, Linux, and so on.

Typical SSH applications include remote access (login) to computer resources over the Internet or via some other untrusted network where you want to perform one of the three core SSH capabilities:

- Secure command shell
- Secure file transfer
- Secure port forwarding

Although remote login is the primary use of SSH, you can use the protocol as a general-purpose cryptographic tunnel capable of copying files, encrypting email connections, and triggering remote execution of programs. Your company's remote access security policy should require SSH and disallow Telnet for secure remote access to company systems and partner extranets.

SSH Versus Telnet

Telnet is quite insecure for so many reasons: It has no protection, encryption, or any way to protect your password or any activity you conduct via Telnet. When this book was first published, a simple Google search revealed 53,400 hits when searching for "telnet vulnerabilities." The same search done today revealed 147,000 hits—now automated tools, both licensed and commercial, regularly exploit them.

SSH is better than Telnet because of its built-in encryption; however, the benefits of SSH do not stop there. SSH offers additional features and benefits as follows:

- Denies IP spoofing of packets, thereby ensuring you know the host that is sending the packets
- Encrypts packets to prevent the interception of clear text passwords and other data by intermediate hosts
- IP source routing by preventing a host from pretending that an IP packet comes from another, trusted host
- Prevents the manipulation of data by people in control of other devices along the route of your packets

A much simpler way to look at this is that SSH doesn't trust any device other than the one with which it is trying to establish a secure connection. On the other hand, Telnet trusts anyone and pretty much anything by default.

Perhaps showing you just how easily people with malicious intent could use easily downloadable tools to gather even basic information about your network and then place themselves in a position for an ARP poisoning attack would help make the point. It took me 10 minutes to find, download, and install ettercap—a software suite for man-in-the-middle (MitM) attacks on a LAN. The software supports various platforms: Linux, FreeBSD, OpenBSD, Mac OS, Windows 2000/XP/2003, and Solaris 2.x, just to name a few. Using this freeware tool, hackers, script kiddies, or disgruntled employees can position themselves for a MitM.

Note An MitM is a form of active eavesdropping where the hacker makes connections with two nodes (one being the user, the other the target) and intercepts data between the two, making them believe that they are talking to each other over a private connection.

When installed, I can open ettercap and scan the network segment I'm on and view all hosts currently attached to the network (see Figure 6-4).

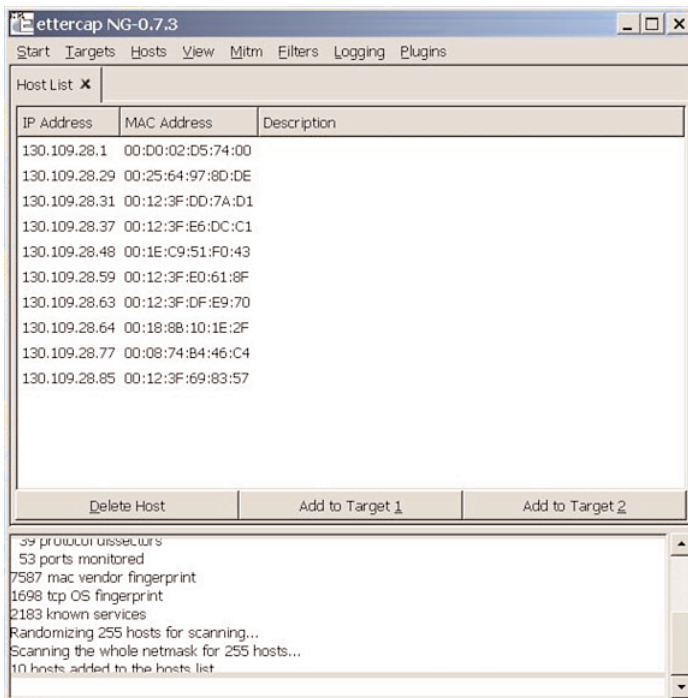


Figure 6-4 ettercap Reveals Attached Hosts

After seeing and selecting the hosts I'd like to intercept traffic between, be it an SSH or a Telnet session, I can then pick which type of MitM attack I want to accomplish (see Figure 6-5). At this point, if communication were with a Telnet session, I could gather usernames and passwords and then wreak havoc on routers and servers alike. Because the use of ettercap to wreak havoc is not the focus of this book, I'll skip that exercise. But just know and understand that SSH is secure and encrypted, whereas Telnet is not and is susceptible to MitM attacks.

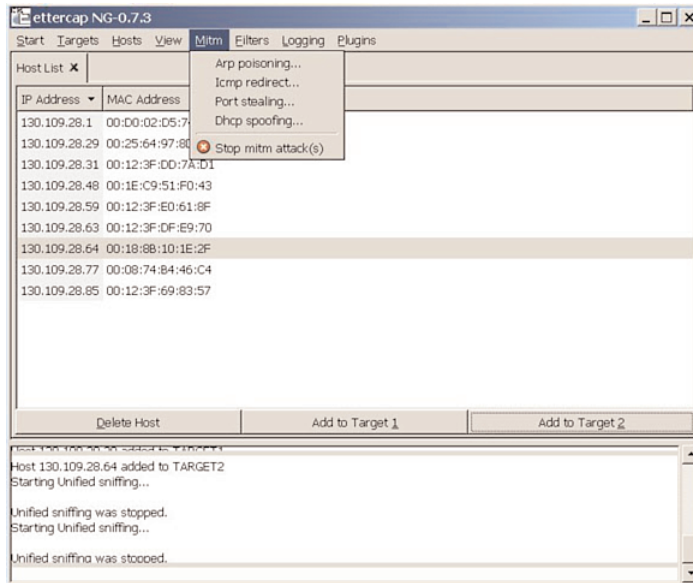


Figure 6-5 Man-in-the-Middle Attack Vectors Within the ettercap Application

Note ettercap version 0.7.3 is a multipurpose sniffer/interceptor/logger for switched LAN. It supports active and passive dissection of many protocols (even ciphered ones) and includes many features for network and host analysis. If you are interested in learning more about ettercap, go to <http://ettercap.sourceforge.net>.

SSH Operation

To review, SSH is used to connect two different hosts using an encrypted communication session. In its simplest mode of operation, SSH uses TCP to connect to a host and authenticates using a username and password; authentication is successful, and SSH begins encrypting data.

Depending on the version of SSH, a variety of different *encryption methods* can be available for use, as shown in Table 6-1.

Table 6-1 *SSH Encryption Methods*

Encryption Method	SSH1	SSH2
3DES	Yes	Yes
IDEA	Yes	No
BLOWFISH	Yes	Yes
TWOFISH	No	Yes
ARCFOUR	No	Yes
CAST-128-CBC	No	Yes

Connecting keys are used during the authentication phase of SSH. Depending on the version of SSH used, either an RSA or DSA key is used with a pair created, one public and one private. Depending on the version of SSH, a variety of different authentication methods can be available for use, as shown in Table 6-2.

Table 6-2 *SSH Authentication Ciphers*

Authentication Cipher	SSH1	SSH2
RSA	Yes	No
DSA	No	Yes

The private key is stored encrypted while the public key is stored on the users' machine if they authenticate properly. This enables SSH software clients to automatically connect because the key is stored for use any time the user starts a connection.

Tunneling and Port Forwarding

SSH brings an interesting feature to the realm of information security: the concept of forwarding certain traffic (identified by port number) via SSH in a tunnel. The two most common protocols to take advantage of this feature are FTP and X Window. This forwarding feature provides SSH with the capability to use these other protocols for conducting operations on the host terminating the SSH connection. Perhaps the other end is a web server, and you want to upload new files over the Internet, or you want desktop type access to the device using X Window.

Note The best SSH client I have found and recommend is SecureCRT, from VanDyke Software (www.vandyke.com). The most recent version is version 6.5.4. This GUI tool provides for some excellent built-in benefits, such as automatic logging, customizable scripts, and adjustable buffers.

Figure 6-6 shows easy configuration of port forwarding. Notice also that X Window can be easily forwarded because it is so commonly used with SSH.



Figure 6-6 SecureCRT Allows SSH Tunneling

Limitations of SSH

SSH version 1 (SSH1) had several bugs and problems, so choose SSH version 2 (SSH2) if you have a choice. Choosing version 2 eliminates most of the limitations and network inconsistencies by version and policy.

SSH does not help you protect any of your internal systems (PCs, servers, and so on); therefore, after an attacker gains access to one of those systems, he has access to SSH and it, too, can be subverted for his use.

The ability to tunnel through an SSH connection might make you think that it would be a good alternative to a VPN, but that is not the case. A better solution would be tunneling via SSH through a VPN connection—now that is a more secure connection!

This section of the chapter covered SSH rather broadly and gave you enough information to understand what is happening. If you would like to learn more about SSH, check out the following book:

Barrett, Daniel J. and Richard Silverman. *SSH, The Secure Shell: The Definitive Guide*. Cambridge, MA: O'Reilly & Associates; 2001.

In summary, SSH is a popular and powerful tool/client for encrypting TCP sessions over a network. It is most commonly used for remote login but also has other uses for increasing your company's information security posture.

SNMP v3

Simple Network Management Protocol Version 3 (SNMP v3) is defined by RFC 3411–RFC 3418. It has primarily added security and remote configuration enhancements to SNMP. As of 2004, the Internet Engineering Task Force (IETF) recognizes SNMP v3 as

the de facto standard version of SNMP. The IETF deemed it a full Internet standard, the highest maturity level for an RFC, and consider all previous versions of SNMP to be obsolete. The IETF cut them off like Michael Corleone said to his brother, "...you're nothing to me now. You're not a brother; you're not a friend. I don't want to know you or what you do...."

SNMP v3 is derived from, and builds upon, the original SNMP and SNMP v2. All versions share the same basic structure and components: managed nodes, an agent (software that runs on managed devices), and a network management system (NMS) (software that runs on the manager).

A *managed node*, or device, is a device that implements an SNMP interface that enables unidirectional (read-only) or bidirectional access to information specific to itself. Managed devices exchange node-specific information with the NMS. These managed nodes can be any type of device; examples are routers, access servers, switches, IP Phones, IP video cameras, computers, or printers.

An *agent* is a network-management software component that resides on a managed device. An agent has local device-specific knowledge of management information and translates that information into an SNMP-specific form.

An NMS executes applications that monitor and control managed devices. An NMS provides the bulk of the processing and memory resources required for network management.

Security Built In

SNMP v3 provides confidentiality, integrity, and authentication to the SNMP suite; confidentiality by encrypting the packets to prevent snooping by unauthorized personnel; message integrity to ensure that a packet has not been tampered with in transit from the managed device; and authentication by verifying the message is from a valid source.

Figure 6-7 and Figure 6-8 show the configuration of SNMP (versions 1 and 3) on the same Cisco ASA. Figure 6-7 shows the configuration of a basic SNMP v1 implementation, whereas Figure 6-8 shows the SNMPv3 implementation. The differences in the built-in security are quite extensive, as the figures reveal.

With the original version of SNMP, the most security that you could do was establish a different community string and port number. The hope was that the hacker trying to get into your network got distracted by something shiny. With SNMPv3 (see Figure 6-8) you have options to set up various authentication algorithms (MD5 and SHA) and then encryption algorithms for the transfer of data (DES, 3DES, or AES). As mentioned several times before, do not use DES. In this instance you want to use 3DES or AES. In this case, we have chosen to use 3DES. You also have the option to send your password encrypted or clear text—do not do this. Even though you set your authentication and encryption levels high, it is still highly suggested you encrypt your password.

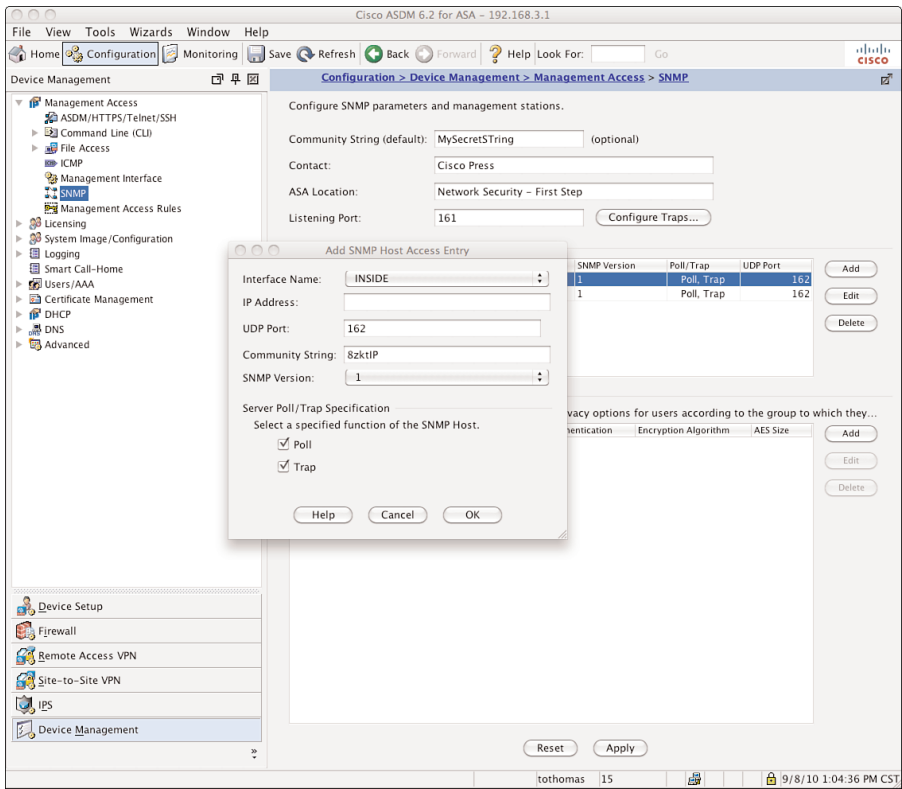


Figure 6-7 *SNMPv1 Implementation/Configuration*

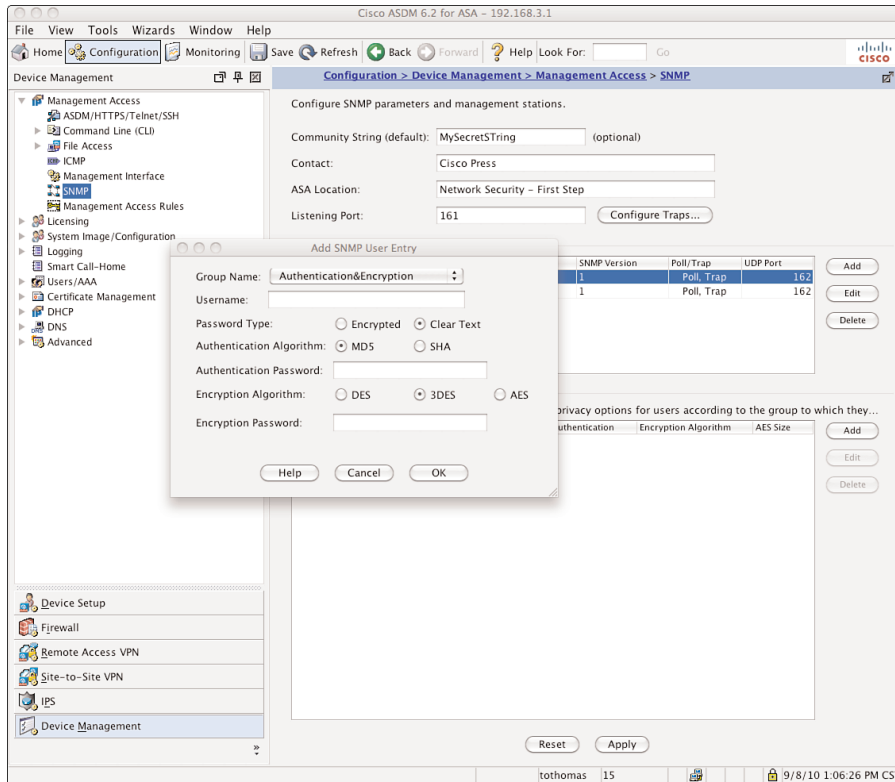


Figure 6-8 *SNMPv3 Implementation/Configuration*

Even with all the built-in security enhancements for SNMPv3, the protocol still has several security implications that you should be aware of:

- It is subject to brute force and dictionary attack tools for guessing the community strings, authentication strings, encryption strings, and encryption keys.
- It is mostly used over UDP, which is a connectionless protocol and vulnerable to IP spoofing attacks.

Our advice to you is to implement this protocol only if absolutely necessary. Otherwise, go into every device that you have and rename the public and private community strings and disable SNMP if you can—this is standard practice for all U.S. military and governmental agencies. This is a critical step for your company to consider and implement, making sure that it is supported in your policy infrastructure.

Chapter Summary

This chapter discussed the importance and functionality of the DES and 3DES encryption algorithms. You saw the complex math involved in each encryption technology to demonstrate the difficulty of cracking them—unless simple passwords are used, such as passwords with all letters.

You can tunnel and protect traffic within a network in several ways, and this chapter covered just two methods: PPTP and L2TP. It discussed each of the benefits and recommended L2TP because it combines the best aspects of both Microsoft and Cisco technologies. This chapter concluded with a discussion of SSH and the value it brings to your information security posture.

Chapter Review Questions

The following questions reinforce the concepts that were covered in this chapter.

1. How long, in bits, is the DES key?
2. True or False: In 3DES, the same key is used to encrypt at each of the three stages.
3. Define a hash in your own words.
4. What is used to create a digital signature?
5. Define authentication and provide an example.
6. Define authorization and provide an example.
7. A hash check occurs at what point in the operation of MD5?
8. Of the security protocols covered in this chapter, which of them use generic routing encapsulation (GRE)?
9. Describe several security benefits of L2TP.
10. What are the three core SSH capabilities?

Firewalls

“Courage is resistance to fear, mastery of fear—not absence of fear.”—Mark Twain

By the end of this chapter, you should know and be able to explain the following:

- Who needs a firewall, and why firewalls are used to protect network resources
- How a firewall is a technological expression of your organization’s written security policy
- When a DMZ is appropriate and the security benefits you gain by deploying a firewall with a DMZ

Answering these key questions enables you to understand the overall characteristics and importance of network security. By the time you finish this book, you will have a solid appreciation a firewall’s role, its issues, how it works, and why it is so important to the security of your network.

The Internet is an exciting and wonderful place to browse and explore. It has been likened to the Wild West, The Great Frontier, and other grandiose achievements of mankind. In reality, the World Wide Web is merely a collection of routers and servers that make up the largest WAN in recorded history. This collection of networking gear provides mail servers, websites, and other information storage and retrieval systems and is all connected to the Internet and accessible to every person who is also connected. It has even been said that the Internet will contain the collective institutional knowledge of mankind, eventually. Entire books have been written on the Internet’s potential and its impact on our lives—rest assured that this is not one of those books. But it does make you ponder just how much of your life is out there already that you might or might not be aware of.

We are concerned with a network’s security, so we must ask what kinds of safeguards are in place to protect such an unbelievable amount of information. Is there some organization that polices the Internet much in the same way that law enforcement cruises the highways? How about a governmental agency that snoops around and double-checks every

possible device connected to the Internet? The answer to these questions is no; there is no unifying organization responsible for protecting the Internet.

The job of securing and protecting the gateways of the Internet's knowledge is left up to the person or persons responsible for the Internet connection and network hardware/software, such as the router, firewall, switch, server operating systems, application, and so on. This person or persons are tasked with the job to ensure that hackers (the bad guys) do not make a mess of the carefully stored and catalogued information in question. And just how can you protect a website, mail server, FTP server, or other information sources accessible from the Web?

The answer is one word—firewall. The sole purpose of these dedicated hardware devices is to provide security for your network. A *firewall* is a security device that sits on the edge of your Internet connection and functions as an Internet border security officer. It constantly looks at all the traffic entering and exiting your connection, waiting for traffic it can block or reject in response to an established rule. The firewall is the law and protection in the lawless wild wild web. A firewall is ever vigilant in its mission to protect the network resources connected to it.

The Internet has made so much information available to individual users as, over the years, access to this information has evolved from an advantage to an essential component for both individuals and businesses. However, making your information available on the Internet can expose critical or confidential data to attack from everywhere and anywhere in the world—the Internet is literally a worldwide network. This means that, when you connect to the Internet in Madison, Mississippi, you can be subject to attacks from Europe, Asia, and Russia—literally any device connected to the Internet anywhere on the earth, which is kind of disturbing. Firewalls can help protect both individual computers and corporate networks from hostile attacks from the Internet, but you must understand your firewall to correctly use it.

This 24-hour/365-day-a-year “electronic Robocop” has an important job: to keep the bad guys out and let the good guys get to the resources they need to do their jobs. Sounds simple, right? On paper, it sounds like a walk in the park, but in reality, properly configuring a firewall is far from easy.

In some cases, a badly configured or feature-inadequate firewall can be worse than no firewall at all. This is difficult to believe, isn't it? Nonetheless, it is true. This chapter dissects a firewall's duties to understand what makes a firewall operate and how it does its job.

Firewall Frequently Asked Questions

Before looking at the overall operation of a firewall, the following sections examine and answer some of the fundamental questions about them.

Who Needs a Firewall?

This is perhaps the most frequently asked security question. If you plan to connect to the Internet, you need a firewall. It does not matter whether you connect from home or your company connects—you need a firewall, period! The increased penetration of broadband Internet services to the home and their always-on Internet connections make home security even more important.

Why Do I Need a Firewall?

You read about security threats in the papers or hear about them on the evening news almost every day: viruses, worms, denial-of-service (DoS) attacks, hacking, and new vulnerabilities to your computer. For example, Code Red, Slammer, and other threats/vulnerabilities are changing with the prevalence of malware and botnets.

It is no secret that hackers are out there, and they are out to get you. Often, you do not know who they are, but you do know where they are and where you do not want them to be (in your network). Like pirates of old who roamed the seas, hackers freely roam the open expanses of the Internet. You do not want them to enter your network and roam among the computers that connect to it, and that is where a firewall becomes a requirement.

You know that you must protect your network from these attackers, and one of the most efficient methods of protecting your network is to install a firewall. By default, any good firewall prevents network traffic from passing between the Internet and your internal network. This does not mean that the firewall can stop all traffic—that defeats the purpose of being on the Internet. It does mean that the firewall is configured to allow only web browsing (HTTP/port 80) to access it from the Internet. Along the way, the firewall provides Stateful Packet Inspection (SPI) rules to every incoming packet (as discussed previously in Chapter 2, “Security Policies.”)

The alternative to having a firewall is allowing every connection into your network from anyone, anywhere—there wouldn't be any sort of packet inspection to determine whether an attack is hidden within one of the incoming packets. Not having a firewall is ill-advised and will make your organization wide open to everyone on the Internet.

Do I Have Anything Worth Protecting?

I often hear people say, “I understand that if I had something worth protecting, I would definitely need a firewall. However, I do not have anything an attacker would want, so why should I worry about a firewall?”

Networks and their resources are important to the way our society conducts business and operates. In practical terms, this means that there is value to your network and having it effectively operate. This increased role of networks means that you definitely have something worth protecting to some degree, as documented in the following list:

- **Downstream liability:** This sounds like a confused Bassmasters fishing show title, but it is perhaps the next big step in the legal evolution of the Internet. Downstream lia-

bility involves allegations that an attacker has taken control of a target computer (yours) and used it to attack a third party. Assume that it is your company's computer that has been compromised by a hacker. Your company's failure to protect its own systems has resulted in the damaging of a third party; the attacker used your computer as a weapon against the third party. Your company is therefore negligent due to lack of due diligence because it failed to protect against reasonable risks—specifically, no firewall was in place, or it was improperly configured, which is just as bad.

The prudent person's responsibility for security here is to use reasonable care. You can find a more detailed definition in Prosser, Wade, and Schwartz's *Cases and Materials on Torts*: "...requiring the actor to conform to a certain standard of conduct, for the protection of others against unreasonable risks." Who says Hollywood liberalism doesn't contribute to society?

- **Lost data:** You have probably heard the stories of companies that lost all their business data in hurricanes such as Katrina or the September 11 attacks, and many companies did not recover. What if your company experienced the same loss of data because you did not have a firewall and an attacker deleted your data because he could? What would happen to your business? Would it cost money to re-create everything? Would you suffer lost sales? Would you still be employed the next day?
- **Compromise confidential data:** Every organization has data it considers confidential and, if lost, might cause financial problems, legal difficulties, or extreme embarrassment. These things might be caused by the loss of customer information such as credit card numbers, secret plans for the new weight loss formula, or secret product plans that end up in the hands of a competitor. The list goes on, and when you have been hacked, you must assume the worst. Perhaps this is why most cybercrimes go unreported—it is embarrassing, and admitting to being hacked is a sign of weakness that could affect the reputation and brand of a company.
- **Network downtime:** Have you ever gone to an ATM machine or a grocery store to get cash and paid with your cash card in the swipe card readers? The networks enabling these devices to operate usually work fine; however, if they were not protected, an attacker might cause them to go down. The loss of revenue from these networks can quickly grow if they are unavailable. Downtime is the bane of any network, and a cost is always associated with these types of events.

Ultimately, everyone has something worth protecting, and failure to do so is ill-advised; it is just a matter of time before something happens. The next question is, "*What does a firewall do to protect my network?*"

What Does a Firewall Do?

A firewall examines traffic as it enters one of its interfaces and applies *rules* to the traffic—in essence, permitting or denying the traffic based on these rules. Figure 7-1 shows a firewall filters both inbound and outbound traffic.

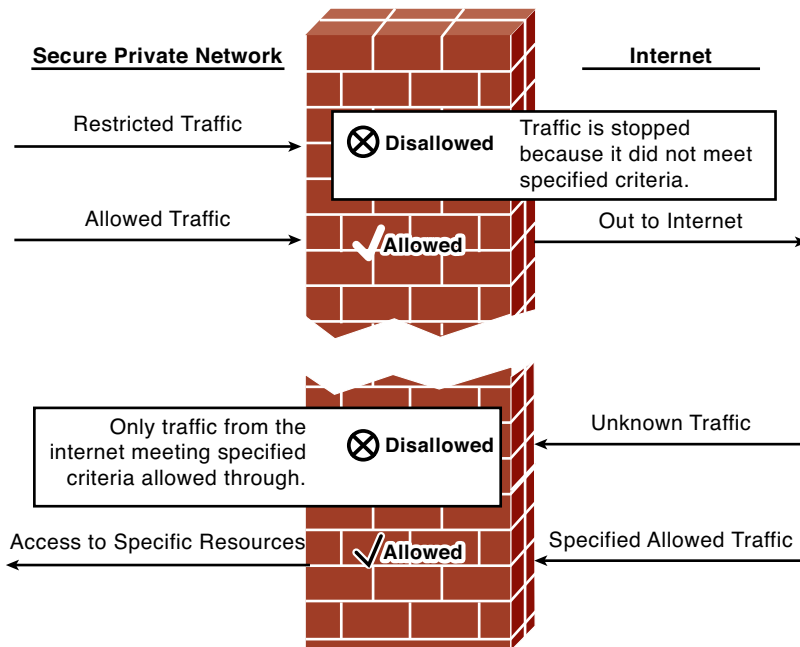


Figure 7-1 Firewall in Operation

Firewalls use access control lists (ACLs) to filter traffic based on source/destination IP addresses, protocol, and the *state* of a connection. In other words, normally you might not allow FTP/21 *into* your network (via the firewall), but if a user inside your network begins an FTP session out to the Internet, it is allowed because the session was *established* from inside the network. By default, firewalls trust all connections to the Internet (outside) from the trusted internal network (inside).

A firewall can also log connection attempts with certain rules that might also issue an alarm if they occur. Finally, firewalls enable you to perform Network Address Translation (NAT) from internal private IP addresses to public IP addresses. The section “Firewall Operational Overview” discusses the roles of a firewall; however, here you can tie the firewalls back to Chapter 2’s security policy discussions by examining how a firewall enforces your security policy.

Firewalls Are “The Security Policy”

What kind of traffic is allowed into or out of your network? How do you secure your network against attacks? What is your security policy? What happens to the people who do not follow the security policy? Who is responsible for writing and updating the security policy?

All these questions are valid, and they all deserve answers. Having a network that connects to the Internet via a firewall is only the first step to security; because this book is

about first steps, this would be a perfect place to start. You should now know that the security policies form the basis of how firewall rules are determined and then implemented into a production network.

Do you remember the old saying, “No job is ever finished until the paperwork is done?” Well, no security solution is complete until you establish a written narrative of the rules and regulations that govern your organization’s security posture. This written version of your security rules and regulations is known as a *security policy*. Now, this policy document is different in nature and scope than a security plan, so be sure that you understand what makes a policy unique from every other security document an organization maintains. And just what is it that makes a security policy different from a security plan? Drum-roll please....

PUNISHMENT! That is correct; a security policy includes what is permissible and what will happen to you if you do not live by the law of the land. If you do not follow the rules, you can be

- Fired or dismissed
- Demoted
- Demoted and fined
- Fired, dismissed, and demoted
- Demoted, dismissed, and even punked!
- All the above

All kidding aside, the security policy document spells out in clear language exactly what the regulations and expectations are, who enforces them, and what happens to you if you break them. A security policy is all about the consequences of user actions coupled with audit in the form of AAA usually.

Having said that, how can a firewall be the security policy? Simple—a firewall does what it does by following the rules configured by a network engineer or information security officer (ISO). These rules should perfectly align with a written narrative version found in the security policy document you have on your shelf, next to the box of CDs at the back of the server room or sitting useless in some manager’s office. Grab that old dusty binder and check it out. You should see that the security policy document contains information and a listing of the network rules (refer to Chapter 2). The interesting thing is that all the rules in the policy document form the basis of what you must configure on the firewall.

Note Wait a minute! We have a hand in the front row. Yes...you with the confused look on your face. Your question is, “Why is the binder that contains the security policy so dusty and located in such an obscure place?” As strange as that might sound, go ahead and put your hand down. I will tell you the answer to that question is that most organizations either do not have a security policy set, or the set that they have is so old that it was written during a previous presidential administration.

The configuration rules entered on a firewall should perfectly align with the rules outlined in an organization's security policy. If you were to examine the firewall's configuration file, you might see something like Example 7-1, which is a portion of a Cisco Adaptive Security Appliance (ASA) configuration.

Example 7-1 *Sample Cisco ASA Firewall Rules*

```
access-list OUTSIDE extended permit tcp any object-group HTTPS-SERVERS eq https
access-list OUTSIDE extended permit tcp any object-group WEB-SERVERS eq www
access-list OUTSIDE extended deny ip host 90.84.x.x any
access-list OUTSIDE extended permit icmp any any time-exceeded
access-list OUTSIDE extended permit icmp any any unreachable
access-list OUTSIDE extended permit icmp any any echo-reply
access-list OUTSIDE extended permit tcp any host 12.238.x.x eq ftp
access-list OUTSIDE extended permit tcp any host 12.238.x.x eq ftp-data
```

The **access-list permit** statements in Example 7-1 are most likely in keeping with some security policy statement that dictates what services are allowed, by name, to enter the protected network and the destinations to which those services are allowed to access. Specifically, this example shows the customer having web servers (www-80), secure web servers (https-443), and an FTP-21 server. These permit entries in your firewall's configuration are your network's security plan, and the security policy defines what they are and why they are present.

To expand on the firewall to security policy analogy, examine some additional security policy bullet points and how a firewall aligns with them:

- A security policy outlines what action will be taken in response to circumstances that arise.
- A security policy document is constantly evolving and changing to meet new security needs.
- A security policy dictates both acceptable and unacceptable usage parameters.

If you perform a point-by-point comparison of a security policy with a firewall configuration, you see that firewalls act with a written security policy document, as shown in Table 7-1.

Table 7-1 *Comparing Security Policies and Firewall Configurations*

	Security Policy	Firewall Configuration
Ability to respond to circumstances	Yes	Yes
Constantly evolving	Yes	Yes
Dictates behavior	Yes	Yes

The intention of this section is not to convince you that a firewall is a replacement for a security policy document, but to get you thinking about security as an all-encompassing philosophy of plans, policies, and security devices. You must put a great deal of thought into a complete solution—not simply rely on a single aspect to protect your network. When you are ready to plan your firewall's configuration and develop the rules permitting or denying traffic, you should use your security policy as the starting point. Firewalls are the physical and logical manifestations of your security policy.

We Do Not Have a Security Policy

The reality is that not every company has a security policy set (yet), and although it is important, you can still secure your network without one. Presume that you have a firewall already in place and functional. The best advice is to slowly start the process of implementing security in your network. This means carefully reviewing the business needs (very important) of each rule that you currently have in your firewall and writing down each need. Documenting *why* something was done will be helpful later if there is a security incident or when the network changes, providing justification on removing the entry. Certainly this advice is also true for anything new that needs to be accessed; you can plan on new things given the ever-forward marching of technology. If this book helps you keep your business and family safer, you have done something to be proud of...now go write those security policies!

Firewall Operational Overview

Every long journey begins with the first step. Before delving too deeply into other areas of security appliance behavior, it is essential to understand how a firewall performs its magic.

Most firewalls (most, not all) rely on Stateful Packet Inspection (SPI) to keep track of all outbound packets and the responses these packets might generate. Keeping track of the hosts on the protected network that are generating outbound packets keeps rogue or unsolicited WAN packets from entering an external interface.

In other words, a firewall that uses SPI, as discussed in Chapter 5, “Overview of Security Technologies,” watches all traffic that originates from an inside host, tracks the conversation from that host to the desired destination, and ensures that the inbound response to that request makes it back to the host that started the whole thing in the first place.

Note A firewall that is not stateful in design and configuration is incomplete and should not be used to protect your network. The importance of the stateful tracking of connections is critical to the security of any network. This chapter focuses on firewalls that track the state of a connection. As a reference point, all Cisco ASA and PIX firewalls are considered stateful packet inspection firewalls.

The critical dual purposes of *packet inspection and filtering* (blocking) of packets is one of the most fundamental responsibilities of a firewall. The following list includes the most common rules and features of firewalls:

- **Filter *incoming* network traffic based on source or destination:** Blocking unwanted incoming traffic is the most common feature of a firewall and is the main reason for a firewall—stopping unwanted traffic from entering your network. This unwanted traffic is usually from attackers, thus the need to keep it out.
- **Filter *outgoing* network traffic based on source or destination:** Many firewalls can also screen network traffic from your internal network to the Internet. For example, you might want to prevent employees from accessing inappropriate websites. You might also place a firewall between your network and a business partner with rules to keep each of you safe.
- **Filter network traffic based on content:** More advanced firewalls can screen network traffic for unacceptable content. For example, a firewall integrated with a virus scanner can prevent files that contain viruses from entering your network. Other firewalls integrate with email services to screen out unacceptable email.
- **Detect and filter malware:** The rise and proliferation of botnets and malware have driven firewall manufacturers to implement features designed to detect infected hosts through packet inspections. This is a good example of how security is ever changing and the security of the network must continue to advance as well because what was secure yesterday might not be tomorrow.
- **Make internal resources available:** Although the primary purpose of a firewall is to prevent unwanted network traffic from passing through it, you can also configure many firewalls to enable selective access to internal resources, such as a public web server, while still preventing other access from the Internet to your internal network. In many cases, you can accomplish this by using a DMZ, which is where the public web server would be located. (DMZs are discussed later in the section “Essentials First: Life in the DMZ.”)
- **Allow connections to internal network:** A common method for employees to connect to a network is using virtual private networks (VPN). VPNs enable secure connections from the Internet to a corporate network. For example, telecommuters and traveling employees can use a VPN to connect to the corporate network. VPNs can also connect branch offices to each other over the Internet, saving on WAN costs.
- **Report on network traffic and firewall activities:** When screening network traffic to and from the Internet, you need to know what your firewall is doing, who tried to break in to your network, and who tried to access inappropriate material on the Internet. Most firewalls include a reporting mechanism of some kind. A good firewall can also log activity to a syslog or other type of archival storage receptacle. Perusing firewall logs after an attack occurs is one of a number of forensic tools you have at your disposal.

Firewalls in Action

These might be new concepts for you, and hopefully you are not thoroughly confused at this point. Look at Figure 7-2 for a bit more clarity of this process. Please refer to the list, which explains the steps a bit more in depth.

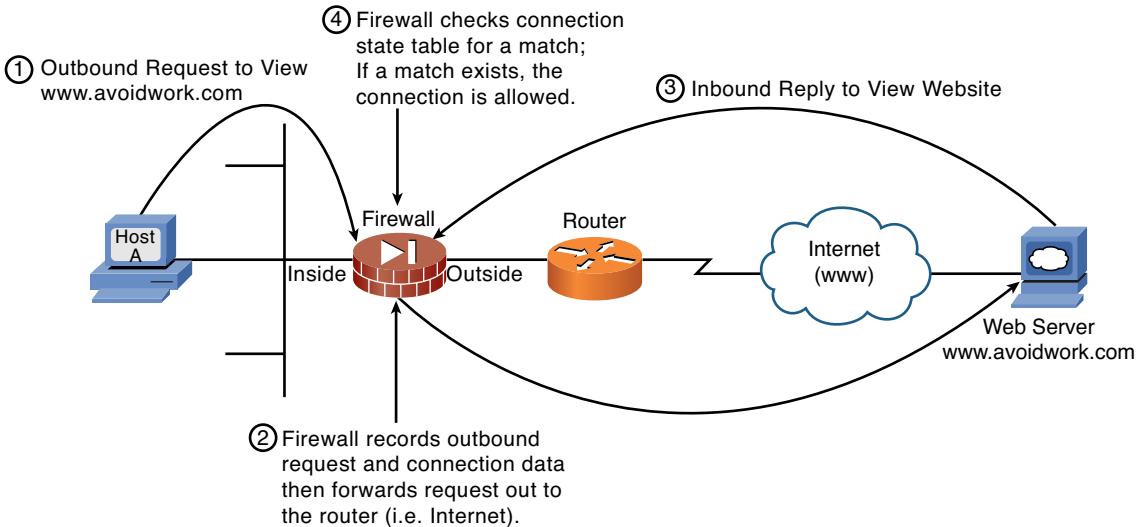


Figure 7-2 Firewall in Operation

Before looking at the list of steps, you need to know that many firewalls have only two physical interfaces, and 99 percent of them are based on Ethernet. These interfaces are called *inside* (protected) and *outside* (unprotected) and are deployed in relation to *your network*; some have DMZ interfaces as well. Thus, in practice, the outside interface connects to the Internet and the inside interface connects to your internal network:

Figure 7-2 shows a high-level view of the following:

1. Host A is an Apple Macbook Pro that opens a web browser and wants to view a web page from the `www.avoidwork.com` web server. This action causes Host A to send the request to view this web page out through the firewall across the Internet and to the web server.
2. The firewall sees the request originated with Host A and is destined for `www.avoidwork.com`.
 - a. The firewall records (tracks) the outbound request and expects that the reply will come only from the `www.avoidwork.com` web server.
 - b. A session marker is placed in the firewall's session state table that tracks the communication process from start to finish.

- c. Connection metrics, such as time opened and so forth, are also placed with the marker in the session state table record maintained by the firewall for this conversation.
3. The Avoidwork.com web server replies to the web page request from Host A, which is then transmitted back through the Internet and to the firewall.
 4. The firewall checks its session state table to see whether the metrics being maintained for this session match the outbound connection. If all the stored connection details match exactly, the firewall enables the inbound traffic.

The information contained in the firewall's state table records and tracks information such as who needed www information from the avoidwork.com server, when they asked for it, how they asked for it, and so forth. This provides an added level of protection over and above the "can I enter or not" rules because if a certain traffic type is allowed in but the host did not ask for it (attack), it's denied. Because a firewall maintains connection state information about inbound and outbound connections, the possibility of a hacker "spoofing" or "forging" a packet with the intention of penetrating your network becomes more difficult. When attackers try to send packets to get through a firewall, incorrect or missing connection state information means that the session is terminated and most likely logged for later review.

Note Many firewalls examine the source IP addresses of packets to determine whether they are legitimate. An attacker would conduct an IP spoofing attack to try to gain entry by spoofing the source IP address of the packets sent to the firewall. If the firewall thinks the packets originated from a trusted host because they had the correct source IP address, the firewall might let the packets through unless other criteria fails to be met. This reinforces the principle that technology alone does not solve all security problems. In addition, you need the involvement of your company's management and, you guessed it, a security policy. Cisco firewalls use an adaptive security algorithm as a method of dynamically appending a random number to the translated session to make it even more difficult for a hacker to intercept.

Implementing a Firewall

The choice of firewalls is almost mind-boggling these days; they come in every shape, size, and capacity. When I am designing a firewall solution for a customer, the first thing I want to know is what will the firewall's responsibilities be?

The type of firewall you install depends on your exact requirements for protection and management, and the size of your network, or what is to be protected by the firewall. Firewalls usually fall into one of the following categories:

- **Personal firewall:** A personal firewall is usually a piece of software installed on a single PC to protect only that PC. These types of firewalls are usually deployed on home PCs with broadband connections or remote employees. Of course, any time

someone wants to deploy a firewall, it is a good idea. You can find some of the more well-known personal firewalls at these websites:

www.zonealarm.com

www.firewallguide.com

Operating system manufacturers such as Apple and Microsoft have responded to this need by integrating personal firewalls within them. Apple's OS X comes with an IP firewall and Windows has a similar firewall, it is just not as secure as the one in OS X. Most antivirus companies have expanded their products to include all sorts of protection through the use of their product suites.

- **All-in-one firewall/routers:** These kinds of firewalls are widely used by broadband (cable or DSL) subscribers who have the benefit of a single device that offers the following features and functionality: router, Ethernet switch, wireless access point, and a firewall. If this type of firewall appeals to you, ensure that you take care to determine the firewall's capabilities, and be skeptical of the security you can gain from these devices, regardless of who makes them. **WARNING:** Do not be tricked into assuming that a home router has a good firewall built into it; do your research first. I especially advise people to check on how the manufacturer supports what it makes; for example, if it does not take phone calls, you might want to continue shopping.
- **Small-to-medium office firewalls:** These firewalls, such as the Cisco ASA 5505 and 5510 or the older PIX 501 and 506, are designed to provide security and protection for small office home office (SOHO) types of requirements. In most cases, they have expansion slots allowing for additional network connections or advanced feature cards to be installed.
- **Enterprise firewalls:** These firewalls, such as the Cisco ASA 5520 and up, are designed for larger organizations with thousands of users. These larger models are needed when there are demands for larger numbers of connections, capacity, and features. As a result, they have additional features and capacity, such as more memory and extra interfaces along with slots for advanced feature cards to be added. An example in some cases would be an IPS module.

Normally, a firewall is installed where your internal network connects to the Internet. Although larger organizations also place firewalls between different parts of their internal network that require different levels of security, most firewalls are placed to screen traffic passing between an internal network and the Internet. For example, if a large organization enables business partners to connect directly to its network, you typically find a firewall controlling what is allowed into its network from the partners. This placement of an internal firewall is definitely considered best practice.

Note No matter what type of firewall you choose, you must define the traffic filters that will support your security policy. Cisco firewalls all run the same version of an operating

system that has the same reporting and management capabilities, regardless of the model, which is helpful when administering them.

Determine the Inbound Access Policy

As network traffic passes through a firewall, the traffic is subject to the rules defined within the firewall. Because 99 percent of all networks use private IP addresses on the inside of their networks, you can expect almost every firewall to be using Network Address Translation (NAT)—as discussed in Chapter 5.

Note Packets coming in from the Internet in response to requests from local PCs (users) are addressed to the firewall's outside interface. The firewall is likely using NAT and tracking the state of each inside user request. The firewall is dynamically allocating port numbers on the outside interface using NAT. Thus, allowing multiple users to use a public IP address so their requests can be routed on the Internet is the essence of NAT. The use of a single IP address and port numbers to translate addresses is known as *port address translation (PAT)*. These port changes are also rapidly made, making it difficult for an attacker to make assumptions about which port numbers to use.

If all your LAN traffic were destined for the Internet, the inbound access policy would be straightforward in its design. The firewall permits only inbound traffic in response to requests from hosts on the internal LAN. The firewall tracks all outbound requests in its state table, as previously discussed.

However, there will come a time when specific requests from the outside must be allowed and controlled through the firewall. Notice that we did *not* say that this was a good idea or that you should do it, we are just acknowledging that it's a business function that a security professional must support.

Note The realities of the real world make companies want to have their own email or web servers without spending money on a new firewall that has a DMZ interface, which is where you place these servers whenever possible. The section “Essentials First: Life in the DMZ” discusses the purpose and role of a DMZ interface.

Allowing direct access from the Internet (outside) through your firewall is perilous but common practice. The key to security in these types of implementations is to strictly define the traffic types you will allow and the port number. For example, permitting IP to any location inside your network is inappropriate. For example, you should permit only inbound traffic from the Internet HTTP (port 80) traffic to your web server (IP address: 10.10.10.10). Allowing only HTTP (port 80) traffic to the web server from the Internet is much smarter than allowing every kind of TCP/IP protocol and port.

A strongly recommended best practice is to add layers of security in the form of a personal firewall, intrusion detection system (IDS), and antivirus software. Also, before you implement these devices as layers, make sure your security policies outline the best practices and what steps are needed to maintain security. A layered security model should be used to protect your network; the more layers, the harder it is for an attacker to penetrate your network. The use of layers is sort of like the joke told between hunters. When you see a hungry and angry bear in the woods start to charge you, as you begin to run remember you do not have to be faster than the bear, just faster than the other hunter! Layering network security definitely helps make your network less appealing than your competitors. Another layer would be to integrate an IPS in a firewall, making a layered defense.

Determine Outbound Access Policy

All firewalls screen traffic coming into a firewall from the Internet, but a well-implemented and designed firewall also screens outgoing user traffic. Spoiled employees are not going to like this, but the truth of the matter is that companies pay for Internet connections in support of their business, NOT to let employees surf, watch video, stream music, or look at pictures they are not supposed to.

You might also want to use your firewall to control what IP addresses are allowed to exit; specifically, you should allow only IP addresses that are found on your internal network out, thus preventing spoofing of IP addresses.

Perhaps there are also certain places on the Internet where you do not want users to go. Alternatively, you might want to specify the locations they are allowed to go because every other destination will be denied by default. Recall the earlier discussion of proxy servers and how they can be used to control and monitor traffic that leaves your network. They are a good example of a device that defines an outbound access policy. Remember, employees and contractors are bound to rules, whether they be policies or service-level agreements (SLA), and good behavior is not optional—it's mandatory—and so are accurate logging and event correlation.

In addition, recall the earlier discussion about placing a firewall between your network and connections to business partners. This type of firewall usage and placement is also where you would apply and control traffic bound from your network to theirs. The next section looks at the next aspect of firewall and network security: the Demilitarized Zone (DMZ).

Essentials First: Life in the DMZ

The *Demilitarized Zone (DMZ)* is a term used in the military to define a buffer area between two enemies. Perhaps the most commonly acknowledged DMZ in the world is the DMZ between North Korea and South Korea, which separates them because they have not yet signed a permanent peace treaty since the Korean War. Perhaps this is an interesting piece of military and political trivia that you did not know, but how does it relate to securing your network and firewalls?

If your company has a self-hosted public website complete with email servers, you might consider using a two-interface (inside and outside) firewall and have the firewall create translation rules that direct the inbound traffic to the correct servers on your private network. Although this might seem like a safe thing to do, it could be disastrous if a talented hacker sets his sights on you. Connecting web, mail, and FTP servers located on the inside of your network to the Internet can be dangerous and, in some cases, simply not recommended. Secure FTP is also an option but the same rules apply.

Well, some smart people got together a long time ago and said, “Hey—let’s put a third interface on the firewall and call it a DMZ.” Sending traffic from the Internet inbound directly to your private network is a *bad* idea. Adding the third interface to a standard firewall made things both easier and quite a bit safer when deploying Internet accessible servers and services (www, email, and so on). If you were going to sell computers out of your house, you would not want people coming inside your house to buy one, would you? Of course not; you would want to set up a little shop in the garage or on the front porch, thus preventing people that you do not know from wandering all over your house and tampering with your comic book collection or going into your fridge to make a sandwich.

A DMZ is an interface that sits between a trusted network segment (your company’s network) and an untrusted network segment (the Internet), providing physical isolation between the two networks enforced by a series of connectivity rules within the firewall. The physical isolation aspect of a DMZ is important because it enables Internet access only to the servers isolated on the DMZ and not directly into your internal network, as shown in Figure 7-3.

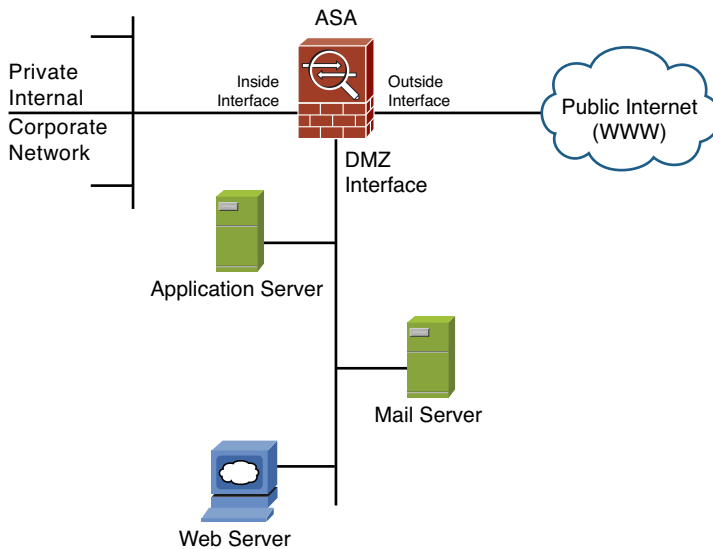


Figure 7-3 *DMZ Placement and Function*

In Figure 7-3, the segment connected to the DMZ interface contains the mail, web, and application servers. Rules applied to the DMZ interface prevent traffic from the Internet from going beyond the segment attached to it.

The biggest benefit to a DMZ is in isolating all unknown Internet requests to the servers on the DMZ and no longer allowing them into your internal network. However, some additional benefits to deploying a firewall with a DMZ can help you better understand what happens in your network and thereby increases security:

- Auditing DMZ traffic
- Locating an IDS on the DMZ
- Limiting routing updates between three interfaces
- Locating DNS on the DMZ

This section discussed what a DMZ is and provided a general example of how to use one. The following case studies examine a requirement for a DMZ and why you should use one in a network given a specific set of criteria.

Case Studies

This chapter presented several interesting aspects of how firewalls operate and how they can be deployed in networks. The introduction of this information needs to be reinforced with some real-world case studies that provide some answers to questions you might still have and clarify the important aspects of what has already been covered.

Case Study: To DMZ or Not to DMZ?

Carpathian Corporation has grown and is in need of increased security and additional capacity in the form of a new firewall; this time it wants to use a dedicated DMZ. If the Carpathian Corporation wants to continue with its proposed plan for self-hosting, it needs to consider the security-related issues relevant to the suggested DMZ solution. It is taking the right steps by asking what security ramifications should be addressed prior to making the purchase. The Carpathian IT staff needs to take a good look at the risk factors involved with providing for its own Internet services (web servers) and where the pitfalls might occur:

- **Question/Security Issue #1:** Can Internet traffic travel to servers on the private network, or is there another solution?

Answer: The web and mail servers will be attached to the DMZ segment. They will not be dual homed or have conflicts of security in its implementation because they will be physically separated from inside hosts.

- **Question/Security Issue #2:** How can the IT staff ensure that inbound network traffic will stay confined to the segment containing the web and mail servers?

Answer: The DMZ interface rule set will not allow external traffic to reach the private network, by nature of configured connectivity rules. This will keep the inbound Internet traffic confined to the DMZ segment only.

- **Question/Security Issue #3:** What measures can be taken to hide the private network from the inbound network traffic?

Answer: The DMZ interface will not have routes or dual-homed NIC cards that would normally enable this to occur.

The Carpathian IT staff is in the “If we self-host, we must use a DMZ” frame of mind. This frame of mind is correct, and that should be obvious at this point: Use a firewall with a DMZ interface—always!! A DMZ is another layer of security and defense for your network, as shown in Figure 7-4.




Figure 7-4 Firewall Deployment with Web Server in a DMZ

Cisco lists a variety of configuration settings when viewing their devices’ configuration files. Example 7-2 shows several configuration files for clarity purposes. To illustrate the case study, comments are made surrounding key configuration entries; however, not every command is discussed because that is beyond the scope of this book. You can find additional information at Cisco.com.

Example 7-2 Firewall with Self-Hosted Internal Web Server (No DMZ)

```
Cyberwall(config)# sh run
: Saved
ASA Version 8.5
```

```

!
hostname CyberWall
domain-name CarpathianCorp.com
enable password <ChangeMe> encrypted
passwd <ChangeMe> encrypted
names
!
!
interface Vlan1
  description SECURE INSIDE LAN [do not change]
  nameif INSIDE
  security-level 100
  ip address 192.168.0.1 255.255.255.0
!
interface Vlan2
  description OUTSIDE UPLINK TO SERVICE PROVIDER [do not change]
  nameif OUTSIDE
  security-level 0
  ip address 209.164.3.2 255.255.255.0
!
interface Vlan3
  description DMZ INTERFACE FOR INTERNET FACING SERVERS [alter with care]
  nameif DMZ
  security-level 50
  ip address 10.10.10.1 255.255.255.0
!
!... These commands name and set the security level for each vlan or interface, the
ASA 5505 uses vlans to assign inside and outside whereas all other models have
physical interfaces. Through these commands, the firewall knows which interface is
considered untrusted (outside), trusted (inside) and DMZ. Notice the numeric values in
this configuration example. Here we have the least secure interface outside assigned a
security value of 0, as it should be. The inside interface is considered secure, so it
has a value of 100, with the DMZ being somewhere in between at 50.
!
interface Ethernet0/0
  description OUTSIDE INTERFACE [do not change]
  switchport access vlan 2
!
interface Ethernet0/1
  description INTERFACE FOR THE DMZ WEB SERVER [do not change]
  switchport access vlan 3
!
interface Ethernet0/2
  description RESERVED FOR INTERNAL HOST [alter with care]
!

```

```

interface Ethernet0/3
  description RESERVED FOR INTERNAL HOST [alter with care]
!
interface Ethernet0/4
  description RESERVED FOR INTERNAL HOST [alter with care]
!
interface Ethernet0/5
  description RESERVED FOR INTERNAL HOST [alter with care]
!
interface Ethernet0/6
  description RESERVED FOR INTERNAL HOST [alter with care]
!
interface Ethernet0/7
  description RESERVED FOR INTERNAL HOST [alter with care]
!
! --- An access list is created called "OUTSIDE" allowing WWW (http) traffic from
anywhere on the Internet to the host at 10.10.10.212 (the web servers REAL IP address
on the DMZ). Add additional lines to this access list as required if there is a email
or DNS Server. This is the first step in creating a rule set that permits traffic into
our network if it is destined for a specific IP Address.
!
access-list OUTSIDE extended permit tcp any host 10.10.10.212 eq www
!
! --- For purposes of this example we are not going to add anything else. Any
additional entries needing to be placed in the access list must be specified here. If
the server in question is not WWW, replace the occurrences of WWW with SMTP, DNS,
POP3, or whatever else might be required, like the ability to ping the server from the
Internet.
!
logging enable
logging timestamp
!
<<<output omitted for brevity>>>
!
! --- The following NAT commands specify that any traffic originating inside from the
ASA on the 192.168.0.0 /24 network will be NAT'd (via PAT because of the dynamic
interface command) to the ASAs public IP address that is assigned to the OUTSIDE
interface.
!
! --- The ASA NAT rules changed completely the new way is to define the subnets you
wish to NAT using object groups, the next four lines we have defined them as needed
for the INSIDE corporate as well as the DMZ.

```

```

!
object network OBJ_NAT_CORP
  description inside "corporate" subnet that must have internet access
  subnet 192.168.0.0 255.255.255.0
!
object network OBJ_NAT_DMZ
  description DMZ subnet that must have internet access
  subnet 10.10.10.0 255.255.255.0
!
! --- Once the subnets are defined in an object group we assign the type of NAT we
wish to perform as well as the direction. In the following examples we are permitting
the INSIDE and DMZ subnets to access the Internet using PAT via the ASAs outside
interface IP Address for both. This is shown in the command NAT (source interface,
destination interface) dynamic interface. The dynamic keyword means PAT to the ASA.
One of my favorite ways to check if this is working after configuring it open a web
browser and go to www.ipchicken.com this website will tell you the public IP Address
you are coming which should be the ASAs outside IP Address. Yes I know it's a goofy
name but that's what makes it easy to remember plus it makes people smile when you
tell them it.
!
object network OBJ_NAT_CORP
  nat (INSIDE,OUTSIDE) dynamic interface
!
object network OBJ_NAT_DMZ
  nat (DMZ,OUTSIDE) dynamic interface
!
! --- The last remaining NAT we must perform is for the Internet accessible Web
server that is on our DMZ. Once again we create an object group but this time we
specify a single host, which is the real IP address of the web server.
!
object network OBJ_NAT_WEBSERVER
  description real ip address assigned on the web servers nic card
  host 10.10.10.212
!
! --- Now that the object group is created identifying the servers real IP Address
we assign a NAT in the same format as we previously did with the difference being
after the direction (inside,outside) we define this as a STATIC NAT and give the
public IP Address to use. In practice what will happen is as packets reach the ASA
if they pass the access-list the ASA will check what their destination IP Address
is. Should the destination address be 209.164.3.5 (web server public IP Address)
the ASA will NAT those packets to the real IP Address of the server of 10.10.10.212
and forward them to the server on the DMZ.

```

```

!
object network OBJ_NAT_WEBSERVER
  nat (INSIDE,OUTSIDE) static 209.164.3.5
!
!
access-group OUTSIDE in interface outside
!
! --- There is only one access list allowed per interface per direction (for example,
inbound from the Internet on the outside interface) as we have shown here.
!
route outside 0.0.0.0 0.0.0.0 209.164.3.1
!
! --- Set the default route to be via the WAN routers Ethernet interface
!
<<<output omitted for brevity>>>
!
dhcpd dns 192.168.0.10 192.168.0.11
dhcpd domain mydomain.com
dhcpd address 192.168.0.2-192.168.0.125 inside
dhcpd enable inside
!
<<<output omitted for brevity>>>
!
! --- The last major functionality of an ASA show in its configuration is that of the
"inspects". Generally an inspect statement in the following section represents a
protocol that the ASA will be taking extra steps on the packets the statement
represents. For example many attacks are based on altering DNS replies so the ASA has
been configured to inspect DNS packets to help protect your network. Two inspects that
might be of importance to you are "inspect esmtp" and "inspect sip", depending on your
email server configuration and version the presence of esmtp may cause user issues
with emails, try removing it if this occurs. Regarding SIP when NATing a SIP
connection to an internal voice gateway you will want this statement as it provides
functionality that enables NAT to be done correctly and SIP to work, gotcha is it
depends on the provider. Inspects are very helpful and can be adjusted to offer very
granular security, please see www.cisco.com for more information.
!
policy-map type inspect dns preset_dns_map
  parameters
    message-length maximum client auto
    message-length maximum 512

```

```

!
policy-map global_policy
class inspection_default
inspect dns preset_dns_map
inspect ftp
inspect h323 h225
inspect h323 ras
inspect rsh
inspect rtsp
inspect esmtp
inspect sqlnet
inspect skinny
inspect sunrpc
inspect xdmcp
inspect sip
inspect netbios
inspect tftp
inspect icmp
inspect icmp error
inspect ip-options
!
service-policy global_policy global
prompt hostname context
Cryptochecksum:88251e3c18c7d99dfa33f70b90228b63
: end
Cyberwall(config)#

```

Firewall Limitations

A firewall is a crucial component of securing your network and is designed to address the issues of data integrity or traffic authentication (via stateful packet inspection) and confidentiality of your internal network (via NAT). Your network gains these benefits from a firewall by receiving all transmitted traffic through the firewall. Your network gains these benefits from a firewall by receiving all transmitted traffic through the firewall. The importance of including a firewall in your security strategy is apparent; however, firewalls do have the following limitations:

- A firewall cannot prevent users or attackers with modems from dialing in to or out of the internal network, thus bypassing the firewall and its protection completely.
- Firewalls cannot enforce your password policy or prevent misuse of passwords. Your password policy is crucial in this area because it outlines acceptable conduct and sets the ramifications of noncompliance.

- Firewalls are ineffective against nontechnical security risks such as social engineering, as discussed in Chapter 1, “There Be Hackers Here.”
- Firewalls cannot stop internal users from accessing websites with malicious code, making user education critical.
- Firewalls cannot protect you from poor decisions.
- Firewalls cannot protect you when your security policy is too lax.

Note The FBI’s arrest of the phone master’s cracker ring brought several of these security issues to light. These hackers were accused of breaking into credit-reporting databases belonging to Equifax, Inc. and TRW, Inc. and the databases of Nexis/Lexis and Dun & Bradstreet. They also broke into many of the world’s providers. In doing so, these hackers did not use any high-tech attack methods. The phone masters used a combination of social engineering and dumpster diving, both techniques used by attackers that have little technical skill (refer to Chapter 1).

Chapter Summary

This chapter covered the world of firewalls and their role in securing a network. Not everyone believes in the value of these devices, and the discussions answered these naysayers and showed them the folly of their ways. Further proof of the importance of firewalls was provided by expanding on their pure technical aspects, while expressing the fundamental truth that firewalls are the manifestation of a company’s security policy.

One of the online resources that may assist you in determining the direction and policy of your network security is www.opengroup.org/jericho/about.htm. The Jericho Project was formed by a group of corporate security officers who saw the ever-decreasing security being driven by the concept of deperimeterization. In 2004, the Forum set out to drive and influence development of secure architectures, technology solutions, and implementation approaches, for the deperimeterizing IT world, to enable safe, secure collaborative interworking, globally between enterprises—business partners, customers, suppliers, and out-workers—and to encourage development of open standards that would underpin these solutions.

Operationally, this chapter covered how firewalls function, where and when to implement them, and how to design the access policies necessary to define access into your network. Furthermore, the chapter introduced the DMZ interface as an evolution in firewalls and how they provide special locations for various Internet servers. The chapter concluded with several brief case studies demonstrating firewalls in action, followed by some of their limitations.

Chapter Review Questions

The following questions assist in reinforcing the concepts covered in this chapter:

1. Who needs a firewall?
2. Why do I need a firewall?
3. Do I need a firewall?
4. How is a firewall an extension of a security policy?
5. What is the name of the table in a firewall that tracks connections?
6. What fundamental role does a DMZ fulfill in network security?
7. What are four benefits of a DMZ?
8. Can firewalls enforce password policies or prevent misuse of passwords by users?
9. Do firewalls guarantee that your network will be protected?
10. Are all firewalls created equal?

Router Security

*“Faith is being sure of what you hope for and certain of what you cannot see”
—Hebrews 11:1*

By the end of this chapter, you should know and be able to describe the following:

- The major components of Zone Based Firewall (ZFW) for routers
- The value of using the IOS-based intrusion detection functionality and the Cisco Firewall Feature Set (FFS)
- The breadth and scope of techniques used to secure your router to include a secure router template
- Securing your “routing” protocol: OSPF

Answering these key questions will enable you to understand the overall characteristics and importance of network security. By the time you finish this book, you will have a solid appreciation for network security, its issues, how it works, and why it is important.

Everyone is getting online as rapidly as possible in whatever way they can; if you are reading this book, you are probably the person your family calls to “fix” the Internet. Perhaps the best T-shirt I never bought was the one that read, “*No, I will not fix your computer*” from ThinkGeek.com, as shown in Figure 8-1. Think Geek is a website worth visiting; you can find funny and useful gear there.

The point is that most people do not understand that the Internet operates because of routers. They think that individuals have more control and security than they do because their PC connects to the Internet. Of course, this is not the case—there are no guarantees on the Internet, which is a wild and fast place. Of course, fast is relative; just as a reminder: Everyone does realize there are no guarantees on the Internet, which means it is slow at times and there is nothing anyone can do. Companies and especially ISPs try to do a good job, but unexpected events do occur.



Figure 8-1 *No, I Will Not Fix Your Computer*

As people and organizations seek to leverage the unparalleled possibilities of Internet communications, they need secure solutions that

- Protect internal networks from intrusion
- Provide secure Internet and remote access connections
- Enable network commerce through the World Wide Web

Today, the Internet is the focus of powerful, new technologies that dramatically enhance communications with remote customers, suppliers, partners, and employees. Users must be confident that network transactions—especially over public networks such as the Internet—are secure and sensitive information is protected.

Cisco IOS Software runs on more than 80 percent of Internet backbone routers and an equally high percentage of corporate network routers that connect to the Internet. Cisco IOS Software provides complete network services and enables networked applications. Cisco IOS security services offer many options for building custom security solutions for the Internet, intranet, and remote access networks to provide end-to-end network security.

A critical part of an overall security solution is a network firewall, which monitors traffic crossing network perimeters and imposes restrictions according to security policy. As discussed in Chapter 7, “Firewalls,” firewalls are not routers, and they connect the Internet to your corporate network. Routers that connect to the Internet are known as *edge routers*; they form the outermost perimeter of your network. In other words, they are the first layer of security.

Perimeter routers are found at any network boundary, such as between private networks, intranets, extranets, or the Internet. Firewalls most commonly separate internal (private) and external (public) networks.

The Cisco IOS Firewall Feature Set, available as a Cisco IOS Software option, provides an advanced security solution that protects networks from security violations. This integrated router security solution provides one element in a system of security solutions available from Cisco.

This chapter discusses the use of routers, the purpose of a firewall IOS, and what it is. Where within your network will you be applying this type of protection? This chapter explains the use and placement of this type of security technology and its advantages and disadvantages.

A *firewall* is a security device that sits on the edge of your Internet connection and functions as an Internet border security officer by constantly looking at all the traffic entering and exiting your connection, looking and waiting for traffic to block or reject in response to an established rule. The firewall plays the role of law and protection in a lawless global web, ever vigilant in its mission to protect the internal network resources that connect to it.

In contrast, the edge router provides connectivity between you and your service provider and this to the Internet for businesses. Most people view a router as a necessary device that provides them with connectivity. Having a router, however, means that it handles (routes) every single packet that wants to enter or leave the network. It is the role of the firewall to determine what is permitted or denied. However, if you have a router as the first layer into your network, shouldn't you use that router as part of your layered security strategy?

Of course you should. You have paid for the router and spent time configuring it; however, blindly trusting that it is inherently secure is a mistake. Even if your company spent tens of thousands of dollars on other security solutions, the router handling everything might not have had its configuration hardened to protect it and your network. The router is essentially in the default out-of-the-box (OOB) condition. Consider that if an attacker gained control of your router, he could rather easily shut down your entire network's capability to connect to the Internet. This means no email in or out, no e-commerce on your website, perhaps losing connectivity to critical business partners, and so on.

The perimeter router literally *sees* every single IP packet. What might the attacker learn? What might the attacker then be able to do? The router is a smart network device that holds a key position and handles crucial information. Network security is often thought of in terms of servers, firewalls, VPNs, and how to protect IT resources. This chapter covers how to protect any router and then expand its capabilities to further protect your network with an additional layer of security through the use of the Cisco Firewall Feature Set IOS. This specialized IOS provides greatly enhanced security features and functionality for the perimeter router. By securing the router and thus increasing your network's security, you can accomplish the following:

- Prevent routers from unintentionally leaking information about your network to attackers.

- Prevent the disabling of your routers (and thus your network) by attackers or accidental misconfiguration.
- Prevent the use of your routers as platforms to launch an internal attack or to be used to attack others.
- Reduce the load on the firewall and internal network as bad packets and thus stop associated attacks at the edge of your network.
- Quickly activate an additional layer of security to further protect your network.

These accomplishments revolve around the security and functionality of the router and your network. Not everyone wants to spend the money, time, effort, or expertise needed to correctly configure the firewall functionality on a router. The reality is that many companies enable the firewall to be the stateful packet inspection device and not the perimeter router. However, *everyone* should use a router as a layer in the defense of his network. The discussion and debate should center not on *if* but *how* the router should be configured. Following are three ways to configure your perimeter router:

- **Edge router with basic configuration:** Get the router, put a basic configuration in it, connect it to your LAN and the Internet, and you are finished. There is nothing fancy here, and absolutely no security or value to your network! Please don't do this, you're just begging for problems!
- **Edge router as a choke point:** As discussed in Chapter 5, "Overview of Security Technologies," all routers come with the capability to filter traffic based on access control lists (ACL). Access lists can be developed to filter traffic based on the packet type and destination at the perimeter router turning it into a prescreening layer of security. For example, if you host no web servers at your site, why would you ever allow HTTP requests? You wouldn't! Remember that when using ACLs, if the traffic is not permitted, it is implicitly denied! This is the minimum that should be accomplished! We recommend double-checking your access-list logic to ensure you don't inadvertently block legitimate HTTP traffic.
- **Edge router as a packet inspector:** To have the router perform more advanced filtering, this type of router is deployed with the firewall feature set on it. This router is the best of the three, and it is also the most difficult to achieve. Anything in life that is worth having is never free—you must work for it!

This chapter does not cover what a router secures or protects with a basic configuration because the answers to those questions change with every network. Instead, this chapter focuses on how a router functions as a layer of security in your network through the use of static access lists and as a screening device through more advanced access lists.

Edge Router as a Choke Point

A choke point came to the world of networking courtesy of the Internet's military heritage. A *choke point* refers to a single point at which everything will try to either enter or leave your network. A great example of this concept is the real history of the Battle of Thermopylae, which as portrayed in the movie *300* showed a narrow point, the pass of Thermopylae, aka The Hot Gates, being defended by a unified force of Greeks against an enormous Persian host.

In today's world of network security, the term choke point means the edge router. The edge router is the single point from which the entire Internet gains access to your network. The router then is also a single point of failure, but that is an entirely separate discussion.

Edge routers that operate as choke points increase your network's security by restricting the flow of data between your network and the Internet (or another network). A successful network security implementation of an edge router as a choke point is based on understanding what is happening in your network. This knowledge forms the basis of what should be used to filter network activity so that inappropriate activity can be identified. Network activity should be restricted to permit acceptable services only. Chokes provide a great way of implementing a coarse level of control and monitoring that can be fine-tuned using intelligent filters, such as proxy and stateful firewalls.

The value of edge routers configured as choke points is that they can prevent access to specific devices and applications in a performance-friendly way. This increase in security is typically provided through the use of standard and extended access control lists that can address traffic concerns at Layers 2, 3, and 4 of the OSI reference model. Because their performance does not *normally* suffer results from the fact that the router must read the contents of the IP packet anyway to make a decision on where to forward the packet. It does not take much more work to toss out the packet or permit it into the network.

The use of ACLs gives network engineers a high degree of control and filtering capabilities over packets traversing the router. Figure 8-2 demonstrates a common example of the rules and placement of an edge router. As a side note, can you figure out what network device is missing from this figure?

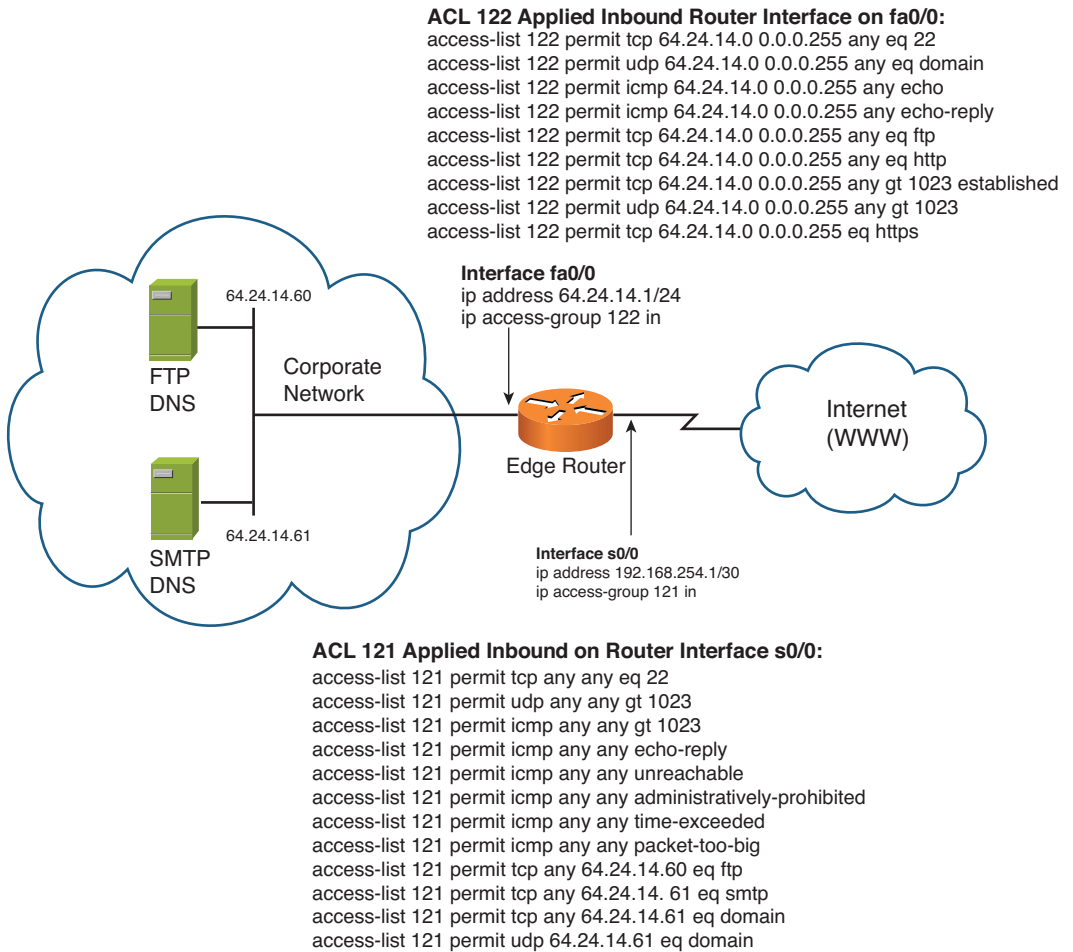


Figure 8-2 Edge Router as a Choke Point

This edge router acting as a choke point into the corporate network permits only the following traffic into the corporate LAN:

- Inbound mail delivery to the email (SMTP) server at IP address 64.24.14.61.
- FTP file transfers to the FTP server at IP address 64.24.14.60.
- DNS (zone transfers via UDP and name lookup requests via TCP) to the DNS server at IP address 64.24.14.61.
- TCP and UDP traffic above port 1023 to allow outbound connections from the private network to function.
- Only specific types of ICMP.
- All other traffic is denied access to the edge router.

As a user on the corporate LAN, the edge router permits only you to establish connections out to the Internet as follows:

- SSH (TCP port 22)
- DNS (UDP port 53)
- FTP (TCP ports 20 and 21)
- HTTP (TCP port 80)
- HTTPS (TCP port 443)

The use of a choke point router to limit access (both in and out) for known services (below port number 1023) leaves the network largely exposed. Because the majority of today's applications use ports above 1023 and not all IP stack and application implementations follow the 49152 through 65535 dynamic/private port guidelines, filtering above 1023 can affect the operation of applications that you want to function and, therefore, cannot deny this port range.

Did you discern what piece of equipment was missing in Figure 8-2? If you guessed a firewall performing stateful packet inspection (SPI) and network address translation (NAT), you were correct and are making progress toward becoming a security guru! But did you also notice that given the importance of email, you are not permitting it in your ACLs? Knowing whether a site needs or has email is important. Again, know the network's traffic needs before you begin implementing any ACLs.

Limitations of Choke Routers

Choke routers are useful and can protect your network as previously demonstrated; however, they are only part of the solution and are likely to stop only a script kiddie or someone who has already read this book and understands that the network is not completely protected. Remember, this security is another layer the attacker must overcome. Some of the limitations of choke routers are as follows:

- Choke routers running regular IOS cannot look at the higher layers of the OSI reference model (Layers 5–7). However, use of enhanced IOS facilitates the use of Network Based Application Recognition (NBAR), which enables a router to detect and block many of the more common worms. You can find a useful article regarding this point at <http://certcities.com/editorial/columns/story.asp?EditorialsID=76>.
- Choke routers do not adequately address protocol and application security concerns; you would have no idea whether your connections were being spoofed.
- Choke routers can have the capability to perform Stateful Packet Inspection (SPI) based on zones; however, this might require upgrades to support this feature, which is found in Cisco IOS 12.4(6)T and later.

Although choke routers do not address the preceding concerns, they are quite valuable for implementing broad network and service access policies (that is, what users on the Internet can access).

Caution Creating static ACLs require some thought and a lot of testing. A poorly written ACL can have adverse effects on the network in terms of performance and service availability. A strongly recommended practice is to write the ACL out on paper first to ensure that you have it designed to accomplish your filtering goals. Also, do not ask for help on your ACLs until you have repeated the mantra—there is an implicit deny all at the end of every ACL that the Cisco IOS does not display. The Cisco TAC will thank me for including that requirement because everyone forgets it—myself included!

Routers Running Zone Based Firewall

By now, you should see the value of prescreening traffic on your edge router and readily agree that using your edge router as a part of your layered security strategy will bring benefits to your network. Using the edge router as a choke point is certainly useful; however, there are some limitations to its use that might be important to you. Perhaps your company is involved in government contracts, so you must have the highest possible level of security. Or perhaps you work for the government. Regardless, the next level up in security is the use of Cisco Zone Based Firewall (ZFW) on the edge router.

Cisco IOS Software Release 12.4(6)T introduced ZFW, a new configuration model for the Cisco IOS Firewall feature set. This new configuration model offers intuitive policies for multiple-interface routers, increased granularity of firewall policy application, and a default deny-all policy that prohibits traffic between firewall security zones until an explicit policy is applied to allow desirable traffic.

Nearly all classic Cisco IOS Firewall features implemented before Cisco IOS Software Release 12.4(6)T are supported in the new zone-based policy inspection interface:

- Stateful packet inspection
- VRF-aware Cisco IOS Firewall
- URL filtering
- Denial-of-service (DoS) mitigation

Note Routers that perform firewall-like operations are sometimes referred to as *routerwalls*. Remember, the key term here is *firewall-like*, which means that it can sort of function as a firewall but should not be used to replace a bona fide, dedicated firewall appliance such as a Cisco ASA.

Cisco IOS Software Release 12.4(9)T added ZFW support for per-class session/connection and throughput limits and application inspection and control:

- HTTP
- Post Office Protocol (POP3), Internet Mail Access Protocol (IMAP), Simple Mail Transfer Protocol/Enhanced Simple Mail Transfer Protocol (SMTP/ESMTP)
- Sun Remote Procedure Call (RPC)
- Instant Messaging (IM) applications:
 - Microsoft Messenger
 - Yahoo! Messenger
 - AOL Instant Messenger
- Peer-to-Peer (P2P) File Sharing:
 - Bittorrent
 - KaZaA
 - Gnutella
 - eDonkey

Cisco IOS Software Release 12.4(11)T added statistics for easier DoS protection tuning.

Some Cisco IOS Classic Firewall features and capabilities are not yet supported in a ZFW in Cisco IOS Software Release 12.4(15)T:

- Authentication proxy
- Stateful firewall failover
- Unified firewall MIB
- IPv6 stateful inspection
- TCP out-of-order support

ZFW generally improves Cisco IOS performance for most firewall inspection activities. Neither Cisco IOS ZFW nor Classic Firewall includes stateful inspection support for multicast traffic.

Zone-Based Policy Overview

Cisco IOS Classic Firewall stateful inspection (formerly known as Context-Based Access Control, or CBAC) employed an interface-based configuration model, in which a stateful inspection policy was applied to an interface. All traffic passing through that interface received the same inspection policy. This configuration model limited the granularity of the

firewall policies and caused confusion of the proper application of firewall policies, particularly in scenarios when firewall policies must be applied between multiple interfaces.

Zone-Based Policy Firewall (also known as Zone Policy Firewall, or ZFW) changes the firewall configuration from the older interface-based model to a more flexible, more easily understood zone-based model. Interfaces are assigned to zones, and inspection policy is applied to traffic moving between the zones. Inter-zone policies offer considerable flexibility and granularity, so different inspection policies can be applied to multiple host groups connected to the same router interface.

Firewall policies are configured with the Cisco Policy Language (CPL), which employs a hierarchical structure to define inspection for network protocols and the groups of hosts to which the inspection will be applied.

Zone-Based Policy Configuration Model

ZFW completely changes the way you configure a Cisco IOS Firewall inspection, compared to the Cisco IOS Classic Firewall.

The first major change to the firewall configuration is the introduction of zone-based configuration. Cisco IOS Firewall is the first Cisco IOS Software threat defense feature to implement a zone configuration model. Other features might adopt the zone model over time. Cisco IOS Classic Firewall stateful inspection (or CBAC) interface-based configuration model that employs the `ip inspect` command set is maintained for a period of time. However, few, if any, new features are configurable with the classical command-line interface (CLI). ZFW does not use the stateful inspection or CBAC commands. The two configuration models can be used concurrently on routers but not combined on interfaces. An interface cannot be configured as a security zone member and be configured for `ip inspect` simultaneously.

Zones establish the security borders of your network. A zone defines a boundary where traffic is subjected to policy restrictions as it crosses to another region of your network. ZFW's default policy between zones is deny all. If no policy is explicitly configured, all traffic moving between zones is blocked. This is a significant departure from the stateful inspection's model, where traffic was implicitly allowed until explicitly blocked with an ACL.

The second major change is the introduction of a new configuration policy language known as CPL. Users familiar with the Cisco IOS Software Modular quality-of-service (QoS) CLI (MQC) might recognize that the format is similar to QoS's use of class maps to specify which traffic will be affected by the action applied in a policy map.

Rules for Applying Zone-Based Policy Firewall

Router network interfaces' membership in zones is subject to several rules that govern interface behavior, as is the traffic moving between zone member interfaces:

- A zone must be configured before interfaces can be assigned to the zone.
- An interface can be assigned to only one security zone.

- All traffic to and from a given interface is implicitly blocked when the interface is assigned to a zone, except traffic to and from other interfaces in the same zone and traffic to any interface on the router.
- Traffic is implicitly allowed to flow by default among interfaces that are members of the same zone.
- To permit traffic to and from a zone member interface, a policy enabling or inspecting traffic must be configured between that zone and any other zone.
- The self zone is the only exception to the default deny all policy. All traffic to any router interface is allowed until traffic is explicitly denied.
- Traffic cannot flow between a zone member interface and any interface that is not a zone member. Pass, inspect, and drop actions can be applied only between two zones.
- Interfaces that have not been assigned to a zone function as classical router ports and might still use classical stateful inspection/CBAC configuration.
- If it is required that an interface on the box not be part of the zoning/firewall policy, it might still be necessary to put that interface in a zone and configure a pass all policy (sort of a dummy policy) between that zone and any other zone to which traffic flow is wanted.
- From the preceding it follows that, if traffic is to flow among all the interfaces in a router, all the interfaces must be part of the zoning model (each interface must be a member of one zone or another).
- The only exception to the preceding deny by default approach is the traffic to and from the router, which will be permitted by default. An explicit policy can be configured to restrict such traffic.

Designing Zone-Based Policy Network Security

A security zone should be configured for each region of relative security within the network so that all interfaces assigned to the same zone will be protected with a similar level of security. For example, consider an access router with three interfaces:

- One interface connected to the public Internet
- One interface connected to a private LAN that must not be accessible from the public Internet
- One interface connected to an Internet service demilitarized zone (DMZ), where a web server, Domain Name System (DNS) server, and email server must be accessible to the public Internet

Each interface in this network will be assigned to its own zone, although you might want to allow varied access from the public Internet to specific hosts in the DMZ and varied

application use policies for hosts in the protected LAN. This concept is demonstrated in Figure 8-3.




Figure 8-3 *Basic Zone Firewall Topology*

In this example, each zone holds only one interface. If an additional interface is added to the private zone, the hosts connected to the new interface in the zone can pass traffic to all hosts on the existing interface in the same zone. In addition, the hosts' traffic to hosts in other zones is similarly affected by existing policies.

Typically, the example network will have three main policies:

- Private zone connectivity to the Internet
- Private zone connectivity to DMZ hosts
- Internet zone connectivity to DMZ hosts

Because the DMZ is exposed to the public Internet, the DMZ hosts might be subjected to unwanted activity from malicious individuals who might succeed at compromising one or more DMZ hosts. If no access policy is provided for DMZ hosts to reach either private zone hosts or Internet zone hosts, the individuals who compromised the DMZ hosts cannot use the DMZ hosts to carry out further attacks against private or Internet hosts. ZFW imposes a prohibitive default security posture. Therefore, unless the DMZ hosts are specifically provided access to other networks, other networks are safeguarded against any connections from the DMZ hosts. Similarly, no access is provided for Internet hosts to access the private zone hosts, so private zone hosts are safe from unwanted access by Internet hosts.

Using IPsec VPN with Zone-Based Policy Firewall

Recent enhancements to IPsec VPN simplify firewall policy configuration for VPN connectivity. IPsec Virtual Tunnel Interface (VTI) and GRE+IPsec enable the confinement of VPN site-to-site and client connections to a specific security zone by placing the tunnel

interfaces in a specified security zone. Connections can be isolated in a VPN DMZ if connectivity must be limited by a specific policy. Or, if VPN connectivity is implicitly trusted, VPN connectivity can be placed in the same security zone as the trusted inside network.

If a non-VTI IPsec is applied, VPN connectivity firewall policy requires close scrutiny to maintain security. The zone policy must specifically enable access by an IP address for remote sites' hosts or VPN clients if secure hosts are in a different zone than the VPN client's encrypted connection to the router. If the access policy is not properly configured, hosts that should be protected can end up exposed to unwanted, potentially hostile hosts.

Intrusion Detection with Cisco IOS

The Cisco IOS Firewall IDS acts as an inline *intrusion detection* sensor, watching packets and communication sessions as they flow through the router and scanning each packet to see whether it matches any of the IDS signatures.

Cisco developed its Cisco IOS Software–based intrusion detection capabilities in the Cisco IOS Firewall Feature Set with flexibility in mind so that individual attack signatures could be disabled in case of false positives. Also, although it is preferable to enable both the firewall and intrusion detection features of the FFS CBAC security engine to support a network security policy, each of these features can be enabled independently and on different router interfaces.

The Cisco IOS Firewall Feature Set includes intrusion detection technology in addition to basic firewall functionality. The Cisco FFS IOS acts as a *limited* inline intrusion detection sensor, watching packets and sessions as they flow through the router. (This is the inline aspect of its operation—scanning each packet to determine whether the contents match any of the IDS signatures it knows about.) When the router detects suspicious activity—in other words, when it believes that a packet contains an attack signature—it responds accordingly before network security can be compromised and logs the suspicious activity by using syslog and by communicating directly with a server running the Cisco Secure IDS Software.

Note System Message Logging (syslog) provides a means for the system and its running processes to report various types of system state information. There are three classes of system state data: error, informational, and debug. Cisco IOS Software provides an extensive system message and error reporting facility. IOS uses more than 500 service identifiers known as facilities to categorize system state data for error and event message reporting. System logging data is an important resource in diagnosing problems in general and, when issued by the firewall feature set, it enables the reporting of events.

Cisco routers running IDS functionality all have the signatures of attacks; these signatures are the reference to which the IDS will compare packets to determine whether there is an attack. It is critical that these signatures be as accurate and up to date as possible. Starting with Cisco IOS 12(4)11T and later, signatures are separated from the IOS version. This

means the signatures and the IOS can be upgraded (or updated) independent of each other. Furthermore, the IDS signatures found on the routers are the same as those on Cisco IDS appliance. These are huge improvements and are definitely recommended should you run older IOS where the IOS was tied to the signatures and vice versa.

In practice, this means that the engineers responsible for your network's security must ensure that the attack signatures are always as current as possible. We recommend that regular updates be applied to any IDS or security device.

The network administrator can configure the IDS-enabled router to choose the appropriate response to various threats. When packets in a session match a signature, the IDS system can be configured to take one, two, or all the following actions:

- Send an alarm to a syslog server or a Cisco Secure IDS Director (centralized management interface).
- Drop the offending packet.
- Reset the TCP connection.

Security best practice procedures recommend that you use the drop and rest actions together. In practice, this would mean that when the FFS IDS receives a packet that matches its IDS attack signatures; the packet is dropped, thereby preventing it from reaching the targeted device in your network. Because attacks come in the form of multiple packets, simply dropping only one packet is not enough to protect your network. The FFS IDS will proactively send a tcp *reset* to the device that sent the offending packet, thereby causing the connection to drop (reset). This combination response is effective because the specific packet and the communication session are dropped.

When to Use the FFS IDS

Cisco IOS Firewall IDS capabilities are ideal for providing additional visibility at intranet, extranet, and branch-office Internet perimeters. Networks of all sizes and complexity will enjoy a more robust protection against attacks on the network and can automatically respond to threats from internal or external hosts.

The Cisco IOS Firewall with intrusion detection is intended to satisfy the security goals of all customers and is particularly appropriate for the following scenarios:

- Enterprise customers who are interested in a cost-effective method of extending their perimeter security across all network boundaries—specifically branch-office, intranet, and extranet network perimeters
- Small and medium-sized businesses looking for a cost-effective router that has an integrated firewall with intrusion-detection capabilities
- Service provider customers who want to set up managed services, providing their customers with firewalling and intrusion detection, all housed within the necessary function of a router

FFS IDS Operational Overview

By now, it should be apparent that understanding packets is important in networking. This is a realization that comes slowly for some people; however, after you accept this truth, networking should become much easier to understand. Everything is a packet, and all network devices are designed to do something with a packet. Sometimes, this is forwarding the packet to its destination, inspecting it, or even altering it in some way to accomplish a goal. This understanding is something that many hackers have figured out, and they use this knowledge to serve the dark side. That is melodramatic but truthful because it is no fun rebuilding a server at 3:00 a.m. because it has been compromised, or dealing with a rampant virus, botnet, or attacks that can bring a network to its knees. This book is not designed to make you an expert at packets, but it introduces you to many of the fundamental truths of network security that provide a solid understanding of how the real world functions. If you need to learn more, you can build on this beginning. That being said, you do not need to live at the packet level; simply knowing that it is there and that it functions is the basis for everything networking.

Note Perhaps the person I respect the most who educates people about living at the packet level is Laura Chappell. Visit her website and her many online resources at www.packet-level.com/.

Living at the packet level is an excellent mindset for troubleshooting, especially if you can “be the packet” and follow its course. From an IDS perspective, packets are the meat and potatoes of everything they look at. Cisco developed its Cisco IOS Software-based intrusion detection capabilities in the Cisco IOS Firewall with flexibility in mind so that individual attack signatures could be disabled in case of false positives. Also, although it is preferable to enable the CBAC security engine’s firewall and intrusion detection features to support a network security policy, each of these features can be enabled independently and on different router interfaces.

The Cisco IOS Intrusion Detection System (IDS) acts as an inline intrusion detection sensor, watching packets as they traverse the router’s interfaces and acting upon them in a definable fashion.

The Cisco IOS IDS identifies the most common attacks using signatures to detect patterns of misuse in network traffic (*attack signatures*). The Cisco IOS Firewall feature set’s intrusion detection signatures were chosen from a broad cross-section of intrusion detection signatures. The signatures represent severe breaches of security, the most common network attacks, and information gathering scans.

In Cisco IOS IDS, signatures are categorized into four types:

- **Info atomic:** Detect patterns as simple as an attempt to access a specific port on a specific host, such as a port scan.

- **Info compound:** Detect complex patterns, such as a sequence of operations distributed across multiple hosts over an arbitrary period of time. In general, both kinds of informational signatures detect attackers' information-gathering activities.
- **Attack atomic:** Detect patterns where an attacker is attempting to access a single host device.
- **Attack compound:** Detects complex attack activities spread 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 that represent the most common network attacks and information gathering scans not commonly found in an operational network.

The following describes the packet auditing process with Cisco IOS IDS:

1. You create an audit rule, which specifies the attack signatures that should be applied to packet traffic and the actions to be taken when a match is found. An audit rule can be as flexible and specific as needed to meet the goals of your security policy. A sample rule follows in which you suspect or want to prevent the spamming of email messages, so the IDS is configured to audit all SMTP traffic and ensure that there are no more than 100 recipients:

```
ip audit smtp spam 100
```

2. You apply the audit rule to an interface on the router, specifying a traffic direction (*in* or *out*). The following example applies the audit rule to look at all inbound SMTP traffic to the router:

```
ip audit smtp in
```

3. If the audit rule is applied to the *in* direction of the interface, packets passing through the interface are audited before the inbound ACL has a chance to discard them. This enables an administrator to be alerted if an attack or information-gathering activity is underway, even if the router would normally reject the activity. It is considered best practice to apply IDS audit rules inbound because they are inspected.
4. If the audit rule is applied to the *out* direction on the interface, packets are audited after they enter the router through another interface. In this case, the inbound ACL of the other interface might discard packets before they are audited. This could result in the loss of IDS alarms, even though the attack or information-gathering activity was thwarted.
5. Packets going through the interface that match the audit rule are audited by a series of modules, starting with IP; then either ICMP, TCP, or UDP (as appropriate); and finally, the application level.
6. If a signature match is found in a module, the following user-configured actions occur:
 - If the action is **alarm**, the module completes its audit, sends an alarm, and passes the packet to the next module.

- If the action is **drop**, the packet is dropped from the module, discarded, and not sent to the next module.
- If the action is **reset**, the packets are forwarded to the next module, and packets with the reset flag set are sent to both participants of the session, if the session is TCP.

If there are multiple signature matches in a module, only the first match fires an action. Additional matches in other modules fire additional alarms, but only one per module. IDS can reset only a TCP-based connection because this protocol has a SYN ACK and the all-powerful RST, which the IDS can send back to the attacker's TCP-based session and shut down that application. UDP is not connection-oriented, so this is not something that can be reset—thus the need for ACLs on a blocking device such as a router or PIX Firewall.

FFS Limitations

CBAC enhances the effectiveness of IOS routers as security devices. Used with other available security enhancements, you can use IOS routers for more than packet forwarding, thus increasing their ROI and allowing administrators to cost-effectively implement more secure networks. Of course, there is no perfect security device. Following are some operational issues and limitations to CBAC of which administrators should be aware:

- Intrusion detection's performance impact depends on the configuration of the signatures, the level of traffic on the router, the router platform, and other individual features enabled on the router, such as encryption, routing, and so on. Enabling or disabling individual signatures does not significantly alter performance; however, signatures configured to use ACLs have a significant performance impact because the more you ask the router to inspect a packet, the greater its effect on router performance.
- For auditing atomic signatures, there is no traffic-dependent memory requirement. For auditing compound signatures, CBAC allocates memory to maintain the state of each session for each connection because by definition compound signatures are going to multiple machines. Memory is also allocated for the configuration database and for internal caching.
- CBAC inspection is not performed on packets with the source or destination address of the firewall interfaces. This impacts the router's operation two different ways:
 - vty (that is, Telnet) sessions between administrators and the firewall are not inspected.
 - Management, authorization, and accounting (TACACS/RADIUS) traffic is not inspected because it, too, is destined to the router's interface.
- Encrypted packet payloads, such as those used in VPNs, are not inspected unless the router is the encrypted link endpoint.

In general, having the more advanced functions available does increase the security of your router and network. However, these functions do not address the best practices in

making the router a secure device when you do not employ them. The following section discusses this aspect of securing a router because given the cost and effort needed to maintain the FFS, you are likely going to deploy it only at the edge of your network; therefore, protecting the inside devices is covered next.

Secure IOS Template

So far, this chapter has covered the different ways to secure your router and use it as a supplement to a dedicated firewall. This section explores how to harden your router and some of the best practices available for making the router a more secure device on your network. For the sake of brevity, you will not see coverage of every single ACL and command possible to secure your router. Following are a couple reasons for this choice:

- The physical constraints of this book do not allow it, so content must be prioritized. Some items left out are specific to certain businesses in networking (ISPs, for example); most networks easily use the remaining items.
- Certain parts of the recommendations, such as TACACS and RADIUS, are covered in previous chapters, so there is no need to cover them again here.

This section is not meant to teach you how to secure your router with brief explanations so that you can decide which commands are appropriate for your network. You can apply these commands and suggestions today! There are many websites that offer all sorts of templates; however, this section discusses only a few of the options available to you. Definitely use this section as a starting point and find the templates that best match your security needs and policy.

Tip The Cisco SDM Security device manager is a mature, sound GUI tool that is now a shipping standard with the security/VPN routers. This GUI offers a robust setup and configuration of VPN and CBAC; it also does a router analysis and locks down the router, but before you start that process, you can find some suggestions to the user about how to do it at www.cisco.com/en/US/products/sw/secursw/ps5318/index.html.

The configuration commands in Example 8-1 are in **bold** text so that they stand out from the supporting comments, which are highlighted for readability. The secure template assumes the topology in Figure 8-4.

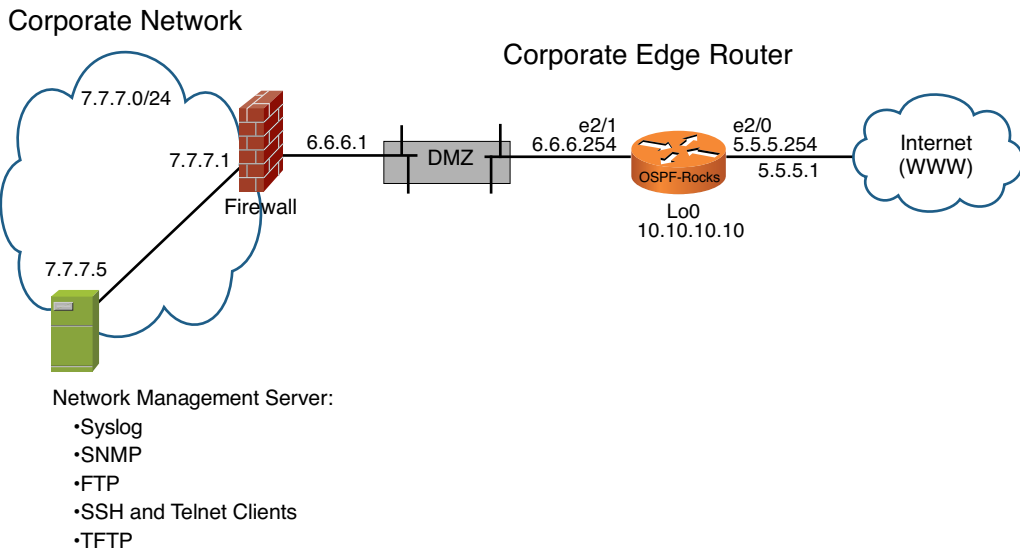


Figure 8-4 *Secure IOS Template Topology*

Example 8-1 *Secure IOS Template*

! The very first step before beginning is to ensure that your IOS is upgraded to the latest stable version. This will ensure that any older security or operational issues are resolved as the best that can with this simple. Remember any operating system will have bugs and flaws, so minimizing them is definitely best practice.

!

The Nagle congestion control algorithm is something that many companies turn on to improve the performance of their Telnet session to and from the router. When using standard Telnet, via TCP to send keystrokes between machines; TCP tends to send one packet for each keystroke typed. On larger networks, many small packets use up bandwidth and contribute to congestion. John Nagle's algorithm (RFC 896) helps alleviate the small-packet problem in TCP. In general, it works this way: The first character typed after connection establishment is sent in a single packet, but TCP holds any additional characters typed until the receiver acknowledges the previous packet. Then the second, larger packet is sent and additional typed characters are saved until the acknowledgment comes back. The effect is to accumulate characters into larger packets (chunks) and pace them out to the network at a rate matching the round-trip time of the given connection. Keepalives ensure that no TCP connections to the router get hung.

!

service nagle

!

This command will disable the auto loading of configuration files from a network server that is disabled, except on systems without NVRAM or with invalid or incomplete

information in NVRAM. In these cases, auto loading of configuration files from a network server is enabled automatically.

!

no service config

!

Attackers will often map a network using ICMP packets with the source route option turned on. Normal traffic does not require source route reporting. This command will stop the router from providing that information.

!

no ip source-route

!

Enabling the two services below allows the router to monitor TCP keepalives on incoming connections and ensures that any sessions left hanging by remote system if it crashes or disconnections abruptly will not block or use up the available router vty (Telnet) ports, thereby denying legitimate connections. In a sense, you could be locked out of the router as a denial of service.

!

!

service tcp-keepalives-in

service tcp-keepalives-out

!

By default, log messages are not time stamped or marked in anyway that would allow you to know when they occurred. You should activate time stampings in all debug messages and log entries down to the millisecond to ensure that you can determine the relevance of each message and ensure that your router's clock is set properly—otherwise it will not be very effective! The following setting will produce entries that are similar to the following:

```
Sep 4 23:58:11.437: %LINK-3-UPDOWN: Interface FastEthernet0/10, changed state to up
```

The command line options in the timestamps command are as follows:

- debug: all debug information is time stamped

- log: all log info is time stamped

- datetime: the date and time is include in the syslog message

- localtime: the local time of the router is used in the log message

- show-time zone: the time zone defined on the router is included (useful if the network crosses multiple time zones and we suggest standardizing on single time zone if this is the case)

- msec: time accuracy to milliseconds - useful if NTP is configured.

!

service timestamps debug datetime msec show-timezone localtime

service timestamps log datetime msec show-timezone localtime

!

By default, a syslog message contains the IP Address of the interface it uses to leave the router. You can require all syslog messages to contain the same IP Address,

regardless of the interface they use. Many large enterprise networks or ISPs use the loopback IP Address to more clearly identify the routers in their network. This keeps their syslogs consistent and allows them to enhance the security of their syslog server. You can also set this interface destination to be any active interface on the router if you do not have a loopback interface configured; however, loopbacks are recommended as it helps you ensure each router is sending information from an address you specify.

!

logging source-interface loopback0

!

The 'service password-encryption' command provides minimal security for user, line, ppp, radius and assorted other passwords and keys that must be stored in the IOS configuration file. The command causes passwords in the config file to be encrypted with a reversible encryption that keeps people from finding your passwords by glancing at your configurations. Note that this encryption does not provide real protection; we recommend considering the use of the enable secret password or TACACS/RADIUS controlled logins.

!

!

service password-encryption

!

By default, Cisco has enabled routers to now act as DHCP clients by default; this is really not a necessary service to have running, so shut it off. Also, given the issues with TCP and UDP small servers, make sure they are off! For example, one of the small servers is "Chargen," which is a character generator service that is used to generate a stream of characters for diagnostic purposes. Then there is the "echo" service that merely echoes back every character that is sent to it. Pointing the "chargen" service at the "echo" service creates a loop that causes an enormous amount of traffic to be generated and will eventually overwhelm the router's CPU and RAM resources; thus, we have the makings of a very serious denial of service attack (DoS). The easiest way to prevent this kind of attack from happening is to disable these services on the router. The commands to do so are "no tcp-small-servers"—disables echo, chargen, discard, and daytime; "no udp-small-servers"—disables echo, chargen, and discard.

!

no service udp-small-servers

no service tcp-small-servers

no service dhcp

!

Not all services are bad; in fact, new entry to the service category is quite useful. Essentially, by enabling it, your syslog entries are numbered to ensure that they are not tampered with to hide hacking from you!

!

service sequence-numbers

hostname OSPF-Rocks

!

Logging is a must in almost every case, so turn it on! Plus, with all the logging we are doing in this configuration, it might be a good idea to rate limit the log messages sent per second to not overwhelm your server because the entries can climb rapidly when you are logging ACLs!

!

logging 7.7.7.5

logging buffered 16384 debugging

logging rate-limit ?

!

When a message is sent to the console port of the router, this results in CPU interrupt occurring in order for the log message to be delivered to the console port and, considering the level of logging that is going on, disable console logging until needed. Console logging is very effective when troubleshooting; you are physically attached to the router, so keep this command ready.

!

no logging console

!

Almost all passwords and other authentication strings in Cisco IOS configuration files are encrypted using the weak, reversible scheme used for user passwords. To determine which scheme has been used to encrypt a specific password, check the digit preceding the encrypted string in the configuration file. If that digit is a 7, the password has been encrypted using the weak algorithm. If the digit is a 5, the password has been hashed using the stronger MD5 algorithm. Even though enable secret is used for the enable password; do not forget service password-encryption so that the remaining passwords are stored in the configuration with type 7 encryption rather than in plain text. Of course, the most secure password type is enable secret, so use it with some CAPITAL letters and some Num83r2; it makes brute force attacks harder. The encryption algorithm type 7 used in enable password and service password-encryption is reversible. The enable secret command provides better security by storing the enable secret password using a non-reversible cryptographic function. The added layer of security encryption it provides is useful in environments where the password crosses the network or is stored on a TFTP server.

!

enable secret <PASSWORD>

no enable password

!

Use TACACS+ for AAA login authentication. Ensure that the local account is case-sensitive, thus making brute-force attacks less effective.

!

aaa new-model

aaa authentication login default group tacacs+ local-case

aaa authentication enable default group tacacs+ enable

```

aaa authorization commands 15 default group tacacs+ local
aaa accounting exec default stop-only group tacacs+
aaa accounting commands 15 default stop-only group tacacs+
aaa accounting network default stop-only group tacacs+
tacacs-server host 7.7.7.5
tacacs-server key OSPF-r0ck2
!

```

In the event that TACACS+ fails, use case-sensitive local authentication with a username on the router so you can still access it. If TACACS+/RADIUS is not available in your network then configure AAA to use locally (on the router) stored username and passwords. The use of authentication keeps attackers guessing, and the router more secure; remember, security is all about multiple layers of defense.

```

!
username <USERNAME> password <PASSWORD>
!

```

Do I really need to explain why you should not use the built-in web server? Sometimes Cisco takes the web too far—it is a router, Jim! Let it route!

```

!
no ip http server
no ip https server
!

```

Allows us to use the low subnets and go classless, which are areas that have not typically been used but as we run out of IPv4 addresses this is becoming more and more important.

```

!
ip subnet-zero
ip classless
!

```

Why these services are still on by default and in IOS is anyone's guess; however, for your sanity and the security of your network, turn them off. As Cisco's IOS has evolved, some of these services have become turned off by default; however, it is always considered best practice to ensure that they are turned off.

```

!
no service pad
no ip source-route
no service finger
no ip bootp server
no ip domain-lookup
!

```

TCP intercept helps prevent SYN-flooding attacks by intercepting and validating TCP connection requests. In intercept mode, the TCP intercept software intercepts TCP synchronization (SYN) packets from clients to servers that match an extended access list. The router responds, and they are allowed to communicate if it is a valid

connection.

!

ip tcp intercept list 120

!

IOS watches and manages a TCP connection for 24 hours after no activity. Why? Who knows? Regardless, it should be changed because there is no need to have a router keep that amount of information in its memory.

!

ip tcp intercept connection-timeout 60

!

Keep half-open TCP connection attempts open only 10 seconds instead of the default 30 seconds. This will help the router defeat Denial of Service attacks since it will close half open connections much faster.

!

ip tcp intercept watch-timeout 10

!

These commands determine when TCP intercept should deactivate or activate; in this case, 1500 and 6000, respectively—the defaults are not very realistic at 900 and 1100.

!

ip tcp intercept one-minute low 1500

ip tcp intercept one-minute high 6000

!

Cisco Systems has added a core dump facility to its IOS. This core dump facility operates like many other similar systems. When a router crashes, a copy of the core memory is kept. Before the memory is erased on reboot, the Cisco router can be set up to copy the core dump out to a server. An account (FTP, TFTP, or RCP) and sufficient disk space (equal to the amount of memory on the router per dump) must be set up and allocated. Catch core dumps in case of a router crash; this is very important with a "security router" because a denial-of-service (DOS) attack might have been successful and crashed your router, so it is good to know what happened. We have configured our Network Management server inside our firewall to accept FTP connections from the router. Make sure that you give the core dump files a unique name, as shown in the following lines. It is recommended that access to the "Cisco core dump" account be made as secure as possible. For example, do not send core dumps to the same FTP server as the one used to provide generic anonymous or user FTP accounts.

!

ip ftp username <FTP SERVER USERNAME>

ip ftp password <PASSWORD>

exception core-file <UNIQUE FILE NAME>

exception protocol ftp

exception dump 7.7.7.5

!

TFTP is the most common tool for uploading and downloading IOS upgrades or configurations. The TFTP server's security is critical. That means using security

tools that only allow a TFTP connection to be successful based on the source IP address. Cisco's IOS allows TFTP to be configured to use a specific IP interfaces address. This allows a fixed ACL on the TFTP server based on a fixed address on the router. This fixed IP Address is commonly the loopback interface if it is configured as these interfaces are frequently used in managing a router. However, if you are using loopback interfaces in your network, the interface closest to the TFTP server should be used; the command is shown below. FTP is also included because it was previously configured in this template.

```
!
ip tftp source-interface <SOURCE INTERFACE>
ip ftp source-interface <SOURCE INTERFACE>
!
```

CEF is an advanced, Layer 3 switching technology inside a router. It defines the fastest method by which a Cisco router forwards packets from ingress to egress interfaces. The `ip cef` command enables CEF globally, not all router support CEF so check your docs.

```
!
ip cef
!
```

Set the time zone properly. It is best to standardize on one time zone for all routers and servers, thus making problem tracking easier. I recommend using the time zone where all your network management devices and servers are located so all logs, traps, and events are in sync.

```
!
clock timezone GMT 0
!
```

NTP is the most overlooked feature on many networks. The Network Time Protocol (NTP) is a protocol designed to time-synchronize a network of machines. It provides a precise time base for networked workstation, servers, and other devices on the network. NTP runs over UDP, which in turn runs over IP. An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a timeserver. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two machines to within a millisecond of one another. Many system administrators configure time synchronization for servers but do not continue that first step to include network devices. If you wish to compare the syslog information from devices all over your network, you must synchronize the time on all of them. Comparing logs from various network devices is essential for many types of troubleshooting, fault analysis, and security incident tracking. Without precise time synchronization between all the various logging, management, AAA and security functions, this sort of comparison would be impossible. When activating NTP, synchronize the router's clock with a local (trusted and authenticated) NTP server. The SECRETKEY must be the same on both the router and the NTP server. Note that NTP is slow to get synchronized properly in the beginning; it is a Cisco thing, so be patient!

```

!
ntp authentication-key 6767 md5 <SECRETKEY>
ntp authenticate
ntp update-calendar
ntp server 7.7.7.5

```

! Configure the loopback0 interface as the source of our log messages. This is often used for routing protocols also because a logical interface does not go down; thus, it is very reliable. Assign an IP address that uniquely identifies this router. One trick is to allocate a netblock for use as the router loopback netblock.

```

!
int loopback0
 ip address 10.10.10.10 255.255.255.255
 no ip redirects
 no ip unreachable
 no ip proxy-arp

```

! Configure and thus activate the null0 interface as a place to send naughty packets. This becomes the "roach motel" for packets—they can route in, but they cannot route out.

```

!
interface null0
 no ip unreachable

```

```

!
interface Ethernet2/0
 description Unprotected interface, facing towards Internet
 ip address 5.5.5.254 255.255.255.0
 no ip directed-broadcast
 no ip unreachable
 no ip redirects
 no ip mask-reply
 no ip proxy-arp

```

! Should you run CEF verify? Yes, if the data path is symmetric, but no if the data path is asymmetric. Use the ip verify unicast reverse-path interface command on the input interface on the router at the upstream end of the connection. This feature examines each packet received as input on that interface. If the source IP address does not have a route in the CEF tables that points back to the same interface on which the packet arrived, the router drops the packet.

```

!
 ip verify unicast reverse-path

```

! Apply our template ACL, more on what this ACL is covering later in the configuration, but applying it is crucial to its success. The following command is how an access-list

is applied to an interface.

!

ip access-group 2010 in

!

Rate limiting traffic to protect the router and by default your infrastructure is extremely important. The values might be tweaked to meet your needs but, in general, we recommend the following. Allow UDP to use no more than 2 Mb/s of the pipe; caution, however, if you are running video on demand as it uses UDP packets.

!

rate-limit input access-group 150 2010000 250000 250000 conform-action transmit exceed-action drop

!

Allow ICMP to use no more than 200 Kb/s of the pipe.

!

rate-limit input access-group 160 500000 62500 62500 conform-action transmit exceed-action drop

!

Allow multicast to use no more than 5 Mb/s of the pipe.

!

rate-limit input access-group 170 5000000 375000 375000 conform-action transmit exceed-action drop

!

Disables the sending of ICMP redirect messages to learn routes; let the hackers wonder!

!

no ip redirects

!

Disables the sending of ICMP protocol unreachable and host unreachable messages and, once again there is no reason to allow ICMP to educate hackers about your network.

!

no ip unreachable

!

Dropping IP directed broadcasts makes routers less susceptible to a denial-of-service attack. The configuration command "no ip directed-broadcast" means that the translation of directed broadcast to physical broadcasts is disabled. If enabled, a broadcast to a particular network could be directed at a router interface, producing effects that might be undesirable and potentially harmful. An example of the ill effects of directed broadcasts being enabled is the so-called SMURF attack.

!

no ip directed-broadcast

!

Cisco IOS Software examines IP header options on every packet. It supports the IP header options Strict Source Route, Loose Source Route, Record Route, and time stamp,

which are defined in RFC 791. If the software finds a packet with one of these options enabled, it performs the appropriate action. If it finds a packet with an invalid option, it sends an ICMP Parameter Problem message to the source of the packet and discards the packet. The IP protocol provides a provision that allows the source IP host to specify a route through the IP network. This provision is known as source routing, which is specified as an option in the IP header. If source routing is specified, Cisco IOS forwards the packet according to the specified source route in the IP header. This feature is employed when you want to force a packet to take a certain route through the network. The default is to perform source routing. As a general rule of thumb, if you are not using IP source routing, turn it off. IP source routing is a well-known security vulnerability used in attacks against a system or to bypass firewalls.

!

no ip source-route

!

The configuration "no ip proxy-arp" means that the router does not respond to ARP requests for other hosts on the network connected to this interface if it knows the MAC address of those hosts. Again, this is to prevent undesirable effects on the connected network and potential security problems. In other words, do not have the router pretend to be something its not.

!

no ip proxy-arp

!

Disables the sending of ICMP mask reply messages. The default for Cisco routers is not to do this, but it never hurts to input the command anyway just to be sure.

!

no ip mask-reply

!

Enables IP accounting with the ability to identify IP traffic that fails IP access lists, thereby allowing your router to log all naughty business. Be sure to check it.

!

ip accounting access-violations

!

If you allow multicast in your network or participate in the MBONE, the following multicast filtering steps help to ensure a secure multicast environment. These must be applied per interface.

!

ip multicast boundary 30

!

Keep flow data for analysis. If possible, export it to a cflowd server.

!

ip route-cache flow

!

interface Ethernet2/1

```
description Protected interface, facing towards DMZ
```

```
ip address 6.6.6.254 255.255.255.0
```

```
!
```

Do we run unicast verify? Yes, if the data path is symmetric. No, if the data path is asymmetric. See above interface description for more information on this command.

```
!
```

```
ip verify unicast reverse-path
```

```
!
```

The following commands have been described previously; for additional information, refer to earlier in the configuration file.

```
!
```

```
no ip redirects
```

```
no ip unreachable
```

```
no ip directed-broadcast
```

```
no ip proxy-arp
```

```
ip accounting access-violations
```

```
ip multicast boundary 30
```

```
no ip mask-reply
```

```
ip route-cache flow
```

```
!
```

Source routing allows the path to be specified in a packet. This could allow the packet to bypass firewalls and so on. Disable this feature!

```
!
```

```
no ip source-route
```

```
!
```

This is a default route to the Internet (could be a routing protocol instead) and if you choose a routing protocol, OSPF is highly recommended.

```
!
```

```
ip route 0.0.0.0 0.0.0.0 5.5.5.1
```

```
!
```

Route to network on the other side of the firewall.

```
!
```

```
ip route 7.7.7.0 255.255.255.0 6.6.6.1
```

```
!
```

The following static routes will black hole networks that are not supposed to be routable on the public Internet. Be very careful about enabling these when running TCP Intercept. The TCP Intercept command directs the router to act as a TCP socket proxy. When the router receives the SYN packet, the router (instead of the destination) initially responds with the SYN|ACK. This is where the interaction between TCP Intercept and black hole routes causes a problem. If you create black hole routes for all bogon ranges and point them to the null device, and if someone launches a SYN flood from a bogon range, the router sends the SYN|ACK to the null device. The router is not (yet) intelligent enough to realize that it has done this, and the TCP Intercept queue begins to build quickly. By default, the timeouts are not aggressive

enough to work through this problem.

!

Cisco has introduced an on device command archive command. When enabled, these sets of commands will record every configuration change made on the router and in the example that follows, report it to a syslog server provided one is configured in the logging section. This command is especially useful when coupled with AAA or TACACS as it will also record what user made the change. This is a great way to do internal auditing and can be a wonderful education tool for new engineers. Plus if you configured your syslog server to alert on configuration changes you can know and "see" what is happening. For example, should someone really be changing things outside of a maintenance window?

!

archive

log config

logging enable

notify syslog

hidekeys

!

Cisco routers can now run specialized scripts utilizing the TCL programming language. These scripts can be very powerful and allow for the automation of a variety of tasks and jobs; however, like any tool they can be abused so turn this feature off if not in use.

!

no scripting tcl init
no scripting tcl encdir

!

Export NetFlow data to our NetFlow server, 7.7.7.5. NetFlow provides some statistics that can be useful when tracing back to the true source of a spoofed attack. We also use the source as the loopback interface, which is a best practice.

!

ip flow-export source loopback0
ip flow-export destination 7.7.7.5 2055
ip flow-export version 5 origin-as

!

Log anything interesting to the syslog server. Capture all the logging output sent from the loopback interface; this makes the ID of this router in the various places recording data easy and uniform to identify.

!

logging trap debugging
logging source-interface loopback0
logging 7.7.7.5

!

Do not share Cisco Discovery Protocol (CDP) information from your secure router

because CDP contains crucial bits of information about your network topology, device configuration, network devices that are in use, IP addresses, and so on. This command disabled CDP globally. If you require CDP on an interface, use `cdp run` and `disable cdp` (not `cdp enable`) on the Internet-facing interface. In other words, use CDP only on interfaces where it is needed—never globally. Note that Cisco ships all devices with CDP enabled by default starting with IOS 11.1CA.

!

no cdp run

!

SNMP is very important for network management, particularly in conjunction with MRTG to track usage statistics. To keep SNMP access even more secure, treat the COMMUNITY string as a password; keep it difficult to guess by using a combination of CAPS, lowercase, and numbers. Ultimately a SNMP community string is the password for SNMP Services so the string should follow your corporate password policy. This is important because the community string is not encrypted. Then, further protect access by including an access control list (ACL) that determines what network/hosts can access SNMP, only if they have the proper community string. Now that is a real layered security approach!

If SNMP is going to be used in read/write mode, think very carefully about the configuration and why there is a requirement to do this because configuration errors in this scenario could leave the router very vulnerable. I have developed and seen tools that, through the use of SNMP Read/Write, can automatically reset password and alter configurations. There are very few good reasons to allow read/write access to a device via SNMP, read only is best practice!

!

If possible, put an ACL at the edge of your network to prevent potential attackers from probing your network via SNMP. There are many publicly and commercially available tools that will scan any network on the Internet via SNMP. This could map out your entire network and/or discover a device that has had SNMP left open. When performing security audits and vulnerability assessments, I have done an SNMP Walk on devices and learned a great deal about a person's network.

!

snmp-server community <COMMUNITY> RO 20

snmp-server location Tampa, FL

snmp-server contact Cyberwraith Consulting [networkoperations@cyberwraith.com]

snmp-server host 192.168.254.70 <COMMUNITY_STRING>

!

In the configuration, this ACL would appear at a different location; however, for completeness, I have moved it here for easy reference. Access list 20 permits SNMP access to this device if the requests come from the server (IP Address: 7.7.7.5) and by default if access is not permitted and is then denied when using Cisco ACLs. Notice that I entered the normally implicit deny any command because I have added the `log` keyword at the end. The inclusion of this deny keyword has the router log denied all

SNMP query attempts to our syslog server allowing us to see who might be trying to access our routers.

!

```
access-list 20 remark ACL TO CONTROL SNMP ACCESS
```

```
access-list 20 permit 7.7.7.5
```

```
access-list 20 deny any log
```

!

Protect and set expectations with an appropriately stern banner that reflects the level of security and monitoring applied to your network. It is also important to set everyone's expectations accessing the router and what happens if attacks are made against it. Although we are just showing the Message of the Day (MOTD) Banner, you could apply the same banner to the console port, aux port, AAA Login, and whenever a user accesses EXEC mode.

!

```
banner motd %
```

Warning!!! This system is solely for the use of authorized users and only for official purposes. Users must have express written permission to access this system. You have no expectation of privacy in its use and to ensure that the system is functioning properly, individuals using this system are subject to having their activities monitored and recorded at all times. Use of this system evidences an express consent to such monitoring and agreement that if such monitoring reveals evidence of possible abuse or criminal activity the results of such monitoring will be supplied to the appropriate officials to be prosecuted to the fullest extent of both civil and criminal law.

Unauthorized Access to this system is a violation of Federal Electronic Communications Privacy Act of 1986, and may result in fines of \$250,000 and/or imprisonment (Title 18, USC). All IP traffic is logged and violators will be prosecuted.

%

!

Another type of banner available is the "exec" banner, which is displayed at the time a user has successfully authenticated and logged in when they enter exec mode on the router. Exec mode is analogous to super user (UNIX) or administrator (Windows).

!

```
banner exec ^
```

Please note that this device is part of a production network and all configuration changes need to be approved in advance. All changes should be recorded and the configuration backed up before you make changes.

^

!

Apply a password to the console port of a router. Requiring a password on the physical console port provides another layer of security by requiring anyone plugging into the device to supply a password. Including the transport input disables reverse Telnet and protects the physical ports against access.

The connection timeout value for Console and AUX ports on a router is 10 minutes. This timeout is controlled by the `exec-timeout` command, as shown in the configuration below. VTY (Telnet) sessions do not have an associated timeout value. Leaving the VTY timeout unchanged is generally regarded as bad practice because it will hog the few available ports on the router and could cause maintenance access problems in the time of emergencies. Notice that setting the idle timeout to 0 means that the session is left connected indefinitely.

```
!
line con 0
  exec-timeout 15 0
  transport input none
line aux 0
  exec-timeout 15 0
  transport input none
!
```

Apply an access control list (ACL) to the VTY (Telnet) ports that define which systems, by source IP address, can attempt to access this router via Telnet. Most IOS versions support only five VTY ports; this means that when you look in the configuration and see "line VTY 0 4," there can be a maximum of five Telnet connections if you count 0 as a line (0, 1, 2, 3, 4). In the following example, we are configuring a group of VTY lines (0-3) to all have the same operating parameters. The access list is then applied to the VTY ports through the `access-class` command as shown below. The command `logging synchronous` is an all time favorite of mine; it preserves what you have been typing when the router begins reporting information which by default the router tacks on the information to the line you're typing on causing you to completely lose track of what has been done so far. The `logging synchronous` command lets the router give you the output like normally but not on the line you are typing on!

```
!
line vty 0 3
  access-class 100 in
  exec-timeout 15 0
  logging synchronous
  transport input telnet ssh
!
```

Notice in the preceding configuration lines the use of telnet and SSH as a means to access the VTY lines. Best practice is to only use SSH; however, the IOS version you are running might not allow that. Whenever possible use SSH and not telnet, which would mean the keyword `telnet` would not be included in the preceding statement. Remember to remove it!

```
!
```

The definition of this access list is important to understand and would normally appear much earlier in the configuration; however, for ease of understanding, I have

moved it to the relevant section. access control list 100 will deny everyone access to the router and permit connection attempts from the Network Management server (7.7.7.5) or the firewall (6.6.6.1); only if SSH (port 22) or Telnet (port 23) is used, we log every successful access and this allows us to monitor who is connecting, when, and how. Of course, we also log any denied access attempts to learn the same information. This also serves to create an audit trail of all access to the router through the use of extended ACLs to log some additional data.

!

```
access-list 100 remark DEFINE TELNET ACCESS TO THE ROUTER
access-list 100 permit tcp host 7.7.7.5 host 0.0.0.0 range 22 23 log-input
access-list 100 permit tcp host 6.6.6.1 host 0.0.0.0 range 22 23 log-input
access-list 100 deny ip any any log-input
```

!

Whenever possible, enable SSH connectivity because SSH is much more secure than Telnet. Obviously, you must have an IOS image that supports SSH, and do not forget to generate the key with the `crypto key generate RSA` command.

!

Leave one VTY safe (line #4) for emergency access, just in case. The host 7.7.7.8 is a secure host in your network management operations center. If all the VTYS are occupied, this leaves one VTY available and logging is also happening.

!

```
line vty 4
access-class 105 in
exec-timeout 15 0
logging synchronous
transport input telnet ssh
```

!

NOTE: You can also use AAA during the login process as well and if it is configured properly, you should!

!

```
access-list 105 remark VTY Access ACL
access-list 105 permit tcp host 7.7.7.8 host 0.0.0.0 range 22 23 log-input
access-list 105 deny ip any any log-input
```

!

Although this section covers how to configure a router virtually and how it operates, do not forget about the physical security of your routers. Physical access to network devices usually allows unprecedented levels of control to tap the link, block, jam, inject traffic, and so forth. It makes no sense to install complicated security measures when access to the hardware is not equally secure.

Routing Protocol Security

Any WAN these days runs a dynamic routing protocol; the most common and secure of which is Open Shortest Path First (OSPF). Although an in-depth discussion of how dynamic routing with OSPF works is outside the scope of this book, a brief overview is in order. Dynamic routing provides a means for routers to share knowledge of the networks of which they are aware. OSPF is one of the protocols that can be used for this exchange of information. If attackers were to isolate or inject bogus routing updates into your dynamic routing protocol, they could cause all sorts of issues within your network, resulting in critical data not getting across the network.

To safeguard your network's routing information within your WAN, you can configure route authentication between routers. This section briefly discusses how route authentication in OSPF is done and how it can benefit the security of your network.

Route authentication enables peer routers to positively identify the source of incoming encrypted dynamic routes. This means that attackers cannot forge erroneous routes or tamper with the exchange of routes without detection.

OSPF Authentication

OSPF is responsible for transmitting routing updates and building a routing table to ensure connectivity across a network. OSPF incorporates security within its function as a routing protocol. The authentication capacity provided in OSPF is sufficient to protect the exchange of routing information.

This section describes OSPF authentication as part of a total security plan and explains what neighbor router authentication is, how it works, and why you should use it to increase your overall network security. Following are several topics of importance about this issue that this section discusses:

- Benefits of neighbor authentication
- Conditions for deploying OSPF neighbor authentication
- How neighbor authentication works
- Configuring neighbor authentication

OSPF is responsible for transmitting routing updates and building a routing table enabling data to flow across a network. OSPF authentication was designed to protect only the integrity of the routing information within an OSPF routing domain; in other words, data is not protected—you need additional encryption to accomplish that. You can prevent any OSPF router from receiving fraudulent route updates by configuring the router to use a type of security known as *neighbor router authentication*. Following are two design characteristics that truly define how OSPF authentication operates:

- OSPF authentication is activated for an entire network, or in OSPF terminology, an area.
- The authentication key must match for neighboring routers on the same link.

Following are several different ways that you can deploy this type of security within your OSPF network:

- By assigning the same OSPF authentication key throughout the entire OSPF area
- By assigning a different key for every link within the network

Regardless of which technique you decide to use, the passwords used between neighboring routers must match.

Note This section refers to neighbor router authentication as *neighbor authentication*. Neighbor router authentication is also sometimes called *route authentication*. The use of neighbor enables us to be extremely specific in our discussion.

Benefits of OSPF Neighbor Authentication

When configured, neighbor authentication occurs whenever routing updates are exchanged between neighboring OSPF routers within the OSPF area that has authentication activated. This authentication ensures that a router receives reliable routing information from a trusted source (that is, a neighbor router also running OSPF).

Without OSPF authentication, unauthorized or deliberately malicious routing updates could compromise the integrity of your network traffic. A security compromise could occur if an unfriendly party diverts or analyzes your network traffic. The compromise might not be the result of malicious action.

For example, an unauthorized or compromised device could send a fictitious routing update to convince your router to send traffic to an incorrect destination. This diverted traffic could be analyzed to learn confidential information about your organization, or it could merely be used to disrupt your organization's capability to effectively communicate using the network. OSPF authentication prevents any such fraudulent route updates from being received by your router.

When to Deploy OSPF Neighbor Authentication

You should consider configuring a router for OSPF authentication if that router meets any or all of these conditions:

- It is conceivable that the router might receive a false route update.
- If the router were to receive a false route update, your network might be compromised. This is almost a certainty.
- You deem it necessary as part of your network security architecture.

Remember that if you configure a router for OSPF authentication, you also need to configure the neighboring routers for authentication as well.

How OSPF Authentication Works

When OSPF authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives. This is accomplished by the exchange of an authenticating key (sometimes referred to as a *password*) that is known to both the sending and the receiving router. Following are two types of OSPF neighbor authentication used:

- Plaintext authentication
- Message Digest Algorithm Version 5 (MD5) authentication

Both forms work in essentially the same way, with the exception being that MD5 sends a message digest instead of the authenticating key itself. The message digest is created using the key and a message, but the key itself is not sent, preventing it from being read while it is being transmitted. Plaintext authentication sends the authenticating key itself over the wire. Plaintext authentication is not recommended for use as part of your security strategy as a means of protecting against malicious attacks. The primary use of plaintext authentication is to avoid accidental changes to the routing infrastructure. Using MD5 authentication, however, is a recommended best practice.

As with all keys, passwords, and other security secrets, you must closely guard the authenticating keys used in neighbor authentication. The security benefits of this feature are reliant upon your keeping all authenticating keys confident.

Plaintext Route Authentication

Each participating neighbor router must share an authenticating key. This key is specified on each router during the configuration of OSPF. You can specify multiple keys with OSPF. For example, you can have a different key for each WAN interface on a router running OSPF. The caveat is that the neighbor router off each interface must have a matching key configured on the receiving interface, as shown in Figure 8-5.

In general, when a routing update is sent, the following authentication sequence occurs:

1. A router sends a routing update with a plaintext authentication key (that is, password).
2. The receiving (neighbor) router checks the received authentication key against the same key stored in its own memory.
3. If the two keys match, the receiving router accepts the routing update packet. If the two keys do not match, the routing update packet is rejected.

MD5 Route Authentication

MD5 authentication works similarly to plaintext authentication, except that the key is never sent over the wire. Instead, the router uses the MD5 algorithm to produce a message digest of the key (also called a *hash*). The message digest is then sent instead of the key itself. This ensures that nobody can eavesdrop on the line and learn keys during transmission.

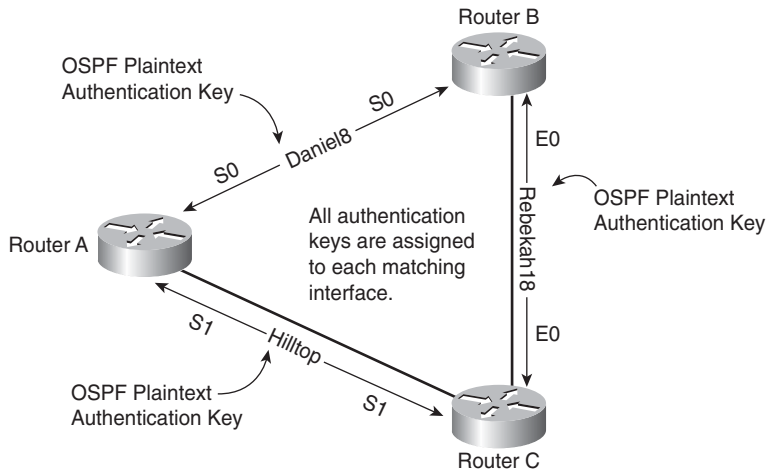


Figure 8-5 OSPF Routing Authentication

One difference between plaintext and MD5 authentication is that MD5 does support defining more than one key per interface (and uses the key number to differentiate the keys). One implication of potentially having more than one key defined is that key management and especially changing keys can be more graceful with MD5 (define a new key on one router [they become out of sync but continue using the old key]; define the new key on the other router [they become in sync and use the new key]; remove the old key from both routers). Contrast this operation to plaintext, where you define a new key on one router, and they become out of sync and terminate the neighbor relationship until the new key is defined on the second router.

Authenticating with Other Dynamic Routing Protocols

Although OSPF is an excellent reliable and secure protocol, it is not the only dynamic routing protocol that you can use. Most modern routing protocols can perform route authentication between neighboring routers to protect the integrity of your network. OSPF was presented as a baseline to introduce the concepts and benefits available to you; if you want to learn more about OSPF, check out the following definitive text on the subject:

OSPF Network Design Solutions, Second Edition
www.ciscopress.com/bookstore/product.asp?isbn=1587050323

Chapter Summary

This chapter discussed ways and places in which you can use a router with a deeper purpose than it might have been implemented with. To this end, the chapter examined how you can use a router to prescreen your network as a *choke point* of entry. The next level was to have the router act as a more advanced packet inspection tool through the use of the Cisco IOS Firewall Feature Set coupled with the intrusion detection feature. Both of these advanced technologies are not a replacement for dedicated devices of the same kind; however, they do offer a higher level of security in your network by adding additional layers of inspection and protection.

Next, the chapter focused on some of the more fundamental methods you can use immediately to secure the router itself. This information was presented in a real router configuration file, thus giving you a point of reference when comparing your router configurations with the suggestions provided here. The chapter concluded with an introduction to securing the routing updates within your network and the best practice methods to do so.

You can find additional resources on security at the following locations:

- **“Increasing security on IP Networks”**: An old but essential document on some of the essentials to security and IP-based networks:
www.cisco.com/en/US/docs/internetworking/case/studies/cs003.html.
- **Cisco Security Intelligence Operations**: An online list at the Cisco website of all its security advisories, including tutorials and details about how to protect yourself from some of the worst vulnerabilities on the Internet today (Cisco.com account required for some features) available at <http://tools.cisco.com/security/center/home.x>.
- **The BRST - Border Router Security Tool**: A web-based utility for generating secure configuration files for Cisco routers in a border configuration. The administrator fills out a web form, clicks submit, and receives a router config file:
<http://sourceforge.net/projects/borderroutersec/> or if you want to try it out already on a web server for you.
- **BRST - Border Router Security Tool Questionnaire**: A web-based utility for generating a secure configuration for Cisco routers. It is primarily designed to be used for border routers in small to medium-sized companies but the concepts can be applied to larger internal routing infrastructures: <http://borderroutersec.sourceforge.net/>.

Chapter Review Questions

1. Because every company that connects to the Internet has a router, should you deploy security on those routers?
2. What is the value of edge routers being used as choke points, and how effective can they be in increasing your network's security?
3. Which four features from classic IOS Firewall features have been implemented in the Zone Based Policy Firewall?
4. What are the two major changes to the way you configure IOS Firewall Inspection, as compared to the Cisco IOS Class Firewall?
5. Can the Cisco IOS IDS have multiple points of packet inspection?
6. Temporary access control lists have timers associated with them. Define how they function based on protocol (ICMP, UDP, and TCP).
7. What is the difference between atomic and compound signatures?
8. What happens when an attacker uses chargen and echo together? How would you stop this from occurring in a Cisco router?

This page intentionally left blank

IPsec Virtual Private Networks (VPNs)

*Change is life giving; it helps us grow into someone greater than we already are.—
Success Stories*

By the end of this chapter, you should know and be able to explain the following:

- The difference between the different types of VPNs
- The benefits and goals of VPN technology and how it should be deployed
- Where the encryption modes are and the functions they play in securing VPNs
- The protocols used during the operation of an IPsec VPN

Answering these key questions will enable you to understand the overall characteristics and importance of network security through the use of several different types of VPNs. By the time you finish this book, you will have a solid appreciation for network security, its issues, how it works, and why it is important.

Workers today are more mobile than ever and are accessing information through laptops and other mobile devices such as smartphones. For telecommuters, Gartner Dataquest predicted that in 2010, 29 percent of workers were telecommuters, up from 27.5 percent in 2009. The number of mobile workers who work from home or other locations continues to rise as well. Mobility enables workers to maintain their productivity, no matter when, where, or how they work.

As connectivity grows and personal mobility increases, the need for networks to adapt and provide services also continues to increase. Users do not understand the security concerns for the remote services that they demand for productively, regardless of location. Users traveling to other countries, in airports, customer sites, and so on demand the ability to connect to corporate resources to fulfill their jobs. With the increased levels of connectivity from T1s and wireless in airports, to Wi-Fi hot spots, and customers with high-speed connections, those people who are responsible for maintaining networks are faced with some difficult decisions. How should they provide the required IT services to users, regardless of their location, in a secure and reasonable manner?

Technology has evolved, and the leading solution for these demands is Internet Protocol Security Protocol (IPsec) encrypted *virtual private networks (VPN)*. Occasionally, a technology's name accurately reflects its function, and this is the case with VPNs.

This chapter discusses the use of VPNs, how they function, the encryption provided by IPsec, and how these technologies can ensure your network's security is maintained while increasing available services to your customers. Everyone has customers to whom they provide some degree of service, regardless of the field. However, for VPNs, *customers* can be defined as anyone with the business need to securely connect to the corporate network to access resources. Customers can be mobile users (sales, system engineers, and so on), power users going online all the time, executives conducting your company's affairs while out of the office, or business partners picking up or dropping off important information. *Resources* are defined here as any device not directly accessible from the Internet; these resources might include email servers, file servers, Citrix servers, or network devices.

The information accessed by mobile workers is not simply limited to business information. Workers from the Millennial—Generation Y group (those born after 1980) typically use the same mobile device to access both personal and professional information. Of the estimated 14 million telecommuters, 69 percent of them report that they use whatever device, software, or site they want, regardless of corporate policies. To continue to foster innovation, enable productivity, and meet the needs of the mobile workforce, companies must adapt to the changing trends in mobility.

Note The National Institute of Standards and Technology (NIST) created AES, which is a new Federal Information Processing Standard (FIPS) publication that describes an encryption method. AES is a privacy transform for IPsec and Internet Key Exchange (IKE) and has been developed to replace the Data Encryption Standard (DES). AES is designed to be more secure than both DES and 3DES, so don't use anything DES.

Arguably the hottest topic in data security today, VPNs are full of promise for businesses seeking to lower cost, increase flexibility and scalability, and ensure the security of their communications.

But what exactly does a VPN do, and how can it affect your business drivers—lowering cost, mitigating risk, and increasing revenue? The popularity of VPN technology is directly related to its potential to bring about significant return on investment (ROI). For businesses paying the often staggering costs of private connections via MPLS or Frame Relay, the cost savings associated with deploying VPNs to replace these costly connections is significant. To understand the value of a VPN to your business, you might want to consider the benefits that VPNs most often bring:

- Site-to-site VPNs can take the place of expensive WAN telco circuits by replacing private line services with site-to-site VPNs that use the Internet instead to connect remote sites.

- Remote access VPNs enable employees who work from home or are out of the office to remain securely connected to organization resources.

If your organization is making significant recurring investments in either WAN telco circuits, a VPN can provide an alternative approach with a big payoff in cost savings and flexibility.

Before entering into a technical overview of the components and possibilities involved in deploying a VPN, you must firmly understand the core operations of VPNs. Analogies work well because they introduce people with vastly different levels of knowledge and experience to a complex subject.

Analogy: VPNs Securely Connect IsLANDs

Your network (LAN) is an island of sanity, order, and user services in an unpredictable ocean known as the Internet. You know thousands of other islands exist within this ocean; when you want to travel from island to island, you would hop on a ferry and travel to the next island that just also happens to be that website or the latest smartphone you had your eye on.

Now, you are on this ferry (using TCP/IP) traveling over the ocean (Internet) to reach something on an island (LAN) that is going to provide you with some sort of service (website). This makes perfect sense, right? Now, how many other people do you see on that ferry—perhaps a few, or perhaps many thousands? The potential problem is that you have no security or privacy traveling from island to island; other people can see everything you see. Now, if you were reading the latest news on www.foxnews.com, who cares if you do not have privacy? However, if you were going to your company's island to check on something, this lack of privacy can have serious ramifications. Do you want anyone looking over your shoulder as you put in your credit card number to make a purchase or to upload the latest sales figures to the corporate server?

Because you are traveling on the worldwide ocean that is the Internet, you have no control over the wires, fiber, routers, or switches that make up the Internet. Nor do you get any guarantees of any sort. In other words, you might reach some website or other server, but there are no guarantees. Remember, connecting to the Internet is a privilege and not a right! Having no control of the Internet means that you are susceptible to security issues, and this becomes especially true if you want to connect two private networks using a public resource such as the Internet, and you want to do this because it is a great cost saver.

Note Once, when conducting a network assessment of a customer's network, I observed that the company had no firewalls at any of its four sites, which were all connected directly to the Internet. This is a serious concern, but what struck me as a real issue is that this customer had configured Microsoft servers at each of the locations to trust one another over the public Internet! All a hacker would have had to do was hijack that trust, and the network would be totally compromised; it had occurred several times but the company refused to

make a change. I had to shake my head in disbelief—do not let this happen to you! Use VPNs to protect your network!

As the person in charge of connecting your island to another, you are directed to connect your island with a new one your organization is getting ready to open. Your island decides to build a bridge to this other island so that there is an easier, more secure, and direct way for people to travel between the two. It is expensive to build and maintain this bridge, even though the island you connect with is close. But the need for a reliable, secure path is so great that you do it anyway.

This situation is a lot like having a private WAN. The bridges (private lines) are separate from the ocean (Internet), yet they can connect the islands (LANs). Many companies have chosen this route because the need for security and reliability drives the connection from their remote offices to their main office.

Your island would like to connect to a second island that is much farther away, but you decide that the cost to build a bridge is simply too high to justify. You quickly learned that, if the offices are far apart, the cost could be prohibitively high, just like trying to build a bridge that spans a great distance. However, the need is still there.

Note Many businesses have a tendency to allow IT to drive the evolution of their business, and although this is appropriate for some, most businesses must reverse this thinking. The needs of the business should drive the evolution of a company's IT infrastructure. To me, this is a fundamental truth because businesses are not in business to build a big IT department or network! Nerds, take note: The days of blindly spending money are over, and reality has unfortunately returned in the form of the proven business model.

Are you wondering when VPNs are going to fit into this analogy? You have established that you need increased security, and the first option was to build a bridge; however, that is too expensive. You could give a submarine to everyone who needs the ability to privately and securely travel between islands. A submarine is a perfect analogy for a VPN because, like a submarine, VPNs have the following amazing properties:

- They can be very fast.
- They are easy to take with you.
- They can hide you from others.

It might not be easy to take a submarine with you; however, I am sure you understand this analogy. There are several different ways to implement VPNs, and the following sections examine the three types of VPNs. Another good analogy would be the concept of the *Stargate* portals from Hollywood. You must get the symbols right on both sides (the SA for VPN), and you must have a stargate on the other side that is “on” for the hyperspace

tunnel to form (the VPN tunnel)....I bet you thought that nothing intelligent comes out of Hollywood; I know I did!

VPN Overview

A VPN is an encrypted network connection that uses a secure tunnel between endpoints via the Internet or other network, such as a WAN. In a VPN, dial-up connections to remote users and leased-line or Frame Relay connections to remote sites are replaced by local connections to an Internet service provider (ISP) or other service provider's Point of Presence (POP). The increasing prevalence of Internet broadband connections to small remote offices and homes makes the use of cheaper access to the Internet attractive. As discussed, after the initial investment in VPNs, the cost to add more sites or users is minimal.

VPNs enable each remote user of your network to communicate in a secure and reliable manner using the Internet as the medium to connect to your private LAN. A VPN can grow to accommodate more users and different locations much easier than a leased line. Scalability is a major advantage that VPNs have over typical leased lines. Unlike leased lines, where the cost increases in proportion to the distances involved, the geographic locations of each office matter little in the creation of a VPN.

A VPN enables a private intranet to be securely extended through IPsec encryption across the Internet or other network service, facilitating secure e-commerce and extranet connections with mobile employees, business partners, suppliers, and customers. Following are three main types of VPNs:

- **Remote Access VPNs:** Enables remote users to securely connect to a central site across the Internet. This type of VPN is a user-to-LAN connection that enables employees who need to do so to connect to the corporate LAN from the Internet. Their systems use special VPN client software that enables a secure link between themselves and the corporate LAN. Typically, a corporation that wants to set up a large remote access VPN provides some form of Internet dial-up account to its users using an ISP. The telecommuters can then connect to the Internet and use their VPN client software to access the corporate network. A good example of a company that needs a remote access VPN would be a large firm with hundreds of salespeople in the field. Remote access VPNs are sometimes referred to as soft (as in software-based) VPNs, virtual private dialup networks (VPDN), or client-based VPNs. Some of the fastest growing uses of them are as follows:
 - VPN-capable mobile devices such as smartphones or tablets.
 - The Cisco VoIP CIPC (Cisco IP Communicator) SoftPhone application also works well over a VPN, turning your PC into a secure telephone.
- **Site-to-site VPNs:** Used to extend a company's private network to other buildings or sites through the use of dedicated equipment so that remote employees at these locations can use the same network services. These types of VPNs are considered active-

ly connected at all times. Site-to-site VPNs are sometimes referred to as hard (as in hardware-based) VPNs, intranet, or LAN-to-LAN (L2L) VPNs.

- Extranet VPNs:** Enable secure connections with business partners, suppliers, and customers for the purpose of e-commerce. Extranet VPNs are a type of site-to-site VPN with the addition of firewalls to protect the internal network. A good example would be companies that work closely with suppliers and partners to achieve common goals such as supply and demand relationships—for example, when one company has a demand for supplies and the supplier fulfills the demand based on the company’s needs. Working across an extranet, these two companies can share information more quickly, and the firewall rules ensure that access is happening only to the shared resource.

All these VPNs aim to provide the reliability, performance, quality of service, and security of traditional WAN environments using lower cost and more flexible ISP or other service-provider connections. Figure 9-1 illustrates the three types of VPNs.

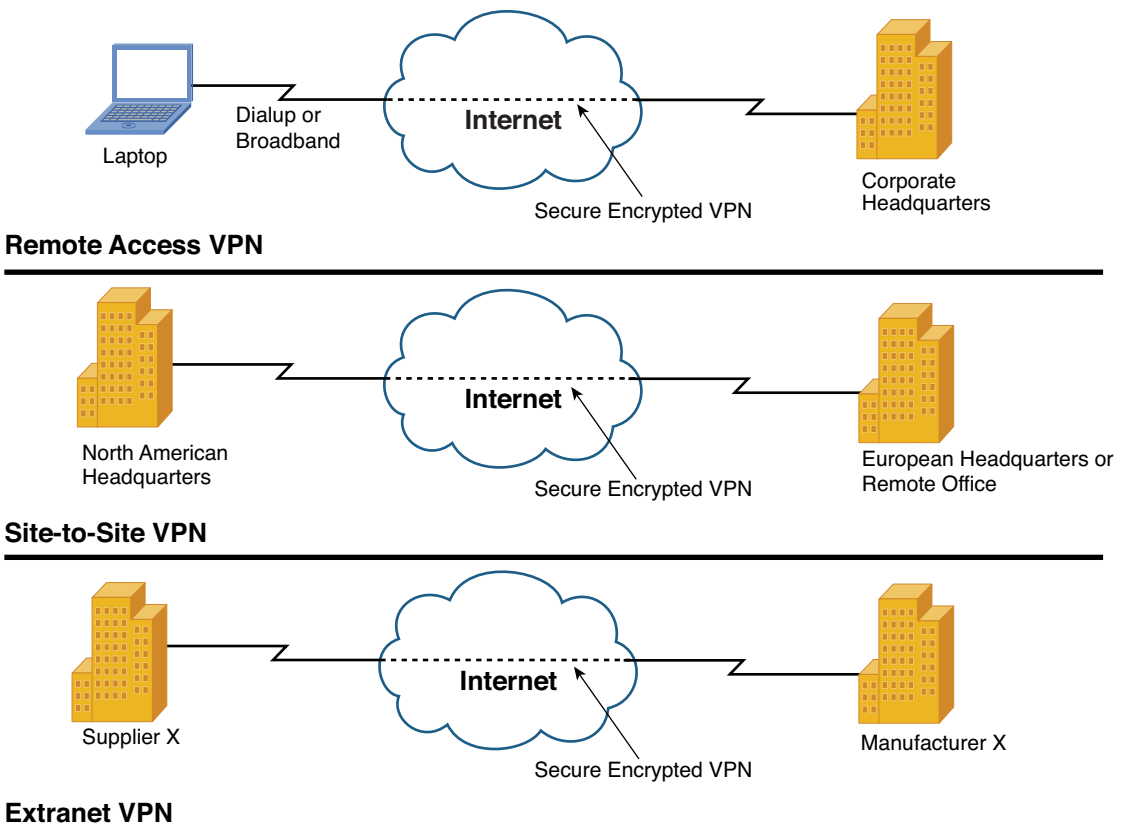


Figure 9-1 Types of VPNs

In Figure 9-1, all the VPNs use the Internet. You can also use VPN technology within your network to provide an additional layer of security to control access to sensitive information, systems, or resources. For example, you can use VPN technology to limit access to financial systems to certain users or to ensure that sensitive or confidential information is sent in a secure way. In this scenario, VPNs can encrypt and further secure traffic to sensitive systems. The following section discusses the placement of VPNs and the specific associated benefits.

VPN Benefits and Goals

A well-designed VPN can greatly benefit any company. Some of the benefits of implementing a VPN in your network include the following:

- Before the advent of VPN technologies, employees in remote locations would need to get an expensive connection such as a Frame Relay T1 to reach their company's network. You might want to reduce telecom costs with local broadband connections to the Internet through which users use a VPN client. Depending on the number of employees in the field, this alone can be a huge cost savings. For many smaller companies with limited financial breathing room, VPNs can be a practical solution when remote access is needed.
- You want to increase the productivity of your users by enabling them to securely access network resources regardless of their geographic location.
- You want to reduce the operational costs associated with dedicated WAN connections by replacing them with direct Internet connections such as business class broadband, through which remote sites connect via a site-to-site VPN.
- You want to simplify your network's topology by adding VPNs strategically throughout your network.
- You want greater flexibility in deploying mobile computing, telecommuting, and branch office networking; easier e-commerce and extranet connections with business partners, suppliers, and customers' external Internet access; and internal intranet and extranet access provided using a single secure connection.
- You want to reduce office costs by having users work from home three days a week. Home users typically have higher production and less stress.

Before implementing a VPN, you should spend some time contemplating what you want to accomplish with your VPN. During this exercise, before choosing a solution provider or hardware and software, you should consider which features are most important. Security, which is mentioned later, is one of the most important features of your VPN.

Caution It is possible to have unencrypted VPNs that rely on some other type of encryption or routing for security—for example, MPLS VPNs. Only under specific circumstances

are these VPNs the appropriate solution for a network. Best practice dictates that you always encrypt your traffic over a VPN; failure to do so could be disastrous, and the responsibility will rest squarely on your shoulders.

VPN Implementation Strategies

VPN implementation strategies are extremely varied because every vendor these days has a “VPN solution” for you! Some of the solutions are what they claim to be, and others have raised concerns among the security community, as discussed in Chapter 8, “Router Security.” This section looks at some of the different potential components available from Cisco, and how you can use multifunction devices such as a Cisco Adaptive Security Appliance (ASA) to fulfill a VPN role:

- **Firewalls:** Firewalls are crucial to the security of your network. If you did not have a firewall in place before reading Chapter 7, “Firewalls,” you probably do now. Today, all Cisco firewalls support the combining of VPNs with stateful packet inspection (SPI). Solutions range from standards-based site-to-site VPNs leveraging the Internet Key Exchange (IKE) to IP security (IPsec) VPN standards. Cisco ASA firewalls encrypt data using 56-bit Data Encryption Standard (DES), 168-bit Triple DES (3DES), or up to 256-bit Advanced Encryption Standard (AES) encryption. Obviously, you want to use AES because it is the most secure! An amazing piece of technology, the Cisco ASA Firewall combines dynamic Network Address Translation (NAT), content filtering, stateful packet inspection, firewall, and VPN termination capabilities into a single piece of hardware. Instead of using Cisco IOS Software, the ASA has a highly streamlined OS that trades the capability to handle a variety of protocols for extreme robustness and performance by focusing on security services.
- **VPN-capable routers:** Cisco routers can be upgraded to have the capability to use VPNs. These upgrades come in some form of the following, depending on the router model in question: IOS, memory, or dedicated VPN hardware. You can gain some unique features with the provision of scalability, routing, security, quality of service (QoS), and dynamic multipoint VPNs. Based on Cisco IOS Software, there is a router suitable for every situation, from small office/home office (SOHO) access through central-site VPN aggregation, to large-scale enterprise needs.
- **Client software:** Simple to deploy and operate, the Cisco VPN Client establishes secure, end-to-end encrypted tunnels to the VPN devices listed here. With this thin design, IPsec-compliant software can be preconfigured for mass deployments, and the initial logons require little user intervention. The client software is available for the following operating systems: Windows 32 and 64 bit, Linux, Solaris, and Mac OS. Cisco has also worked with other manufacturers to include VPN software in many of today’s mobile devices.

Depending on the type of VPN (remote access or site-to-site), you must use specific software and hardware components to build your VPN. However, you should also consider the following:

- **Manageability:** Manageability of a VPN concerns the amount of effort needed to successfully maintain the established network connectivity. Specifically, *PC Magazine* rates manageability by the “ease-of-use factors for remote and local management options, including whether the device provides a browser-based interface or command line access” (*PC Magazine*, 2002).
- **Reliability:** Obviously, if the VPN software or hardware is unavailable when you need it, you are losing productivity and probably money. When choosing a solution, you should request up-time statistics for comparison and build a redundant solution if business needs dictate.
- **Scalability:** As a company’s business grows, so does its IT requirements. To quickly and cost-effectively grow your VPN infrastructure, you need to choose a solution that has scalability in mind. The last thing an IT manager wants to do is start from scratch and replace a VPN infrastructure because of a bottleneck in its growth potential.

When selecting the right device to provide VPN services to your network, you must be aware of the limitations.

Split Tunneling

Many VPN users are already behind firewalls, and they need to access resources only through a VPN. Traditional VPNs do not enable users to also access network resources on their local segment while they connect to their corporate VPN at the same time. This becomes an issue when, for example, these users must access a system via a VPN *and* print to a local network printer. To correct this potential problem, a feature has been introduced known as *split tunneling*.

Split tunneling occurs when remote VPN users or sites are allowed to access a public network (the Internet) at the same time that they accesses the private VPN, without placing the public network traffic inside the tunnel first. This is not always the best feature to enable, however, because it could enable an attacker to compromise a computer connected to two networks. Figure 9-2 illustrates an overview of how split tunneling works.

Overview of IPsec VPNs

IPsec has become the de facto standard for creating VPNs in the networking industry providing excellent security. Several vendors have implemented it and, because the Internet Engineering Task Force (IETF) has defined IPsec in an RFC, interoperability between vendors makes IPsec the best option for building VPNs. IPsec offers a standard means of establishing authentication and encryption services between peers. IPsec is an IETF standard; furthermore, it is FIPS-compliant when used with AES encryption making it the

best option for deploying VPNs. IPsec acts at the network layer of the OSI model, protecting and authenticating IP packets between participating IPsec devices (peers), such as Cisco routers or firewalls. IPsec provides the following network security services:

- **Data confidentiality:** The IPsec sender can encrypt packets before transmitting them across a network. If hackers cannot read the encrypted data, it is of no use to them.
- **Data integrity:** The IPsec receiving endpoint will authenticate all packets sent by the IPsec sending endpoint to ensure that the data has not been altered during transmission.
- **Data origin authentication:** The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- **Anti-replay:** The IPsec receiver can detect and reject replayed packets.

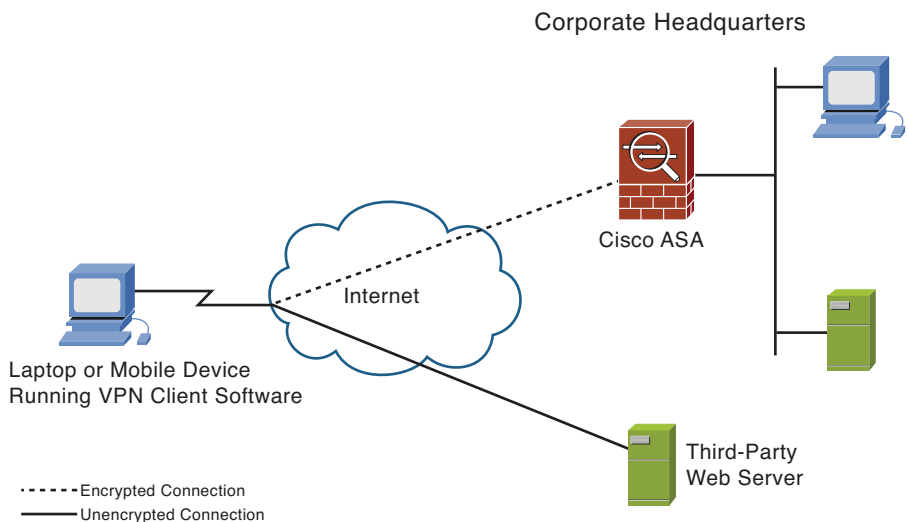


Figure 9-2 *Split Tunneling Overview*

IPsec protects sensitive data that travels across unprotected networks, and IPsec security services are provided at the network layer; therefore, you do not need to configure individual workstations, PCs, or applications. This benefit can provide a great cost savings. Rather than providing the security services that you do not need to deploy and coordinate security on a per-application, per-computer basis, you can simply change the network infrastructure to provide the needed security services. For example, you can load updated VPN client software on a Cisco ASA, which can cause an outdated client to download and install the latest version before connecting. This is a huge time and cost savings if the alternative is manually installing updates on user laptops, plus it enables you to maintain excellent security.

IPsec provides enhanced security features, such as better encryption algorithms and more comprehensive authentication. Corporate networks connected to the Internet can enable flexible and secure VPN access with IPsec. With IPsec technology, customers can now build VPNs over the Internet with the security of encryption protection against wire tapping, eavesdropping, or other attacks that intrude on private communications.

Note Only IPsec-compliant systems can take advantage of this protocol. Also, all devices must use a common key, and each network's firewalls must have similar security policies set up.

IPsec provides authentication and encryption services to protect unauthorized viewing or modification of data within your network or as it is transferred over an unprotected network, such as the public Internet. IPsec can encrypt data between various devices, such as the following:

- Router to router
- Firewall to router
- Firewall to firewall
- User to router
- User to firewall
- Mobile device to firewall

Mobility is increasing exponentially as smartphones and tablet devices gain more and more adoption by the industry. In many cases, users are driving IT departments to give them access from wherever, whenever, and however necessary for them to access the network. Consider the practical applications of an IPsec-based VPN client configuration:

- The VPN client can be preconfigured for mass deployments.
- Requires little user intervention for initial logins, can also be tied in to current internal authentication methods for a single sign-on service for users.
- Supports Cisco Easy VPN capabilities, decreasing network security policy configuration at the remote location; also enables ease of management by your IT staff.
- All current operating systems are supported to include 64 bit.
- Security policies can be centralized and customized as needed to meet your security posture.

IPsec is a framework of open standards defined by the IETF. IPsec provides security for transmission of sensitive information over unprotected networks, such as the Internet. Figure 9-3 shows the three most common types of VPNs.

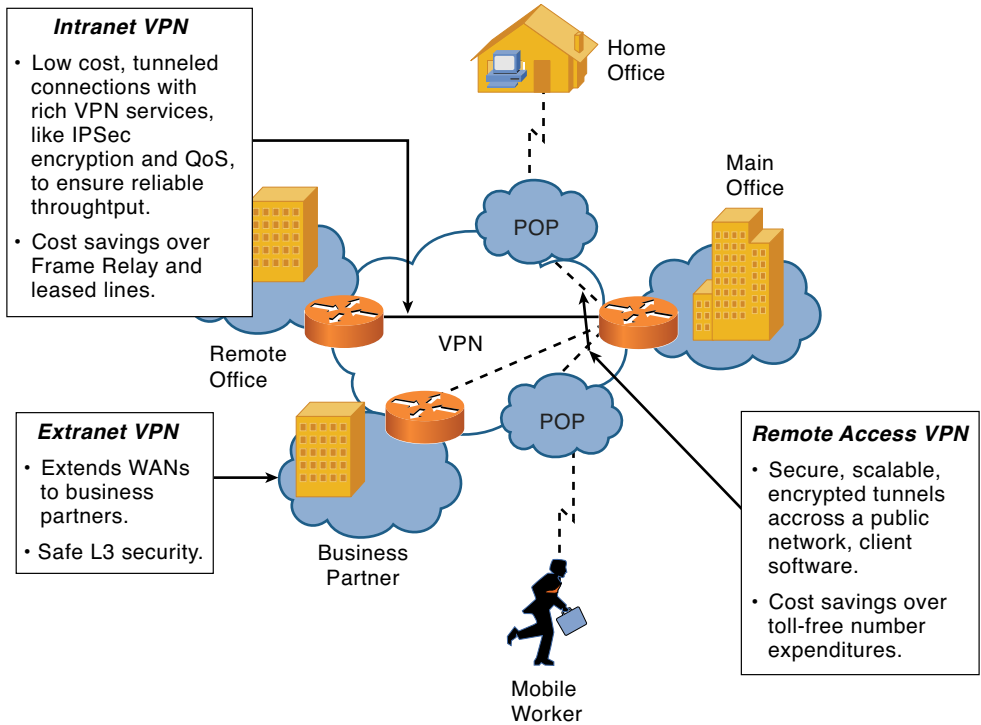


Figure 9-3 VPN Connectivity Overview

Authentication and Data Integrity

To establish trust, *authentication* verifies the identity of the two VPN endpoints and the users sending traffic through the VPN. An endpoint could be a VPN client, firewall, mobile device, or router. Authentication is a process of IPsec that occurs after data encryption and before decryption on the receiving end. It is a necessary function within IPsec to ensure that both the sending and receiving parties are who they claim to be. With IPsec, each peer must be manually configured with a preshared key (usually agreed upon before a connection attempt is made).

Note Users can also be authenticated via digital certificates to an Active Directory server, or you can also require a machine to have a digital certificate to even begin the connection process.

Data integrity is another function within IPsec. *Integrity* means that the packet that the receiving party received has not been altered during transmission. This is achieved via the use of a one-way hash algorithm. A one-way hash is the equivalent of an encrypted checksum. After the sending party encrypts and authenticates a packet, a one-way hash is run on the value of the entire packet. A hash is interesting in that its result will always be a

fixed size, regardless of the input. This is another security mechanism so that hackers cannot know the input field size. The one-way hash creates an encrypted field appended to the message. On the receiving end, the one-way hash value is pulled from the packet, and the receiving end runs its own one-way hash. Because the hash is run on variables within the packet such as time sent, number of bytes, and so on, both ends' hash value must be the same—meaning that the packet has not been tampered with. If the values are different, the packet is discarded and IPsec renegotiates its security parameters.

Tunneling Data

Tunneling is what VPNs rely on to create a private network over the Internet. Basically, this is the process of taking an entire packet of data and encapsulating it within another packet before sending it over a network. This act of encapsulating one packet of data into another is what happens when a packet is encrypted and transmitted through a tunnel. The network must understand the outer packet's protocol for the packet to be routed across the network. Tunneling data within a VPN requires three different protocols to work:

- **Data:** The original data packet, usually an IP packet, which is to be encrypted and transmitted through the VPN. Perhaps the data to be sent through the tunnel is an FTP file transfer.
- **Encapsulating protocol:** The protocol (GRE, IPsec, PPTP, and L2TP) wrapped around the original data (that is, encapsulated). IPsec is the de facto VPN standard used as the encapsulating protocol at this stage, and it enables the data packet to be encrypted and protected.
- **Carrier protocol:** The protocol the network uses over which the encrypted VPN information travels. The original packet (data) is encapsulated inside the encrypting protocol, which is then put inside the carrier protocol's header (usually IP) for transmission over the public network through the use of a routing protocol such as OSPF.

Tunneling works well with VPNs. It enables you to use protocols not supported on the Internet inside an IP packet, and it can still be safely sent. At the beginning of a VPN tunneled transmission, a data packet from the source LAN is encapsulated with new header information that enables intermediary networks to recognize and deliver it. After this is done and the transmission is complete, the tunneling protocol "header" is stripped off at the other end of the tunnel, and the original packet is transferred to the destination LAN for delivery.

Although tunneling enables data to be carried over public networks, tunneling alone does not ensure privacy. To secure a tunneled transmission against any interception and tampering, all traffic over the VPN is encrypted. In addition, VPNs typically include additional features, such as firewalls at the perimeters.

In site-to-site VPNs, the encapsulating protocol is usually IPsec or generic routing encapsulation (GRE). GRE includes information about what type of packet you are encapsulating and about the connection between the client and server. The difference depends on the level of security needed for the connection, with IPsec being more secure and GRE having

greater functionality. IPsec can tunnel and encrypt IP packets, whereas GRE can tunnel IP and non-IP packets. When you need to send non-IP packets (such as IPX) through a tunnel, use IPsec and GRE together.

VPN Deployment with Layered Security

Security in depth is critical, as we have mentioned time and time again, when looking at putting all the various solutions together with regard to VPN. Figure 9-4 demonstrates layering your security.

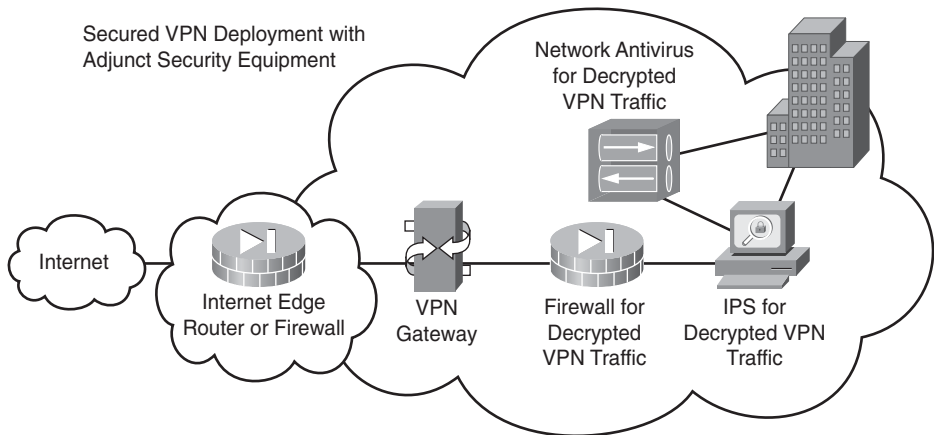


Figure 9-4 *VPNs with Layered Security*

This figure shows several devices; consider their placement and role in the defense of your network:

- **Internet:** All visitors to your websites, email, and VPNs come from the Internet, which is where the dragons live.
- **Internet edge router or firewall:** This device can be many things. Its primary role is to provide the IP connection to the Internet via your service provider. Then it can prescreen traffic and even act as an initial firewall. The key point is to have this device do something to protect your network; it's there, so use it to make another layer of defense.
- **VPN gateway:** This device is what all the VPN clients are terminating against. Its role is to provide the processing power to do all the necessary encryption and decryption. This is also the device that builds permanent VPN tunnels to remote sites and business partners.
- **Firewall for decrypted traffic:** When traffic is trying to enter your network, by the time it gets to this point, it is either regular Internet traffic (unencrypted, of course)

or it is traffic sent via a VPN that is unencrypted. Regardless of the type, it must be subjected to the rules of your firewall and stateful packet inspection. This is the device that perform both functions.

- **IPS for decrypted traffic:** Inbound traffic can now be subjected to inspection by an intrusion prevention or detection system. You know that traffic has passed the firewall rules to reach this, so now look into the packets to ensure there are no embedded attacks within the packets.
- **Network antivirus detection:** This device can perform more than just antivirus. It might also be checking for botnets and redirect host to determine whether it meets the company security policy

IPsec Encryption Modes

IPsec has two encryption modes: tunnel and transport. Each mode differs in its application and in the amount of overhead added to the passenger packet. These different modes of operation are summarized briefly in that tunnel encrypts the packet header and the payload of each packet, whereas transport encrypts only the payload.

IPsec Tunnel Mode

This is the normal way in which an IPsec VPN is implemented between two ASA firewalls (or other security gateways) connected over an untrusted network, such as the public Internet. All discussions involving IPsec are about the tunnel mode because this is the most secure method and is the industry standard. Tunnel mode enables IPsec to encrypt, which then encapsulates the entire IP packet securing the data through the encryption. Because it encapsulates or hides the packets to be successfully forwarded, the encrypting routers themselves own the IP addresses used in these new headers because they are the next hop routing addresses needed. Using tunnel mode results in additional packet expansion of approximately 20 bytes per associated IP header; a new IP header must be added to the original packet, as shown in Figure 9-5.

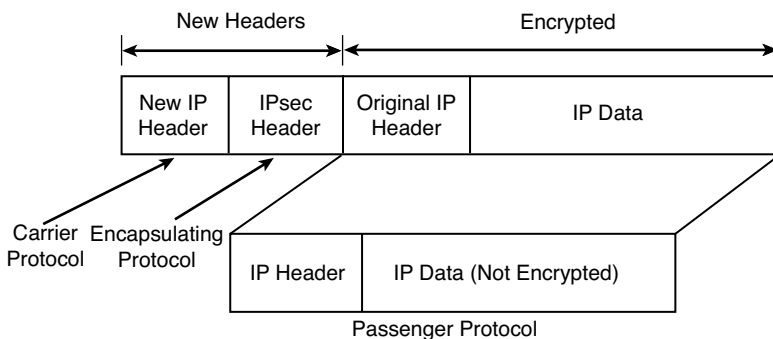


Figure 9-5 *Tunnel Mode*

In tunnel mode, IPsec encrypts the entire packet and writes a new IP header onto the new encrypted packet, which masks the original source and destination IP address information. This information will be used when the packet is decrypted at the other VPN tunnel endpoint.

Transport Mode

This method of implementing IPsec is typically done with L2TP to enable the authentication of Windows VPN clients. Chapter 6, “Security Protocols,” covers this concept, so this chapter focuses on IPsec and tunnel mode. In transport mode, only the data portion of a packet is encrypted, making it less secure than tunnel mode. Tunnel mode is inherently more secure than transport mode (because the entire original packet is encrypted, not just the payload as in transport mode), as shown in Figure 9-6.

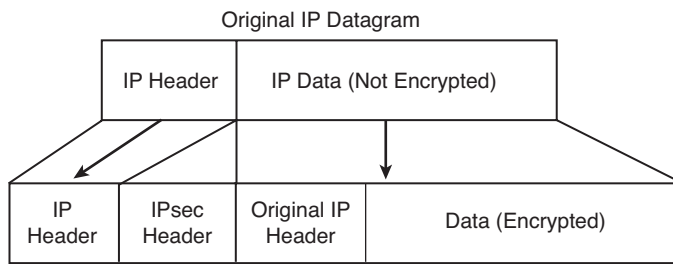


Figure 9-6 *Transport Mode*

IPsec Family of Protocols

IPsec works on the network layer of the OSI model—securing all data that travels between the two endpoints without an association to any specific application. IPsec accomplishes these goals through the use of three main protocols that combined form a cohesive and secure standards-based framework ideally suited for VPNs. The three protocols described in the IPsec standards have various functions within them, as detailed in the list that follows:

- **Internet Security Association Key Management Protocol (ISAKMP):** This protocol is used during the initial phase of negotiating the IPsec connection to establish the VPN between VPN endpoints or peers; Oakley defines the method to establish an authenticated key exchange. This method can take various modes of operation and can also derive keying material via algorithms such as Diffie-Hellman. Within ISAKMP is Internet Key Exchange (IKE), which provides a framework for negotiating security parameters (for example, SA lifetime, encryption type, and so on) and establishing the accuracy of the keys.
- **Encapsulated Security Protocol (ESP):** Provides data confidentiality and protection with optional authentication and replay-detection services. ESP completely encapsulates user data. ESP can be used either by itself or with AH. ESP runs using the TCP protocol on ports 50 and 51 and is documented in RFC 2406.

- **Authentication Header (AH):** Provides authentication and antireplay services (optional). AH provides services to limited portions of the IP header and extended header but does not provide for data encryption by applying a one-way hash to create a message digest of the packet. AH is embedded in the data to be protected and can be used either by itself or with Encryption Service Payload (ESP). (Refer to RFC 2402.) AH has largely been superseded by ESP and is considered deprecated.

Security Associations

Security associations (SA) establish trust between two devices in a peer-to-peer relationship and enable VPN endpoints to agree on a set of transmission rules by negotiating policies with a *potential* peer. Consider a security association such as a contract negotiated enabling the two VPN endpoints to agree to the various parameters for how the VPN tunnel is to be secured.

A security association is identified through an IP address, a security protocol identifier, and a unique security parameter index (SPI) value. The SPI value is a 32-bit number embedded in packet headers.

ISAKMP Overview

Internet Security Association and Key Management Protocol (ISAKMP) is a framework that defines the mechanics of implementing a key exchange protocol and negotiation of a security policy and threat mitigation (for example, DoS and replay attacks). ISAKMP is used for secure exchanges of both SA parameters and private keys between peers in an IPsec environment, and key creation and management, and is defined in RFC 2408.

ISAKMP provides for several methods of key management and provides secure transit of IPsec parameters between peers. It accomplishes this by using similar algorithms used by IPsec for the actual encryption of the data payload. Like IPsec, ISAKMP is not a protocol, but simply an interface to manage various ways of dynamic key exchange. ISAKMP defines various methods—such as digital signatures, certificates, and one-way hash algorithms—to ensure that negotiation of SAs between peers is securely handled.

Currently, the only supported protocol in ISAKMP is the Internet Key Exchange (IKE) protocol. When IKE is actively employed in the encryption process, many features become available to the IPsec communication process. Using public-key cryptography, IKE negotiates security parameters and key exchanges before the IPsec processing begins.

ISAKMP uses UDP port 500 to communicate. In addition, UDP port 4500 is used when NAT is present on a network and because 99 percent of today's networks use NAT, make sure you allow this port through your firewall as well. This port is designated for a specific function referred to as NAT-T, as in *Transparent*, to account for operating IPsec over NAT.

Internet Key Exchange (IKE) Overview

IKE provides negotiation, peer authentication, key management, and key exchange. As a bidirectional protocol, IKE provides a secure communication channel between two devices that negotiates an encryption algorithm, a hash algorithm, an authentication method, and any relevant group information. It uses key exchange based on Diffie-Hellman algorithms, and network administrators can closely tie IKE with policy management systems. To prevent a man-in-the-middle attack—when an attacker sniffs packets from the network, modifies them, and inserts them back into the network.

IKE is the protocol that IPsec uses for completion of Phase 1 of negotiating the VPN tunnel. IKE negotiates and assigns SAs for each IPsec peer, which provides a secure channel for the negotiation of the IPsec SAs in Phase 2. IKE provides the following benefits:

- Eliminates the need to manually specify all the IPsec security parameters at both peers
- Enables you to specify a lifetime for the IPsec SAs
- Enables encryption keys to change during IPsec sessions
- Enables IPsec to provide antireplay services
- Enables CA support for a manageable, scalable IPsec implementation
- Enables dynamic authentication of peers

IKE negotiations must be protected, so each IKE negotiation begins by the peer agreeing on a common (shared) IKE policy. This policy states the security parameters used to protect subsequent IKE negotiations. After the two VPN peers agree on a policy on how to encrypt the tunnel, a security association established at each peer identifies the VPNs security parameters, and these SAs apply to all subsequent IKE traffic during the negotiation. Figure 9-7 illustrates the two modes of operation possible for IKE: main and aggressive. You can see that main is more involved and aggressive consolidates steps—personally, I like main better.

IKE Main Mode

An IKE session begins with the initiator sending a proposal or proposals to the responder. The proposals define what encryption and authentication protocols are acceptable, how long keys should remain active, and whether perfect forward secrecy should be enforced, for example. Multiple proposals can be sent in one offering. The first exchange between nodes establishes the basic security policy; the initiator proposes the encryption and authentication algorithms it is willing to use. The responder chooses the appropriate proposal (assume a proposal is chosen) and sends it to the initiator. The next exchange passes Diffie-Hellman public keys and other data. All further negotiation is encrypted within the IKE SA.

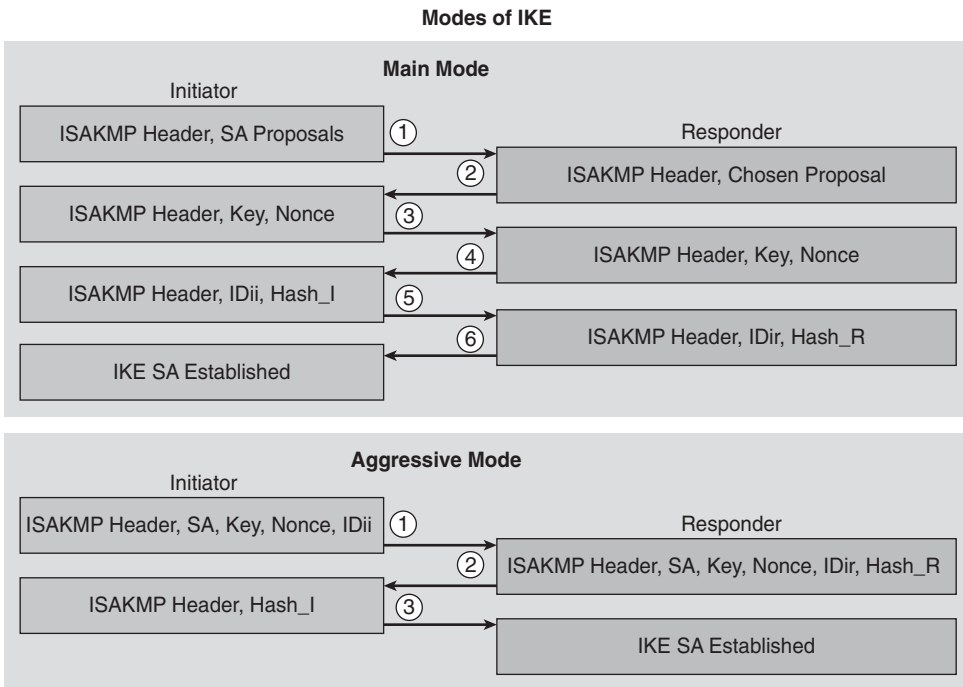


Figure 9-7 Modes of IKE

IKE Aggressive Mode

Aggressive mode squeezes the IKE SA negotiation into three packets, with all data required for the SA passed by the initiator. The responder sends the proposal, key material, and ID and authenticates the session in the next packet. The initiator replies by authenticating the session. Negotiation is quicker, and the initiator and responder ID pass in the clear.

IPsec Security Association (IPsec SA)

IPsec SA is unidirectional and thus requires that separate IPsec SAs be established in each direction. IPsec SA is a two-phase, three-mode procedure. In Phase 1, the basics of the security policy are exchanged. In phase 2, IKE's two modes can be used: main mode and aggressive mode. In Phase 3, the only available mode is called *quick mode*. The third exchange authenticates the ISAKMP session. After the IKE SA is established, IPsec negotiation (quick mode) begins. IPsec negotiation, or quick mode, is similar to an aggressive mode IKE negotiation, except negotiation must be protected within an IKE SA. Quick mode negotiates the SA for the data encryption and manages the key exchange for that IPsec SA.

Both IKE and IPsec use SAs, although the SAs are independent of one another. The security associations define which protocols and algorithms should be applied to sensitive packets and specify the keying material to be used by the two peers. SAs are unidirectional and are established separately for different security protocols (AH and ESP). You can establish IPsec SAs in two ways:

- **Manual SAs with preshared keys:** The use of manual IPsec SAs requires a prior agreement between administrators of the ASA firewall and the IPsec peer. There is no negotiation of SAs, so the configuration information in both systems should be the same for IPsec to process traffic successfully. Manual is easy to configure; however, it is difficult to change preshared keys because the tunnel fails when you do, and the trouble is that preshared keys are usually never changed.
- **IKE-established SAs:** When IKE is used to establish IPsec SAs, the peers can negotiate the settings they will use for the new security associations. This means that you can specify lists (such as lists of acceptable transforms) within the *crypto map* entry, for example, when using a Cisco device.

Note A potential point of confusion is that the acronyms ISAKMP and IKE are both used in Cisco IOS Software to refer to the same thing. These two items are somewhat different, as shown in the next definition.

IPsec Operational Overview

IPsec's main task is to enable the exchange of private information over an insecure connection by negotiating the connection and providing the keys in a secure manner. IPsec uses encryption to protect information from interception or eavesdropping. However, to use encryption efficiently, both parties should share a secret key (password) used for both encrypting and decrypting the information as it enters and exits the VPN tunnel. IPsec uses IKE to establish the secure link, so the VPN forms connecting two endpoints, and data is encrypted securely flowing through the VPN. At a high level, the sequence of events for an IPsec tunnel creation are as follows:

1. One of the IPsec peers receives or generates interesting traffic on an interface that has been configured to initiate an IPsec tunnel when *interesting* traffic is received. Interesting traffic is defined in the endpoint as the type of data to be sent into the tunnel, thus its name.
2. IKE begins to negotiate the VPN security association. IKE either uses main mode or aggressive mode in the creation of an IKE SA between two IPsec peers. Upon success of this step, the ISAKMP SAs are created.
3. IKE now negotiates the IPsec SAs, and the IPsec SAs are created if successful.
4. Data starts passing through the encrypted VPN tunnel with all the encryption done per the parameters in the SAs.

These four seemingly simple steps require some additional explanation for Steps 2 and 3 because they are the critical aspects of the VPN tunnel creation. IPsec operates in two major phases to allow the confidential exchange of a shared secret key, as described in the following sections.

IKE Phase 1

IKE Phase 1 handles the negotiation of security parameters required to establish a secure channel between two IPsec peers. Phase 1 is implemented through the IKE protocol and is primarily concerned with establishing the protection for IKE messages. The sequence of events of IKE Phase 1 is as follows:

1. Phase 1 is the creation of the ISAKMP SA, where peers negotiate and agree upon policies, which are used by IPsec to negotiate and set up the SAs.
2. The key step here is that the VPN peers use their shared secret keys to authenticate with each other using Diffie-Hellman. If either the policies or the keys do not match, Phase 1 fails and the connection halts.

After Phase 1 is complete and a secure channel is established between peers, IKE moves into Phase 2. Figure 9-8 shows the negotiation of the Phase 1 parameters through the use of preshared keys.

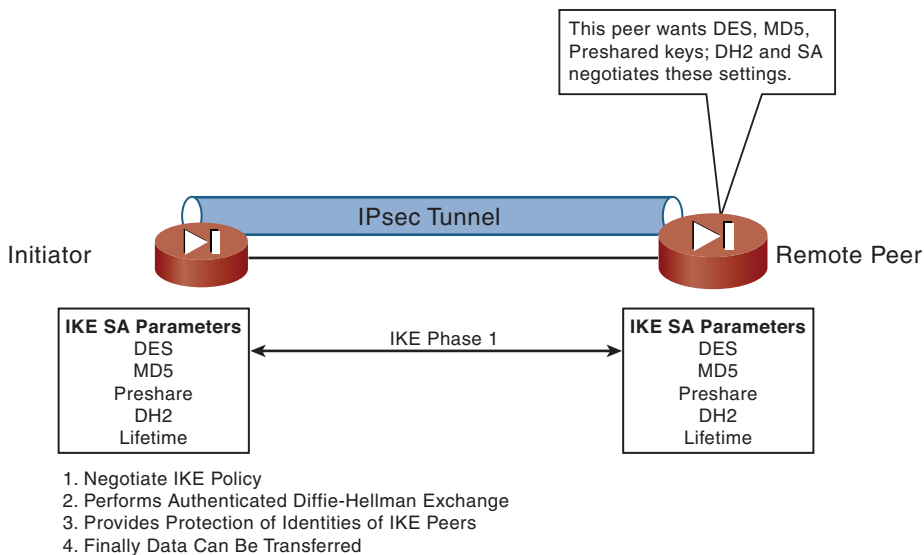


Figure 9-8 IKE Phase 1 Operation

IKE's Phase 1 operation has two modes of operation: aggressive and main mode. Aggressive mode eliminates several steps in the authentication of IKE, reducing it to just three steps, whereas main mode uses the full four steps to authenticate. Although it's

faster, aggressive mode is considered less secure than main mode, for obvious reasons. Cisco devices use main mode by default, but they respond for peers using aggressive mode if configured to do so.

IKE Phase 2

IKE Phase 2 advances the security of the connection by using the secure tunnel established in IKE Phase 1 to exchange the IPsec security parameters required to actually transmit user data (see Figure 9-9).

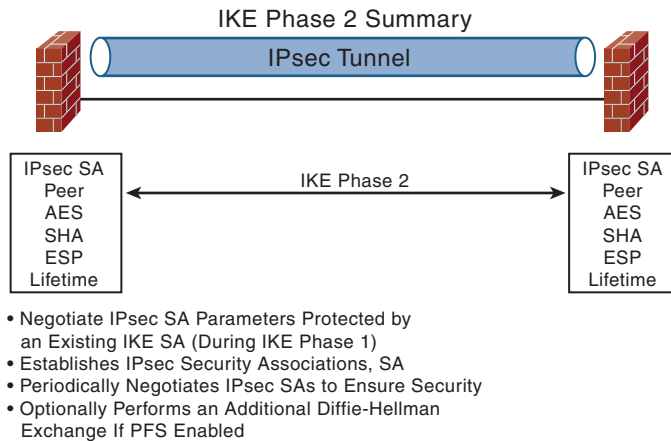


Figure 9-9 *IKE Phase 2 Operation*

In Phase 2, IKE negotiates SAs on behalf of IPsec, according to parameters configured in IPsec. The ISAKMP SA created in Phase 1 protects these exchanges. IKE phase 2 has one mode, called *quick mode*. Quick mode occurs after IKE has established the secure tunnel in Phase 1. It negotiates a shared IPsec policy, derives shared secret keys used by the IPsec security algorithms, and establishes IPsec SAs. Quick mode exchanges nonces that provide replay protection. The nonces are used to generate new shared secret key material and prevent replay attacks from generating bogus SAs.

Quick mode is also used to renegotiate a new IPsec SA when the IPsec SA lifetime expires. Base quick mode refreshes the keying material used to create the shared secret key based on the keying material derived from the Diffie-Hellman exchange in Phase 1.

Perfect Forward Secrecy

If perfect forward secrecy (PFS) is specified in the IPsec policy, a new Diffie-Hellman exchange is performed with each quick mode, providing keying material that has greater entropy (key material life) and thereby greater resistance to cryptographic attacks. Each Diffie-Hellman exchange requires large exponentiations, thereby increasing CPU use and exacting a performance cost. By default, Cisco devices do not have PFS configured.

The secure tunnels used in both phases of IPsec are based on SAs used at each IPsec endpoint. SAs describe the security parameters, such as the type of authentication and encryption that both endpoints agree to use.

Diffie-Hellman Algorithm

The Diffie-Hellman algorithm was the first public-key algorithm and is still considered one of the best. IKE uses public-key cryptography to negotiate security parameters and protect key exchanges. Specifically, the Diffie-Hellman algorithm is used in the IKE negotiations to enable the two peers to agree on a shared secret by generating the key for use. This is why you will see that the Diffie-Hellman algorithm is used several times throughout the process.

In general, here is how the algorithm works: Each peer contains a private key. The Diffie-Hellman algorithm takes that private key and generates a public key. The public key is a product of the private key, but is such that the private key cannot be deduced by knowing the public key. The peers then exchange public keys, as shown in Figure 9-10.

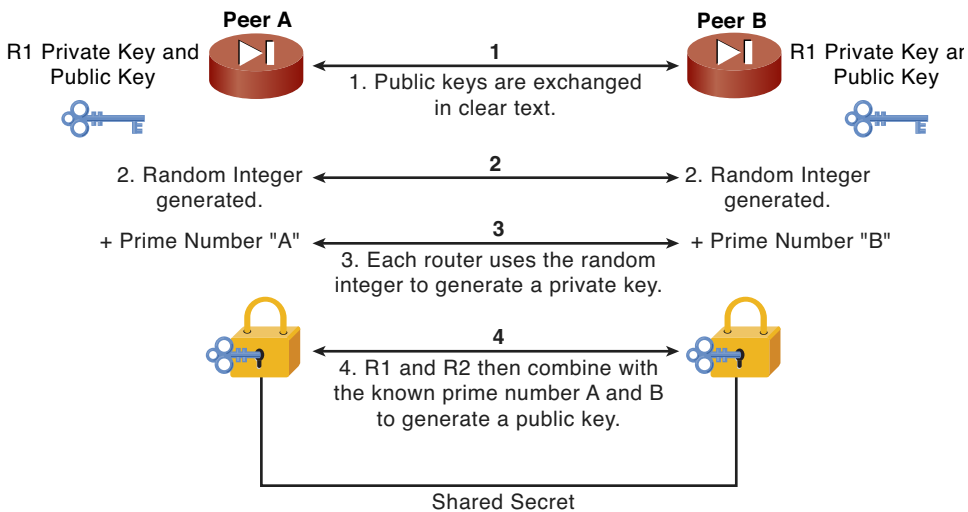


Figure 9-10 *Diffie-Hellman Key Exchange*

Note Symmetric key algorithms use the same key for both encryption and decryption. Symmetric key algorithms offer significant advantages over public key algorithms. The main advantage is speed because only one key is randomly generated, as opposed to two in public key cryptography. The only problem with asymmetric key algorithms is the security involved in sharing the private key between peers over an unprotected link.

If Peer A wants to pass encrypted traffic to Peer B, Peer A encrypts the traffic going to Peer B with Peer B's public key.

Peer B then uses its own private key to decrypt the message because its public key is derived from its private key. This ensures that only Peer B can decrypt the message because only Peer B knows its own private key.

This method enables a secure communications channel to be established (ISAKMP SA) so that subsequent IPsec SAs can securely exchange key information in privacy without having to use a public key algorithm to exchange their own keys every time encrypted traffic is passed. Figure 9-11 shows the various steps in ISAKMP Phase 1 and Phase 2 negotiations.

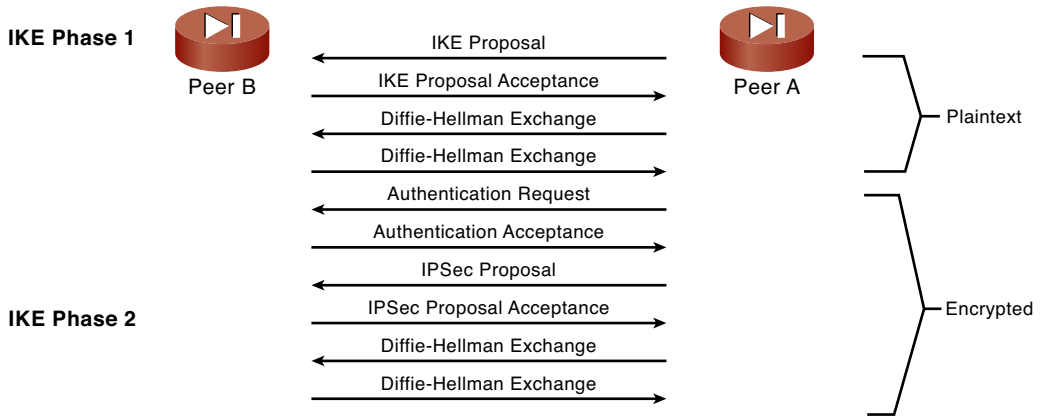


Figure 9-11 VPN Connection Establishment

Figure 9-11 illustrates that traffic is already encrypted before the end of IKE Phase 1. This provides for a secure exchange of the IPsec proposals and keys performed on behalf of IPsec in IKE Phase 2.

In addition to providing a secure mechanism for key exchange and managing IPsec SAs, ISAKMP also provides several other important functions. ISAKMP can be configured to set IPsec SA lifetimes, which enables more control over how often keys are exchanged. It also enables keys to change during communication without removing and re-creating the IPsec SAs. With standalone IPsec, if keys are to change during communication, existing SAs are “torn down” and rebuilt with the new keys. Because ISAKMP negotiates SAs for IPsec and protects them with its own SA, keys can be changed on-the-fly without re-creating SA negotiations. This provides a substantial advantage over IPsec alone. ISAKMP also enables dynamic authentication of peers and data integrity checks via the use of one-way hash algorithms.

Router Configuration as VPN Peer

We wanted to include one of the ways to configure a router with the capability to be part of a site-to-site VPN. This particular configuration is important because Cisco routers make up 80 percent of the routers in operation today. Therefore, it seemed that many networks could have their security greatly increased by using the router to terminate VPNs.

Configuring ISAKMP

IKE exists only to establish SAs for IPsec, but before it can do this, it must negotiate an SA (an ISAKMP SA) relationship with the peer. Because IKE negotiates its own policy, you can configure multiple policy statements with different configuration statements, and then let the two hosts come to an agreement. ISAKMP negotiates the following:

- **Encryption algorithm:** Used to protect user data transmitted between two IPsec peers (DES or 3DES).
- **Hashing algorithm: MD5 or SHA:** This selection specifies the hash algorithm used to ensure data integrity. The default is SHA-1. MD5 has a smaller digest and is considered slightly faster than SHA-1.
- **Authentication:** RSA signatures, RSA encrypted nonces (random numbers), or pre-shared keys. This selection specifies the method of authentication that establishes the identity of each IPsec peer. Preshared keys do not scale well with a growing network, but they are easier to set up in a small network.
- **Lifetime of the SA (in seconds):** The default is 86,400 seconds or 24 hours. As a general rule, a shorter lifetime (up to a point) provides more secure IKE negotiations. However, with longer lifetimes, future IPsec security associations can quickly be set up. However, as with many of the characteristics used in the VPN creation, the VPN tunnel activates and functions if the values do not match.

There is an implicit trade-off between security and performance when you choose a specific value for each parameter. The level of security provided by the default values is adequate for most organizations' security requirements. If you interoperate with a peer that supports only one of the values for a parameter, your choice is limited to the other peer's supported value.

When the IKE negotiation begins, the peer that initiates the negotiation sends all its policies to the remote peer, which tries to find a match. The remote peer checks its policies in order of priority (highest priority first) until a match is found.

A match is made when both policies from the two peers contain the same encryption, hash, authentication, and Diffie-Hellman parameter values, and when the remote peer's policy specifies a lifetime less than or equal to the lifetime in the policy being compared. If the lifetimes are not identical, the shorter lifetime (from the remote peer's policy) is used.

If no acceptable match is found, IKE refuses negotiation and IPsec is not established. If a match is found, IKE completes negotiation and IPsec security associations are created. Currently, there are two methods of configuring ISAKMP:

- Use preshared keys, which have the advantage of being simple to configure.
- Use a centralized *Certificate Authority (CA)*, which is a third-party entity responsible for issuing and revoking certificates. Each device that has its own certificate and public key of the CA can authenticate every other device within a given CA's domain. This solution has the advantage of being scalable throughout a large enterprise network.

Note IKE negotiation is done on UDP port 500. IPsec uses IP protocols 50 and 51. Make sure these are permitted on any access lists you have between the peers.

The following section discusses the use of preshared keys, which is by far the most common method of configuring ISAKMP.

Preshared Keys

If you use the IKE authentication method of preshared keys, you are setting the keys, or in other words, sharing them with the other peer with whom you plan to create a VPN by manually configuring these keys on the device and its peers. You can specify the same key to share with multiple peers, but it is more secure to specify different keys to share between different pairs of peers. To configure a preshared key on the ASA firewall, perform the following steps. Although configuring IKE is simple and you do not use a CA, it does not scale well. To configure IKE, you must do the following:

Step 1. Configure ISAKMP policy options.

Step 2. Configure ISAKMP key.

Configuring the ISAKMP Protection Suite

The following command creates the ISAKMP policy object. You can have multiple policies, but only one is in this example:

```
CYBERWRAITH(config)# crypto isakmp policy 1
CYBERWRAITH(config-isakmp)#
```

With the following **group** command, you can declare what size modulus to use for Diffie-Hellman calculation:

```
CYBERWRAITH(config-isakmp)# group 2
```

Group 1 is 768 bits, and group 2 is 1024 bits. Why would you use one over the other? First, not all vendors support group 2. Second, group 2 is also significantly more CPU-intensive than group 1; therefore, you would not want to use group 2 on low-end routers such as the Cisco 2500 series or less. On the other hand, group 2 is more secure than group 1.

Because security is of primary concern, group 2 is used here. (Make sure the peer is also configured to use group 2.) The default is group 1. If you select the default properties, the group 1 lines do not show up when you show the **configuration** command.

MD5 is the hashing algorithm as configured in the following command. Although implementing SHA and MD5 are both mandatory, not all peers can be configured to negotiate one or the other:

```
CYBERWRAITH(config-isakmp)# hash md5
```

The following command shows the security association's lifetime—in this case, 500 seconds. If you do not set a lifetime, it defaults to 86,400 seconds or 1 day. When the lifetime timer fires, the SA is renegotiated as a security measure.

```
CYBERWRAITH(config-isakmp)# lifetime 500
```

The **authentication pre-share** command tells IKE what key to use:

```
CYBERWRAITH(config-isakmp)# authentication pre-share
```

Two options for the **authentication** command besides the **pre-share** are

- **rsa-encr**: Configures RSA-encrypted nonces
- **rsa-sig**: Configures RSA signature

The **rsa-encr** and the **rsa-sig** options are addressed in the section “Using a CA.” For now, remember that **rsa-sig** is the default.

Configuring the ISAKMP Key

The following commands tell IKE what key to use. Remember that the peer, 192.168.10.38 in this case, must have the same key **Slurpee-Machine** in its configuration. I use this particular key because I am configuring a VPN to my good friend Cary's office and he is addicted to these cold delights.

```
CYBERWRAITH(config-isakmp)# exit
```

```
CYBERWRAITH(config)# crypto isakmp key Slurpee-Machine address 192.168.10.38
```

At this point, you are finished with IKE configuration. For the record, the following lines are the peer's IKE configuration:

```
crypto isakmp policy 1
  hash md5
  group 2
  authentication pre-share
crypto isakmp key Slurpee-Machine address 192.168.10.66
```

Configuring IPsec

Whether you use preshared keys or configure a CA, you still have to set up IPsec after you set up IKE. Regardless of which IKE method you use, the IPsec configuration steps are the same. To configure IPsec, you need to

- Step 1.** Create the extended ACL.
- Step 2.** Create the IPsec transforms.
- Step 3.** Create the crypto map.
- Step 4.** Apply the crypto map to an interface.

Step 1: Create the Extended ACL

The following command is a simple ACL that enables the routers to talk to one another (a Telnet from one router to the next, for example):

```
CYBERWRAITH(config)# access-list 101 permit ip host 192.168.10.38 host
192.168.10.66
```

A more realistic ACL looks like the following command:

```
CYBERWRAITH(config)# access-list 101 permit ip 192.168.3.0 0.0.0.255 10.3.2.0
0.0.0.255
```

This command is an ordinary extended ACL, where 192.168.3.0 is a subnet behind the router in question and 10.3.2.0 is a subnet somewhere behind the peer router. Remember that permit means encrypt, and deny means do not encrypt.

Step 2: Create the IPsec Transforms

A transform describes a security protocol (AH or ESP) with its corresponding algorithms; for example, ESP with the DES cipher algorithm and HMAC-SHA for authentication. A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow. You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry's access list. During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers' IPsec security associations. With manually established security associations, there is no negotiation with the peer, so both sides must specify the same transform set.

Create three transform sets, as done in the following command lines:

```
CYBERWRAITH(config)# crypto ipsec transform-set PapaBear esp-rfc1829
CYBERWRAITH(cfg-crypto-trans)# exit
CYBERWRAITH(config)# crypto ipsec transform-set MamaBear ah-md5-hmac esp-des
```

```

CYBERWRAITH(cfg-crypto-trans)# exit
NRGI(config)# crypto ipsec transform-set BabyBear ah-rfc1828
CYBERWRAITH(cfg-crypto-trans)# exit
CYBERWRAITH(config)#

```

The first set uses only ESP; the second set uses AH combined with ESP; and the last set uses only AH. During IPsec SA negotiation, all three are offered to the peer, which chooses one. Also, use the default tunnel mode for all three transform sets. Transport mode can be used only when the crypto endpoints are also the communication's endpoints. The **mode transport** command under the transform-set configuration can specify the transport mode. Tunnel mode is used primarily for the VPN scenario.

Also note that esp-rfc1829 and ah-rfc1828 are based on the original RFCs for this technology and are obsolete transforms included for backward compatibility. Not all vendors support these transforms, but other vendors support only these transforms. Finally, the transform sets in the commands are not necessarily the most practical. For example, both PapaBear and BabyBear have substandard transform-sets. You should use esp-rfc1829 and ah-rfc1828 together in the same transform-set.

Step 3: Create the Crypto Map

Crypto maps specify IPsec policy. Crypto map entries created for IPsec pull together the various security settings that set up IPsec security associations, including the following:

- Which traffic should be protected by IPsec (per a crypto access list)
- Where IPsec-protected traffic should be sent (who the peer is)
- The local address to be used for the IPsec traffic
- What IPsec security should be applied to this traffic (selecting from a list of one or more transform sets)
- Whether security associations are manually established or established via IKE
- Other parameters that might be necessary to define an IPsec SA

For IPsec to succeed between two peers, both peers' crypto map entries must contain compatible configuration statements. When two peers try to establish a security association, they should each have at least one crypto map entry compatible with one of the other peer's crypto map entries.

Using the **ipsec-isakmp** tag tells the router that this crypto map is an IPsec crypto map. Although only one peer is declared in this crypto map, a given crypto map can have multiple peers. The session key lifetime can be expressed in either kilobytes (after x amount of traffic, change the key) or seconds, as shown in the following commands. The goal is to make a potential attacker's efforts more difficult:

```

CYBERWRAITH(config)# crypto map armadillo 10 ipsec-isakmp
CYBERWRAITH(config-crypto-map)# set peer 192.168.10.38
CYBERWRAITH(config-crypto-map)# set session-key lifetime seconds 4000

```

```
CYBERWRAITH(config-crypto-map)# set transform-set MamaBear PapaBear BabyBear
CYBERWRAITH(config-crypto-map)# match address 101
```

The **set transform-set** command is where you associate the transforms with the crypto map. In addition, the order in which you declare the transforms is significant. You most prefer MamaBear in this configuration, and then the rest in descending order of preference to BabyBear.

The crypto map access list bound to the outgoing interface selects the IPsec packets destined to an IPsec tunnel. IPsec packets that arrive from an IPsec tunnel are authenticated or deciphered by IPsec and are subject to the proxy identity match of the tunnel.

Note What happens if a packet does not meet the requirements for encryption? Simply put, that packet is then discarded into the bit bucket.

The **match address 101** command simply means to use access list 101 to determine what traffic is interesting so that it will be placed into the VPN tunnel. You can have multiple crypto maps with the same name (armadillo, in the following example) and different sequence numbers (10, in the following example). The combination of multiple crypto maps and different sequence numbers enables you to mix and match classic crypto and IPsec. You can also modify your PFS configuration here. PFS group1 is the default in the example given here. You could change the PFS to group2 or turn it off altogether, which you should not do.

Step 4: Apply the Crypto Map to an Interface

The following commands apply the crypto map to the interface. Remember to apply the crypto map to the egress interface, not the ingress one. If you have multiple crypto maps that you want to apply to this interface, you must tack the name onto the list in the **crypto map** command:

```
CYBERWRAITH(config)# int e0
CYBERWRAITH(config-if)# crypto map armadillo
```

Remember that crypto maps and their access lists are direction-based (either inbound or outbound) and that traffic not matching the access list is still transmitted without being encrypted.

Firewall VPN Configuration for Client Access

You can configure Cisco ASA Firewalls to terminate client VPNs, thus allowing users to securely access corporate resources.

Used with IKE, dynamic crypto maps can ease IPsec configuration and are recommended for use in networks where the peers are not always predetermined. You use dynamic cryp-

to maps for VPN clients (such as mobile users) and routers that obtain dynamically assigned IP addresses.

Dynamic crypto maps can be used only to negotiate SAs with remote peers that initiate the connection. They cannot be used to initiate connections to a remote peer. With a dynamic crypto map entry, if outbound traffic matches a **permit** statement in an access list and the corresponding security association is not yet established, the ASA firewall drops the traffic.

A dynamic crypto map entry is essentially a crypto map entry that does not have all the parameters configured. The dynamic crypto map acts as a policy template where the missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a peer's requirements. This allows peers to exchange IPsec traffic with the ASA firewall, even if the ASA firewall does not have a crypto map entry specifically configured to meet all the peer's requirements. Dynamic crypto maps are found for use by VPN clients on PCs.

If the ASA firewall accepts the peer's request at the point that it installs the new IPsec security associations, it also installs a temporary crypto map entry. This entry is filled in with the results of the negotiation. At this point, the ASA firewall performs normal processing, using this temporary crypto map entry as a normal entry, and even requests new security associations if the current ones are expiring (based on the policy specified in the temporary crypto map entry). When the flow expires (that is, all the corresponding security associations expire), the temporary crypto map entry is removed.

Like regular static crypto map entries, dynamic crypto map entries are grouped into sets. A set is a group of dynamic crypto map entries all with the same *dynamic-map-name*, but each with a different *dynamic-seq-num*. If this is configured, the data flow identity proposed by the IPsec peer should fall within a permit statement for this crypto access list. If this is not configured, the ASA firewall accepts any data flow identity proposed by the peer.

You can add one or more dynamic crypto map sets into a crypto map set via crypto map entries that reference the dynamic crypto map sets. You should set the crypto map entries that reference dynamic maps to be the lowest priority entries in a crypto map set. (That is, use the highest sequence numbers.)

Note Use care when using the *any* keyword in permit entries in dynamic crypto maps. If it is possible for the traffic covered by such a permit entry to include multicast or broadcast traffic, the access list should include deny entries for the appropriate address range. Access lists should also include deny entries for network and subnet broadcast traffic, and for any other traffic that should not be IPsec protected.

The procedure for using a crypto dynamic map entry is the same as the basic configuration described in the "Basic IPsec Configuration" section, except instead of creating a static crypto map entry, you create a crypto dynamic map entry. You can also combine static and dynamic map entries within a single crypto map set.

Step 1: Define Interesting Traffic

The VPN device recognizes interesting traffic as defined in the ACL that needs to be sent via the VPN tunnel:

```
access-list VPN_NAME remark ACL DEFINES VPN ACCESS
access-list VPN_NAME extended permit ip source destination
```

Step 2: IKE Phase 1 [udp port 500]

VPN device negotiates an IKE security policy and establishes a secure tunnel using five parameters as defined in the ISAKMP policy statements. These statements are also processed from the lowest ID number to the highest; processing stops when both VPN endpoints agree on the five parameters. The following HAGLE mnemonic might help you remember the five parameters needed here:

Hash: md5 or sha-1 used for data integrity, ensures not altered

Authentication: pre-share or rsa-sig Provides origin authentication

Group (DH): 1 [768 bit] or 2 [1024 bit]

Lifetime: 86,400 secs

Encryption: des [default], 3des, AES

With these in mind, consider the following:

```
!
crypto isakmp policy 10
  authentication pre-share
  encryption aes-256
  hash sha
  group 2
  lifetime 86400
!
crypto isakmp policy 20
  authentication pre-share
  encryption aes-256
  hash md5
  group 2
  lifetime 86400
!
```

Step 3: IKE Phase 2

IKE negotiates IPsec SA parameters to encrypt traffic:

```
tunnel-group PEER_IP type ipsec-L2L
tunnel-group PEER_IP ipsec-attributes
pre-shared-key password
!
```

```
crypto map CRYPTO-MAP_NAME ID# set transform-set XFORMSET-AES-MD5
crypto map CRYPTO-MAP_NAME ID# set security-association lifetime seconds 86400
crypto map CRYPTO-MAP_NAME ID# match address ACL-NAME
crypto map CRYPTO-MAP_NAME ID# set peer PEER_IP
```

Step 4: Data Transfer

IPsec uses two protocols to function securely:

- Encapsulating Security Payload (ESP)^[ip protocol 50] to provide data encryption (confidentiality)
- Authentication Header (AH)^[ip protocol 51] to provide origin authentication

The protocol field in the packet's IP header will be 50 (ESP) or 51 (AH) to indicate the next protocol to be found in the packet.

Step 5: Tunnel Termination

IPsec tunnels are typically terminated when the SA times out after a specified number of seconds has elapsed (or by bytes transmitted). When the SA terminates, the ASA discards the keys and creates new ones if needed. New SAs are usually established before the terminating SAs expire so that a given flow continues uninterrupted.

SSL VPN Overview

Today's remote-access VPN deployments require the capability to safely and easily extend corporate network access beyond managed desktops to different users' devices, while protecting these endpoints and key corporate resources from ever-evolving threats. SSL VPN solutions can be customized for companies of any size and deliver remote access connectivity features and benefits such as the following:

- Lower desktop support costs through web-based access without preinstalled desktop software, which facilitates customized remote access.
- Threat protection provided by integrated security in the platform protects against viruses, worms, spyware, and hackers.
- Flexible and cost-effective licensing.
- Reduced cost and management complexity—both an SSL VPN and IPsec VPN on one device means you do not need other security devices.

SSL-based VPNs provide remote-access connectivity from almost any Internet-enabled location using a web browser and its native SSL encryption. It does not require any special-purpose client software to be pre-installed on the system; this makes SSL VPNs capable of “anywhere” connectivity from company-managed desktops and noncompany-managed desktops, such as employee-owned PCs, contractor or business partner desktops, and Internet kiosks. Any software required for application access across the SSL VPN

connection is dynamically downloaded on an as-needed basis, thereby minimizing desktop software maintenance.

SSL VPNs provide two different types of access: clientless and full network access. Clientless access requires no specialized VPN software on the user desktop. All VPN traffic is transmitted and delivered through a standard web browser; no other software is required or downloaded. Because all applications and network resources are accessed through a web browser, only web-enabled and some client/server applications, such as intranets, applications with web interfaces, email, calendaring, and file servers, can be accessed using a clientless connection. This limited access, however, is often a perfect fit for business partners or contractors who should have access to only a limited set of resources on the organization's network. Furthermore, delivering all connectivity through a web browser eliminates provisioning and support issues because no special-purpose VPN software must be delivered to the user desktop.

SSL VPN full network access enables access to virtually any application, server, or resource available on the network. Full network access is delivered through a lightweight VPN client dynamically downloaded to the user desktop (through a web browser connection) upon connection to the SSL VPN gateway. This VPN client, because it is dynamically downloaded and updated without any manual software distribution or interaction from the end user, requires little or no desktop support by IT organizations, thereby minimizing deployment and operations costs. Like clientless access, full network access offers full access control customization based on the access privileges of the end user. Full network access is a natural choice for employees who need remote access to the same applications and network resources they use in the office or for any client/server application that cannot be delivered across a Web-based clientless connection. Figure 9-12 illustrates the flexibility available using SSL VPNs and the many ways they can be used more securely and flexibly than IPsec client-based VPNs.

Comparing SSL and IPsec VPNs

IPsec-based VPNs are the deployment-proven, remote-access technology used by most organizations today. IPsec VPN connections are established using pre-installed VPN client software on the user desktop, thus focusing it primarily on company-managed desktops. IPsec-based remote access also offers tremendous versatility and customizability through modification of the VPN client software. Using APIs in IPsec client software, organizations can control the appearance and function of the VPN client for use in applications such as unattended kiosks, integration with other desktop applications, and other special use cases.

Both IPsec and SSL VPN technologies offer access to virtually any network application or resource. SSL VPNs offer additional features, such as easy connectivity from noncompany-managed desktops, little or no desktop software maintenance, and user-customized web portals upon login. Table 9-1 compares the two technologies.

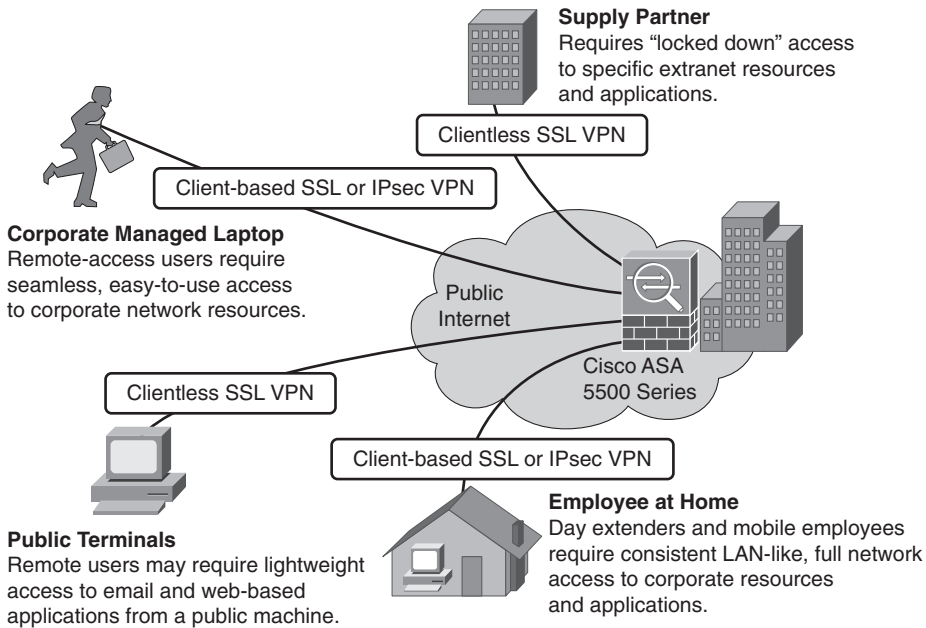


Figure 9-12 SSL VPN Deployment Options

Table 9-1 Comparing IPsec and SSL VPN Technologies

Access Needs	SSL VPNs	IPsec VPNs
Application and Network Resource Access	Both SSL (using full network access) and IPsec VPNs offer broad access to virtually any application or network resource.	
End-User Access Method	SSL VPNs are initiated using a web browser.	IPsec VPNs are initiated using pre-installed VPN client software.
End-User Access Device Options	SSL VPN enables access from company-managed, employee-owned, contractor and business partner desktops, and Internet kiosks.	IPsec VPNs enable access primarily from company-managed desktops.
Desktop Software Requirements	Only a web browser is required for an SSL VPN.	An IPsec VPN requires proprietary pre-installed client software.

Table 9-1 *Comparing IPsec and SSL VPN Technologies*

Access Needs	SSL VPNs	IPsec VPNs
Desktop Software Updates	Basic SSL VPN access can operate without any special-purpose desktop software, so no updates are required. Full network application access is provided using software that automatically installs and updates without any user knowledge or intervention.	IPsec VPNs can automatically update but are more intrusive and require user input.
Customized User Access	SSL VPNs offer granular access policies to define what network resources a user has access to, as well as user-customized web portals.	IPsec offers granular access policies but no web portals.

Which to Deploy: Choosing Between IPsec and SSL VPNs

IPsec is a widely deployed technology that is well understood by end users and has established IT deployment support processes. Many organizations find that IPsec meets the requirements of users already using the technology. But the advantages of dynamic, self-updating desktop software, ease of access for noncompany-managed desktops, and highly customizable user access make SSL VPNs a compelling choice for reducing remote-access VPN operations costs and extending network access to hard-to-serve users such as contractors and business partners. As such, organizations often deploy a combination of SSL and IPsec approaches. IPsec is commonly left in place for the existing installed base. SSL is deployed for new users, users with “anywhere” access requirements, contractors, and extranet business partners. By offering both technologies on a single platform, Cisco remote-access VPN solutions make the choice simple: Deploy the technology that is optimized for your deployment and operating environment. Table 9-2 summarizes the issues to consider when evaluating which VPN technology best fits your companies operating environment.

Table 9-2 *Choosing a Remote-Access VPN Technology*

Feature	SSL VPN	IPsec VPN
“Anywhere” access from non-company-managed devices, such as employee-owned desktops and Internet kiosks	✓	—
Business partner access	✓	—

Table 9-2 *Choosing a Remote-Access VPN Technology*

Feature	SSL VPN	IPsec VPN
User-customized access portals	✓	—
Minimized desktop support and software distribution	✓	—
Greatest flexibility to the end users	✓	✓
Greatest VPN client customizability	—	✓
Reduced administrative burden; that is, no client support	✓	—
Ability to maintain existing IT deployment and support processes	—	✓

Remote-Access VPN Security Considerations

Worms, viruses, spyware, hacking, data theft, and application abuse are considered among the greatest security challenges in today's networks. Remote-access and remote-office VPN connectivity are common points of entry for such threats because of how VPNs are designed and deployed. For both new and existing IPsec and SSL VPN installations, VPNs are often deployed without proper endpoint and network security. Unprotected or incomplete VPN security can lead to the following network threats:

- Enables remote-user VPN sessions to bring malware into the main office network, causing virus outbreaks that infect other users and network servers.
- Enables users to generate unwanted application traffic, such as peer-to-peer file sharing, into the main office network causing slow network traffic conditions and unnecessary consumption of expensive WAN bandwidth.
- Enables theft of sensitive information, such as downloaded customer data, from a VPN user desktop.
- Enables hackers to hijack remote-access VPN sessions, providing the hacker access to the network as if they were a legitimate user.

To combat these threats, the user desktop and the VPN gateway to which the user connects must be properly secured as part of the VPN deployment. User desktops should have endpoint security measures such as data security for data and files generated or downloaded during the VPN session, antispymware, antivirus, and personal firewall. The VPN gateway should offer integrated firewall, antivirus, antispymware, and intrusion prevention. Alternatively, if the VPN gateway does not provide these security functions, separate security equipment can be deployed adjacent to the VPN gateway to provide appropriate protection. It's your responsibility as a security professional to educate your end users through good, consistent user awareness training.

Cisco remote-access VPN solutions offer threat-protected VPN services with full firewall, antivirus, antispyware, intrusion prevention, application control, and full endpoint security capabilities. These security services are integrated into the VPN platform, delivering a threat-protected VPN solution without any additional equipment, design, deployment, or operational complexity.

Steps to Securing the Remote-Access VPN

Technologies required for mitigating malware such as worms, viruses, and spyware and for preventing application abuse, data theft, and hacking exist in the security infrastructure of many organizations' networks. In most cases, however, they are not deployed in such a way that they can protect the remote-access VPN because of the native encryption of VPN traffic. Although additional security equipment may be purchased and installed to protect the VPN, the most cost-effective and operationally efficient method of securing remote-access VPN traffic is to look for VPN gateways that offer native malware mitigation and application firewall services as an integrated part of the product, as shown in Figure 9-13.

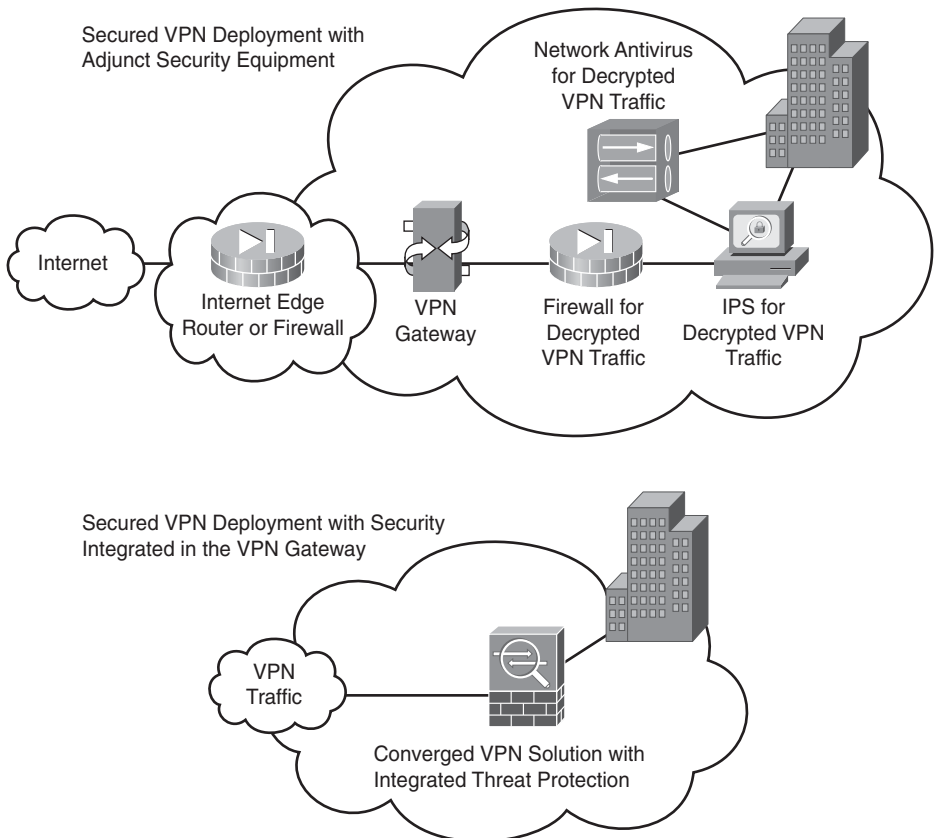


Figure 9-13 *Securely Deploying VPNs*

Cisco AnyConnect VPN Secure Mobility Solution

End users want the flexibility to choose how, when, and where to access both personal and professional information to be productive without being inconvenienced by security checks. The IT support staff, on the other hand, wants to allow access for end users while ensuring that the corporate network and the access remains secure. Employees using their own devices to access the corporate network introduce an additional burden to an organization's IT department. When business data makes its way onto an employee-owned device, it can be a challenge for the enterprise to control its spread or use. To support the increasing number of mobile workers, corporate security administrators must provide context-aware security and policy enforcement, regardless of the end user's location, what device they use, and where the information they access is located. Administrators must also be able to support a heterogeneous set of laptops and mobile devices to encourage choice for their clients—the end users. And finally, they must provide this security unobtrusively to minimize end-user concerns. The Cisco AnyConnect Secure Mobility Solution provides the following:

- Security policy enforcement that is context-aware, comprehensive, and preemptive
- A connectivity experience that is intelligent, seamless, and always on
- Secure mobility across today's proliferating managed and unmanaged mobile devices

Cisco AnyConnect with Cisco ASA 5500 Series Adaptive Security Appliances at the head end provides the remote-access connectivity portion of Cisco AnyConnect Secure Mobility. This AnyConnect VPN is becoming more powerful and flexible with the capability now for it to support any mobile operating system such as Apple's IOS, Android, and the Cisco CIUS platform.

Both the user and device must be authenticated and validated prior to being provided access to the network. After the user is authenticated, the Cisco AnyConnect Secure Mobility Solution can decide which applications and resources the user should have access to. Ideally, this authentication would be transparent to the user. For devices to be authenticated, they must comply with corporate policies and have up-to-date security in place.

The Cisco AnyConnect Secure Mobility Solution enables the connection to simply work and be reliably connected without the user needing to juggle where and how to best connect and persist, even when roaming between networks. As mobile workers roam to different locations, the always-on intelligent VPN in the AnyConnect Secure Mobility client automatically selects the most optimal network access point and adapts its tunneling protocol to the most efficient method.

Mobility changes everything. It changes how and where people work—and it creates new IT security challenges. The Cisco AnyConnect Secure Mobility Solution provides comprehensive and secure remote access. Mobile users can enjoy persistent connectivity back to their corporations, and IT administrators can enable smart, context-aware security policies to protect corporate assets. Figure 9-14 illustrates how AnyConnect VPNs can be successfully deployed within your organization.

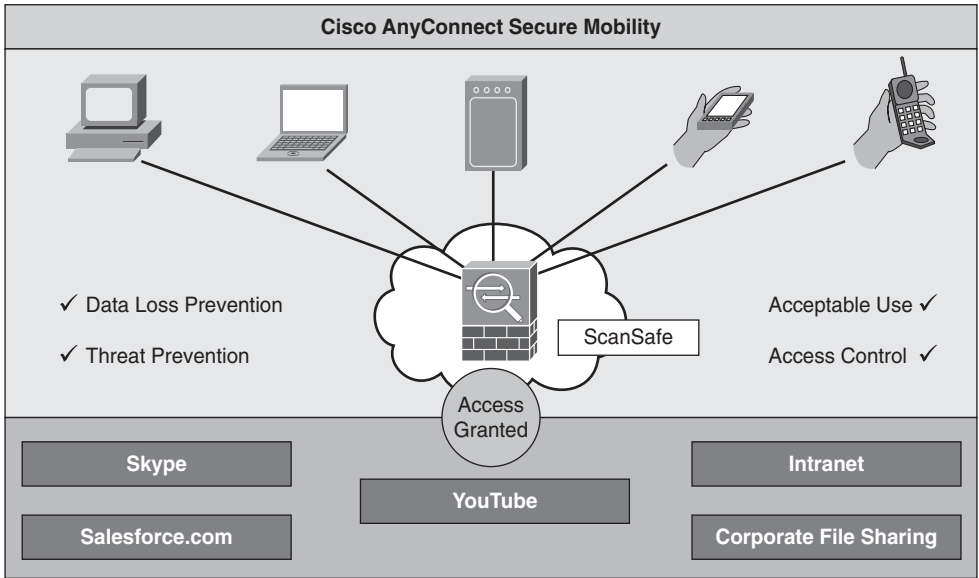


Figure 9-14 AnyConnect Deployment Considerations

Chapter Summary

This chapter discussed what a VPN is and the many benefits that it brings to networks everywhere. The most popular benefit of implementing VPNs is the cost reduction and overall financial savings. The reduction of bandwidth costs has made VPNs one of the best solutions available.

This chapter focused on the best available VPNs: IPsec and SSL/AnyConnect. To understand how they protect your data, the chapter examined all those different levels, phases, and types of processes involved in getting your data packets encrypted into your IPsec-based VPNs. This was a truly amazing task because the subject matter gets complicated quickly. If you are a bit confused, that is understandable; this is a complicated chapter and worth a second read through.

Chapter Review Questions

1. Is it possible to have unencrypted VPNs?
2. What are the three types of VPNs?
3. Select three VPN features and benefits and explain how your organization can directly benefit from each.
4. VPN concentrators are designed for many users—explain how many and when they should be used.
5. Does the VPN Client Software for PCs support Apple's powerful new operating system, Mac OS X?
6. When does split tunneling occur?
7. In relation to a data stream, what role does authentication play in securing it?
8. When tunneling data in IPsec, what three protocols play a role in the process?
9. In site-to-site VPNs, what are the two different encapsulating protocols and what are the differences between them?
10. Name three benefits of IKE.
11. What are three important differences between SSL and AnyConnect VPNs?

This page intentionally left blank

Wireless Security

“Why is the man who invests all your money called a broker?”

—Comedian George Carlin

By the end of this chapter, you should know and be able to explain the following:

- The essentials of wireless LANs, including their benefits and risks
- The major threats to a wireless network
- How to secure a wireless network
- The breadth and scope of possible attacks and exploits available to attackers

Answering these key questions will enable you to understand the overall characteristics and importance of network security within the wireless networking space. By the time you finish this book, you will have a solid appreciation for network security, its issues, how it works, and why it is important.

When was the last time you went on vacation to get away from it all? Perhaps to some remote beach or maybe a getaway to the country? Imagine that you walk out the patio door of your hotel room (an ocean view, of course) and admire the beauty of the sun setting on the ocean. The air is cool, so you decide to sit on the porch in your favorite lounge chair; the seagulls are playing, the waves are breaking in a rhythmic beat, and *beep-beep-beep*—your iPhone begins to go off!

Who could possibly be paging you while you are trying to relax and unplug? What emergency could be so grave that it would require you to be interrupted on this fantasy vacation?

According to the message on the display, there seems to be a problem with the company’s mission-critical firewall/VPN/Exchange server/<insert emergency here>. It looks serious, so you conclude that you need to log in to your office network and take a look.

It is a good thing you chose a hotel with “free” high-speed wireless Internet access. You cannot avoid turning on the laptop that you were not planning to turn on while you were on vacation; you are needed for an emergency.

So, here you are on the patio of your suite (why not a suite? it’s my story!) booting up your laptop and explaining to your wife that it won’t take long. You see the “blinky-blinky” of the wireless NIC’s status lights, you just need to log in. All systems are go!

You fire up Telnet and proceed to log in to the router/firewall and start snooping around to see what the problem could be. This should not take too long, you say to yourself and to your wife. There is still plenty of time to enjoy the rest of the evening and perhaps have a nice dinner. An hour goes by and you have solved the problem. You are quite taken with yourself for being ingenious enough to diagnose and resolve the situation within a few tick-tocks.

*Screeeech...*stop the movie for a second. Unknowingly, the “vacationing uber tech” just caused his company to lose millions of dollars. How, you might ask, did this dashing guy in the movie cause millions of dollars to be lost just by logging in to his company’s router/firewall to fix a problem?

It was not the act of connecting to the router/firewall that caused the problem; it was the fact that he used a wireless connection. You see, the company that uber tech worked for (yes, past tense because he no longer works for them as a result) is a multinational corporation that was about to announce the creation of a new widget that was capable of converting discarded pizza boxes into something truly spectacular we are legally unable to disclose; a competitor of this revolutionary company not only wanted to stop this announcement—but they also wanted a copy of the plans for this widget so they could bring it to market first.

It seems that a hacker employed by the competitor was paid to follow vacationing uber tech and, at a convenient moment, download the contents of his laptop, in hopes that the hacker could find some proprietary information about the widget. Upon “seeing” uber tech boot up his laptop, complete with wireless NIC, the hacker realized that he had struck gold and decided to do some long-distance sniffing and hacking, courtesy of uber tech’s unsecured wireless connection. Long-distance sniffing and hacking—sounds like a script from “Mission: Impossible,” doesn’t it? Too far fetched to actually happen? The truth is that this type of scenario occurs on a daily basis. Bad guys with wireless-enabled laptops steal information right out of the air with little effort. They use tools that are readily available on the Internet and can cause many problems for companies that do not take the time to understand the threats an unsecured wireless connection poses to their corporate network.

This chapter covers several topics related to wireless networking security and helps you identify, understand, and prevent the types of intrusions to which wireless connections are vulnerable from the outside. This chapter focuses on available commercial wireless products and not the home user versions from Cisco subsidiaries such as Linksys. However, do not ignore the advice and suggestions given here when setting up your wireless at home.

Essentials First: Wireless LANs

This chapter discusses the use of wireless LANs (WLAN), which are roaring into use almost every time you turn around—from airports, restaurants, and coffee shops, to people's homes. The growth of personal computers in the 1980s led to the creation of LANs and the Internet in the 1990s; this allowed for connections, regardless of geographic location. WLANs are proving to be the next technology growth area for the 2000s. Businesses are, of course, recognizing the benefits of WLANs and deploying them in ever-increasing numbers. Just as businesses were forced to provide security to PCs and the Internet, so too must businesses understand that, despite the productivity and mobility gains they provide, WLANs have associated security risks that must be addressed.

A WLAN offers a quick and effective extension of a wired LAN. By simply installing access points to the wired network, personal computers and laptops equipped with wireless LAN cards can connect with the wired network at broadband speeds (or greater) from up to 300 yards away from the wireless access point. This means that computers are no longer tied to the infrastructure of wires—rather liberating, isn't it?

The majority of WLAN deployments have used a wireless transmission standard known as 802.11b. The IEEE 802.11b standard operates at the radio frequency of 2.4 GHz—a frequency unregulated by governments. The 802.11b standard offers connectivity speeds of up to 11 Mbps, which provides enough speed to handle large email attachments and run bandwidth-intensive applications such as videoconferencing. The *802.11g* standard with speeds of up to 54 Mbps now dominates the wireless LAN market; other variations of the 802.11 standard are constantly being developed to handle an ever-increasing need for speed. *802.11n* is the latest standard variation, which offers wireless speeds of more than 100 Mbps.

The various wireless standards are targeted to different industry segments, as outlined in Tables 10-1 and 10-2.

Table 10-1 *802.11a—54 Mbp WLAN Standard Characteristics*

Standard	IEEE 802.11a, WLAN
Frequency wavelength	5 GHz
Data bandwidth	54 Mbps, 48 Mbps, 36 Mbps, 24 Mbps, 12 Mbps, 6 Mbps
Security measures	WEP
Optimum operating range	150 feet indoors, 300 feet outdoors
Best suited for a specific purpose or device type	Roaming laptops in home or business; computers when wiring is inconvenient

802.11a never took off; however, the recently ratified 802.11g holds some interesting options to include increased speed and security, as Table 10-2 documents.

Table 10-2 802.11g—54 Mbps/Wi-Fi Standard Characteristics

Standard	IEEE 802.11g, Wi-Fi
Frequency wavelength	2.4 GHz
Data bandwidth	54 Mbps, 48 Mbps, 36 Mbps, 24 Mbps, 12 Mbps, 6 Mbps
Security measures	WEP, AES (in Broadcom 54 g) and possibly WPA/Wi-Fi protected access
Optimum operating range	125 feet indoors and 460 feet outdoors under normal conditions
Best suited for a specific purpose or device type	Roaming laptops in home or business; computers when wiring is inconvenient

When 802.11b clients are granted access to an 802.11g wireless access point, security inevitably must be set (lowered) to allow 802.11b clients access; because of WEP and its problems, the entire wireless network encryption level is reduced to a lowest common denominator. Table 10-3 looks at the specifications for the 802.11n standard.

Table 10-3 802.11n—100+ Mbps/Wi-Fi Standard Characteristics

Standard	IEEE 802.11n, Wi-Fi
Frequency wavelength	2.4/5 GHz
Data bandwidth	150 Mbps, 135 Mbps, 120 Mbps, 90 Mbps, 60 Mbps, 45 Mbps
Security measures	WEP, WPA, WPA2
Optimum operating range	230 feet indoors and 820 feet outdoors under normal conditions
Best suited for a specific purpose or device type	Mobile devices of any sort requiring performance equal to that of wired connections

What Is Wi-Fi?

The term *Wi-Fi* (*Wireless Fidelity*) is often used in discussions of 802.11 networks. Wi-Fi is most certainly the popular marketing word used today when talking about wireless (that is, Wi-Fi hot spots). The term Wi-Fi has become the common way to describe 802.11 wireless networks; it certainly is much quicker and easier to say, so marketing takes the credit for making it the mainstream label.

Wi-Fi also refers to certification by the Wi-Fi Alliance, an international nonprofit association of 802.11 product vendors. 802.11 products that receive Wi-Fi certification have been

tested and found to be interoperable with other certified products. This means you can use your Wi-Fi certified product with 802.11 Wi-Fi certified networks, whether they are Apple computers or Windows-based networks. Although 802.11 products that do not have Wi-Fi certification might work fine with certified devices, the Wi-Fi Certified logo is your assurance of interoperability. You can learn more about the Wi-Fi alliance online at www.wi-fi.org/.

Benefits of Wireless LANs

I had not flown much on airplanes recently, but an important family event—my honeymoon—allowed me the opportunity to fly. Not living near a major airport meant that I had to take a connecting flight to reach my destination. In my travels I experienced several different airports, each of which offered wireless connectivity to travelers, making layovers in airports a more productive time. Businesses of all types (coffee shops, hotels, malls, airports, and so on) all across the world are using this wireless access as a benefit to their customers, and wireless can easily be enabled for a relatively small financial investment. The benefits of deploying wireless LANs can be summarized as the following:

- **Attractive price:** Deploying a wireless LAN can be cheaper than a wired LAN because you do not have the need for wires; simply hook up an access point, and it can provide service to multiple computers.
- **Mobility:** Boost user productivity with the convenience of allowing users to wirelessly connect to the network from any point within range of an access point.
- **Rapid and flexible deployment:** Quickly extend a wired network with the ease of attaching an access point to a high-speed network connection.
- **Application agnostic:** As an extension of the wired network, WLANs work with all existing applications. As discussed previously, the standard protocol is TCP/IP, which is supported over all forms of wireless.
- **Performance:** WLANs offer a high-speed connection that, although equal to Ethernet, is quickly passing it in speed.

The benefits of WLANs are being recognized by individuals and businesses alike; the Gartner Group predicted that by 2010, the majority of Fortune 2000 companies would depend on wireless technology to meet their business and networking needs. I think Gartner got it right, but this event occurred well before 2010, demonstrating the fast-changing and fast-advancing world of wireless and mobility.

Wireless Equals Radio Frequency

The first technical concept you need to grasp when discussing what constitutes a threat to a wireless network is that 802.11 networks use radio frequencies to transmit the data back and forth between endpoints, just like the cordless phones or radios you have at home. The key difference is the frequency at which the signals are transmitted.

Radio waves can travel long distances, depending on the frequency being used. Some frequencies can transmit 300 feet to 400 feet, requiring little power to do so. Most older technology cordless phones and wireless NICs use the 900 MHz frequency as a carrier wave, which can travel quite a bit farther than most people realize. It is not uncommon for a 900 MHz cordless phone to give a user at least one or two city blocks of use before the handset loses its connection to the base unit. One or two city blocks translates roughly to 400 feet to 500 feet.

If your telephone handset can transmit as far as 500 feet, it means that your wireless connection is capable of similar distances. If you have a wireless access point (WAP) installed in your office or home, you can bet that people walking by outside are well within its operational envelope. The same holds true if you have a WAP installed in your small office, home office (SOHO) network. If an average WAP is installed in your living room and you live in an apartment complex, you might already be providing Internet service to most of the complex and not even realize it. You can see this concept in action by scanning for any wireless networks rather easily. One of the most creative and innovative ways of doing this is through the use Meraki Wi-Fi Stumbler, a wireless network detector in a web browser, as shown in Figure 10-1, (<http://tools.meraki.com/stumbler/>) or the compilation of Linux tools under Backtrack 5 from www.offensive-security.com/.

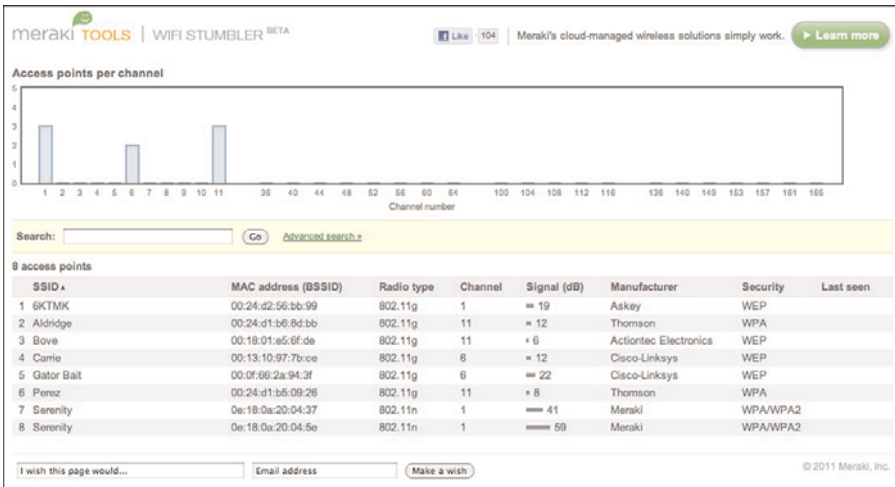


Figure 10-1 Web Based Wi-Fi Network Detector

Wireless Networking

The term *wireless networking* refers to radio technology that enables two or more computers to communicate using standard network protocols such as IP, but without cables. Wireless networking hardware requires the use of underlying technology that deals with radio frequencies and data transmission. The most widely used standard is 802.11, which was produced by the Institute of Electrical and Electronic Engineers (IEEE). This is a standard defining all aspects of radio frequency wireless networking.

802.11b specifies that radios talk on the unlicensed 2.4 GHz band at 11-Mbps transmission rate on one of 15 specific channels. (In the United States, use is limited to only the first 11 of those 15 channels because of government regulations.) Wireless network cards automatically search through these channels to find WLANs, so there is no need to configure client stations to specific channels. When the NIC finds the correct channel, it begins talking to the access point (AP). As long as all the security settings on the client and AP match, communications across the AP can begin, and the user can participate as part of the network.

Note 802.11g is a new high-speed wireless standard that enables users to transmit data at rates of up to 54 Mbps—nearly five times faster than 802.11b technology. Because it operates in the 2.4 GHz frequency band, 802.11g is completely compatible with 802.11b and available for use worldwide. Apple currently has support for 802.11g in all its devices, with Cisco to follow shortly.

Modes of Operation

Two types of wireless networks are possible, and they differ in how wireless devices communicate with each other. WLANs operate either in ad-hoc or infrastructure. Ad-hoc networks have multiple wireless clients talking to each other as wireless peers to share data among themselves without the aid of a wireless access point. An infrastructure WLAN consists of several clients talking to a central device, the an AP, which is usually connected to a wired network such as a corporate or home LAN:

- Infrastructure:** This mode of operation requires the use of a basic service set (BSS); in other words, a wireless access point. The AP is required to enable wireless computers to connect not only to each other but also to a wired network, as shown in Figure 10-2. Most corporate WLANs operate in infrastructure mode because they require access to the wired LAN to use services such as printers and file servers.




Figure 10-2 *Infrastructure Wireless Networking*

- **Ad-Hoc:** Also known as peer-to-peer wireless networking, as shown in Figure 10-3, where a number of wireless computers need to transmit files to each other. This mode of operation is known as independent basic service set (IBSS). You can think of ad-hoc happening without the use of an AP. Each computer can communicate directly with all the other wireless enabled computers. They can share files and printers this way but cannot access wired LAN resources unless one of the computers acts as a bridge to the wired LAN using special software. (This is called *bridging*.)




Figure 10-3 *Ad-Hoc Wireless Networking*

Coverage

Entirely too many wireless access points are available these days to cover them all, so this section focuses on the general coverage levels available. Your mileage might vary, so always check with your manufacturer, and do a little wireless site survey to see what is happening, where you have a good signal, and where you do not.

Every wireless access point has a finite range within which a wireless connection can be maintained between the client computer and the AP. The actual distance varies depending on the environment; manufacturers typically state both indoor and outdoor ranges to give a reasonable indication of reliable performance. Also, when operating at the edge of the range limits, the performance typically decreases because of deterioration of the quality of the wireless signal. Typical ranges are as follows:

- Typical indoor ranges are 150 feet to 300 feet but can be shorter if the building construction interferes with radio transmissions. Longer ranges are possible, but again performance degrades with distance.
- Outdoor ranges are quoted up to 1000 feet, but again, this depends on the location, the environment, and the type of antenna being used.

In most cases, separate APs interconnect via a wired LAN by providing wireless connectivity in specific areas such as offices or classrooms. Depending on the sophistication of

the AP, the range can be modified by adjusting the power level on the AP. This might or might not be an option on some of the lower-end consumer-level APs; however, on many Cisco wireless APs, this is possible. The ranges are commonly 5 mw to 100 mw, which can be a useful method of controlling how far your signal reaches outside your company walls. If the signal goes too far, you increase the risk to your network; too short and you fail to meet the needs of your users or demands of the business.

If a single area is too large to be covered by a single AP, multiple APs can be used. If you choose to go this route, make sure that the APs you want to use have this feature because some do not.

Bandwidth Availability

Bandwidth on an 802.11b network is limited to 11 Mbps per access point. To dispel a lot of confusion, 11 Mbps refers to the *total possible bandwidth* per access point. Many people are used to the wired world, where switches are everywhere and each device gets the full 100 Mbps to the desktop. This is not the case with wireless; the 11 Mbps is divided among all users on that AP. If ten people access the same AP, communication to the wired world will be limited to the equivalent of approximately 1 Mbps per user.

Can you solve the problem by simply adding another access point? I have not used the “it depends rule” since Chapter 6, “Security Protocols,” so its use is way overdue and I am invoking it now. It depends: The 802.11b standard does not contain any specifications for load balancing across multiple APs. Devices that strictly adhere to the standard have no solution to the problem of finding your network becoming overpopulated.

The only way to manage this issue is to add another AP in the same area with a different network name and radio channel, effectively having more than one separate network with a maximum of three in use at the same area. Again, this is if you are using devices that adhere in this regard to the 802.11 standard. In reality, many manufacturers recognized that they would be severely limited in the number of APs they could sell to businesses, so they developed proprietary load-balancing solutions. Additional discussions of these solutions are beyond the scope of this book and should be referred to your wireless vendor.

WarGames Wirelessly

Like many of the beneficial technologies discussed in this book, wireless networks are also susceptible to a variety of threats; however, wireless is still a growing technology, and today you have the opportunity to protect and secure your network. This section takes a high-level look at some of those threats and why you should secure your network.

You might be familiar with the 1983 movie *WarGames*, where a young man (played by Matthew Broderick) finds a back door into a military computer and unknowingly starts the countdown to World War III. The movie’s young hacker executes this mayhem all over a modem, which coined the phrase *wardialing*.

Fast-forward almost 20 years when London-based author Ben Hammersley was writing, and he wanted a cup of coffee or even a bite to eat from the café across the street, but he

still needed to work. Ben installed an access point that gave him the wireless access he wanted; he was a giving man, however, and decided to let his neighbors know that they could have free wireless Internet access as well. Disappointingly, no one took him up on his generosity. Enter Ben's friend, Matt Jones, who posted a set of runes on a website (www.blackbeltjones.com) with the intention of creating a set of international symbols that would let people know that a wireless connection is available. Ben took a piece of chalk and drew these runes on the curb in front of the café and became the first warchalker (see Figure 10-4).




let's warchalk..!	
KEY	SYMBOL
OPEN NODE	ssid  bandwidth
CLOSED NODE	ssid 
WEP NODE	ssid access contact  bandwidth
blackbeltjones.com/warchalking	

Figure 10-4 *Warchalking Symbols*

Shortly after Matt (also known as Black Belt Jones) posted these symbols on the Internet, word spread fast and these two individuals started an Internet phenomenon resulting in new words with such ominous names as warchalking, warspying, warspamming, and wardriving—all ultimately a part of the evolution of wireless access. To clarify, none of these terms enhance the security of your network. They are simply terms that attackers use to describe their activities. The following sections review each of these threats.

Warchalking

If you have ever seen a pirate movie in which a fancifully drawn treasure map displayed a large, red X depicting where the ill-gotten gains were buried, you have some basic idea what role symbology has played in man's pursuit of riches. Much in the same way that the X marked the spot filled with gold, jewels, and silver, so did a series of runes depict areas of danger: which house a policeman might live in, or which houses were considered sympathetic to hobos during the Great Depression. For example, a rune in the shape of the pound sign (#) told fellow hobos that a crime had recently been committed and to avoid the area, or a casually drawn triangle might indicate that there were too many hobos working this area, so pickings were slim.

It was these hobo hieroglyphics from the Great Depression that inspired Ben and Matt to add a new dimension known as *warchalking*. Warchalking is a practice that originated with the intention of telling fellow wireless warriors where they could get a free wireless connection on a corporate or private wireless network. The symbols used by these warchalkers generally indicate whether the wireless AP is considered open or closed, depicted either by two half-circles back to back or a single regular circle, respectively, and what sort of security protects this AP.

Warchalking in its original form turned out to be a momentary cult-like movement that was fascinating for everyone. However, in practice it has changed significantly to reflect the realities of what people are trying to accomplish. Few people walk around drawing marks on buildings; however, people are “chalking” maps using GPSs to show exactly where wireless access can be gained. Searching the Internet reveals quite a few online maps marked for use (www.netstumbler.com/nation.php). One of the added benefits of putting the maps online is that they are not washed away when it rains.

From a security perspective, it is highly unlikely that you will ever *see* the side of your building or sidewalk marked with a warchalk symbol; however, it is likely that if your wireless network is not properly protected, it will appear chalked on someone’s map for anyone to use. You might be wondering how attackers find these APs. Consider the last time you saw anyone walking around with a laptop and a GPS. It does happen, but it might not be obvious because warwalkers typically use backpacks to conceal their activities. In addition to the limitations posed by equipment battery life, all this walking can become tiring. Enter the next wireless threat—wardriving—where converters can power a laptop for as long as the car is running.

Note Wapchalking—A variant of warchalking set up by the Wireless Access Point Sharing Community, an informal group with a code of conduct that forbids the use of wireless APs without permission. The group uses the warchalking marks as an invitation to wireless users to join their community. In warchalking terms, the two half-moon open node mark means that a wireless access device is currently indicating factory default settings and is thus easily detected.

Wardriving

Wardriving makes finding open wireless networks simple and dramatically increases the search area exponentially. The act of wardriving is simple: You drive around looking for wireless networks. Before delving too deeply into this subject, remember that wardriving or LAN jacking an unwary subject’s WAP is possibly illegal, depending on the part of the country in which you live. This book does not imply that you should start security testing outside a sandbox that you own. It merely discusses the technical nature of such a white hat audit. Part of the appeal is that you can now use GPS systems connected to your laptop, which is then powered by your car. This makes the act of wardriving accurate and potentially rewarding for those looking for your wireless network because they can cover a much larger area with a vehicle.

It is disturbing that almost anyone can find your wireless network so easily, isn't it? Vendors turn everything on by default, regardless of network security concerns; this makes it easy for wardrivers. By default, wireless APs broadcast a *beacon frame* that identifies (broadcasts the SSID) the wireless network they are a part of, every 10 milliseconds.

The average antennae on a wireless PCI card NIC is not sensitive enough to do a good job of zeroing in on low- to medium-powered WAP signals, so many wardrivers have resorted to using a USB wireless NIC outfitted with a homemade directional Yagi design antennae hardwired into the USB NIC, as shown in Figure 10-5. Various designs yield better or worse results depending on the signal type of the wireless traffic you are trying to snoop. The wireless network is identified by a 32-bit character known as a *Service Set Identifier (SSID)*. For a wardriver, the easiest networks to find are those broadcasting this SSID. Perhaps you do not have any special applications but only a laptop with Windows. From a security perspective, Windows is wireless-aware and perhaps too friendly because it easily picks up any SSID broadcasts and automatically tries to join any available wireless network. With such a friendly operating system, who needs all the special tools?



Figure 10-5 *Pringles Can Used as a Yagi Antenna*

By default, the SSID is included in the header of the wireless packets broadcast every 10 milliseconds from a WAP. The SSID differentiates one WLAN from another, so all APs and all devices attempting to connect to a specific WLAN must use the same SSID. A device is not permitted to join the wireless network unless it can provide the unique SSID. Because an SSID can be sniffed from a packet in plain text, it does not supply any security to the network, even though it does function as a wireless network password. It is strongly recommended that WAPs have the broadcasting of their SSID disabled.

The presence of an SSID in a wireless network means that those engaging in the search should have more powerful wireless antennas that enable them to pick up and detect wireless signals. For example, if you want to “LAN jack” 802.11b 2.4-GHz wireless network connections, you would most likely opt for a “helix” or “helical” design, which is basically tubular in design with a series of copper wire wrappings around a central core. This custom-made antennae style can be difficult to build because of its exacting standards and rather pricey parts list. On the other hand, a “wave guide” style can be made from rather

inexpensive components such as a Pringles can (as shown in Figure 10-5), coffee can, or juice can.

Depending on your frame of reference (and why you are reading this book), you might be wondering whether wardriving is a crime. Of course, those doing the wardriving do not view it as such; however, those of you who own the wireless networks might have a slightly different perception. While doing research, I stumbled across a quote—supposedly from the FBI—that states its position as follows:

Identifying the presence of a wireless network may not be a criminal violation, however, there may be criminal violations if the network is actually accessed including theft of services, interception of communications, misuse of computing resources, up to and including violations of the Federal Computer Fraud and Abuse Statute, Theft of Trade Secrets, and other federal violations.

Therefore, if you deploy a wireless network, you are likely to have someone try to find it, so your security depends on that individual's understanding that it is his responsibility to ensure that he does not violate any local, state, or federal laws that might pertain to his area. To slightly rephrase: You have gone through all the trouble of purchasing equipment, learning the process, loading the tools, and setting everything up. *Your* wireless network is not secured, and law enforcement expects the wardriver not to do anything illegal. Are you prepared to leave your network vulnerable to those who do not support this law-abiding scenario? If you are, go back to Chapter 1, "There Be Hackers Here," and start reading again!

Warspamming

Everyone has received spam or junk mail; it is a plague on the Internet and, frankly, in my mailbox at home. I believe in free speech; however, that freedom does not give you the right to be heard. Fortunately, lawmakers and politicians around the world are beginning to notice our feelings on this matter and developing laws to penalize spammers. These laws might or might not be effective—time will tell. However, it is becoming more difficult for spammers to source their spam from countries beginning to develop these laws. Also organizations list IP addresses of places where spam has originated from, so what is a spammer to do? Many are now sourcing their spam from other countries; this presents all sorts of logistical problems and additional costs to spammers. As a spammer, what if I could drive downtown or hire someone to find an open wireless network, join that network, and send my spam?

Remember the concept of downstream liability discussed in Chapter 5, "Overview of Security Technologies?" It would be simple to find an open wireless network and join it to send spam. The attacker (spammer) could be sitting in a café across the street, and you might never know. Now fast-forward a bit; the spam is sent to thousands of people who report that they received it, and yet another wrinkle—the spam was pornographic in nature. Yes, it can be even worse than that. (Remember, we are not talking about people who have morals—they are driven by other goals and needs.) A quick check reveals your

network's IP address, which is then blacklisted and reported to your ISP—and do not forget about the new antispamming laws. The result is that all outgoing email from your company is blacklisted. How embarrassing when *your* customers get the bounce message saying that your company is spamming, the ISP shuts off your Internet connection, and law enforcement comes knocking. Also, if you have one of those Internet connections where you are billed by usage, expect a *big* bill this month.

The truth of the matter in warspamming is that your network did spam others, and although it might have been as a result of an attacker, you are now liable because your wireless network was not properly secured. Who do you think is responsible for that, and are they looking for a new job? Expect to see warspamming increase as it becomes more difficult for spammers to operate. Those who want to do questionable things will always find a way; some will stop as it becomes too difficult, and others will not.

Warspying

A nice follow-up to warspamming is warspying, which is a relatively new phenomenon coming to a wireless video network near you. The most popular method of warspying is using those wireless X10 cameras. X10 is the camera featured in pop-up ads all over the Internet, and they invariably have some gorgeous woman in them. X10 is also a means by which to automate your home, as in a smart house; however, that topic is beyond the scope of this book.

Warspying was first documented in the magazine *2600*, an interesting read if you can find the few nuggets of technical worth from the rants it prints. Regardless, it outlined how to make a wireless device that can pick up wireless surveillance systems transmissions. Since then, many people have explored and documented the topic online, and there are now reports of people tapping into all sorts of cameras that are transmitting over a wireless network.

Notice I have completely avoided all discussions of the other nefarious uses into which this could develop. The key is *awareness* and an understanding of how to protect your network.

This section was rather revealing about how wireless networks are found and, to a lesser degree, what some of the threats are. In addition, a variety of more specific threats are possible. Plus, after an attacker joins a wireless network, you have a host of other problems. The following sections examine these topics in more detail.

Wireless Threats

Wireless threats come in all shapes and sizes, from someone attaching to your Wireless Access Point (WAP) without authorization, to grabbing packets out of the air and decoding them via a packet sniffer. Many wireless users have no idea what kinds of danger they face merely by attaching a WAP to their wired network. This section discusses the most common threats faced by adding a wireless component to your network.

The airborne nature of WLAN transmission opens your network to intruders and attacks that can come from any direction. WLAN traffic travels over radio waves that the walls of

a building cannot completely constrain. Although employees might enjoy working on their laptops from a grassy spot outside the building, intruders and would-be hackers can potentially access the network from the parking lot or across the street using the Pringles can antenna (refer to Figure 10-5).

Sniffing to Eavesdrop and Intercept Data

Because wireless communication is broadcast over radio waves, eavesdroppers who merely listen to the wireless transmissions can easily pick up unencrypted messages. Unlike wire-based LANs, the wireless LAN user is not restricted to the physical area of a company or to a single WAP.

But what if you don't have wireless in your network? Not every organization has wireless activated for their users and might feel there is no need to address wireless security. However, almost 100 percent of the laptop computers purchased today come equipped with wireless that is on by default. A hacker could be outside the organization and configure a wireless access point to be wide open; in many cases, these laptops associate by default, enabling the hacker to attack the laptop and gain entry. The moral to the story is that even if you don't have wireless, you need wireless security and an up-to-date wireless security policy to define those often gray boundaries.

The range of a wireless LAN can extend far outside the physical boundaries of the office or building, thereby permitting unauthorized users access from a public location such as a parking lot or adjacent office suite. An attacker targeting an unprotected AP needs only to be in the vicinity of the target and no longer requires specialized skills to break into a network. Any time I do a network assessment for a customer in a shared office building, I almost always find

- A neighboring business that has an open wireless network
- A neighboring user that has joined my customer's wireless network
- One of my customer's employees using their neighbor's wireless

If you want to examine the traffic going out over an Ethernet connection (wired or wireless), the best tool that comes to mind is the ubiquitous *packet sniffer* application. Packet sniffers enable the capture of all the packets going out over a single or multiple Ethernet connection for later inspection. These sniffer applications grab the packet, analyze it, and reveal the data payload contained within. The theft of an authorized user's identity poses one the greatest threats. Figure 10-6 shows a freeware packet sniffer known as Ethereal, which is used on an Apple MacBook Pro over a wireless Ethernet network to capture a mail application transmitting a username and password. (Names and passwords have been changed to protect the innocent, of course.)

The intent here is to show you how packet sniffers can be used against known behavior. In this case, when users start their computers, one of the first things they do is check email.

Many email servers do not require any sort of encryption and, because the wireless network is not transmitting anything encrypted, the data is sent in clear text. Attackers with a packet sniffer could now *steal the user identity* and log in to the mail server as the unaware user anytime because they literally pulled the password out of the air.

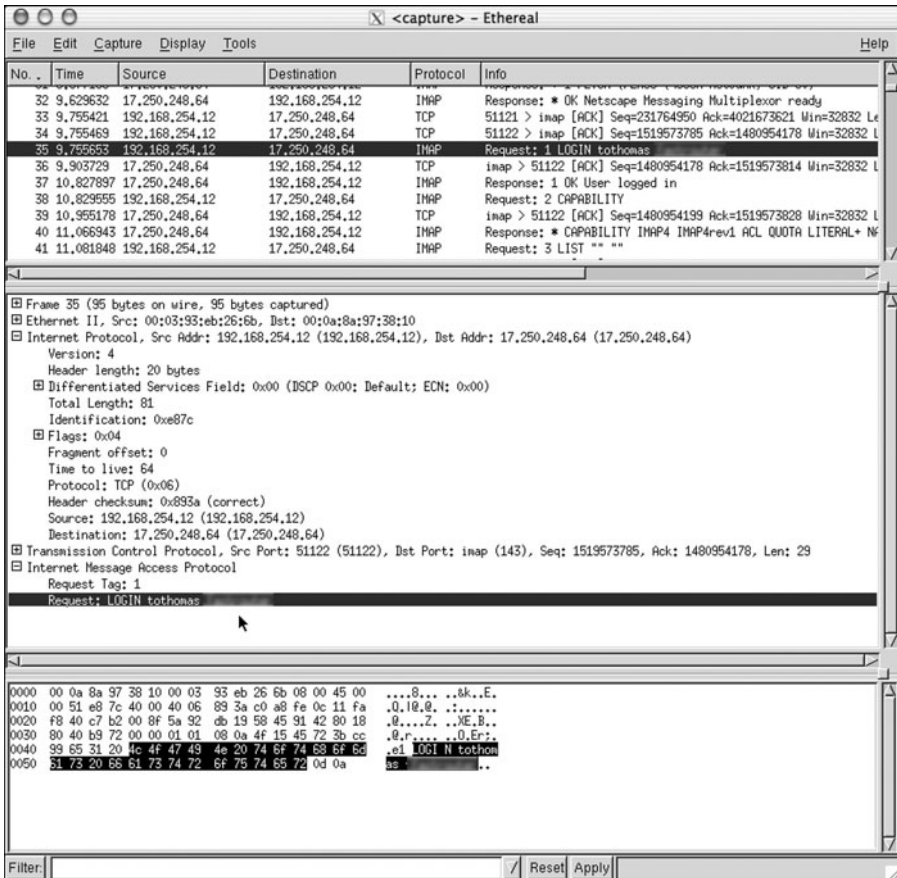


Figure 10-6 Wireless Sniffer Packet Capture

If you have read through packet captures before and are familiar with the information they contain, you should have immediately recoiled in horror at the knowledge that wireless networks are sniffers readily available, and several are free. If this is the first time you have seen a packet capture, you might be in for a shock as you find out the wealth of information contained in a packet's data payload. Imagine if you were a domain administrator logging in to the domain and checking your online bank account or other information that could be critically damaging if someone hijacked it.

Denial-of-Service Attacks

Potential attackers who cannot gain access to your wireless LAN can nonetheless pose security threats by jamming or flooding your wireless network with static noise that causes wireless signals to collide and produce CRC errors. Wireless networks are especially vulnerable to these sorts of attacks, sometimes even unintentionally as *every* wireless network shares the same unlicensed frequencies (channels). These denial-of-service (DoS) attacks effectively shut down or severely slow down the wireless network in a similar way that DoS attacks affect wired networks. Sometimes a DoS is not malicious; it could be a wireless phone that is set on the same frequency causing interference or a microwave oven. Sometimes, though, it could be phony messages to disconnect users or consume AP resources.

This vulnerability is apparent, and being on a wired network does not reduce your vulnerability to viruses, attacks, or in any other way increase security; it will quite likely get worse. However, newer wireless standards such as 802.11n use the 5 GHz frequency, which is a larger range of channels and less crowded, reducing the chance of accidental service interruptions due to channel overlap.

Note Restaurants, hotels, business centers, apartment complexes, and individuals often provide wireless access with little or no protection. In these situations, you can access other computers connected to a wireless LAN, thereby creating the potential for unauthorized information disclosure, resource hijacking, and the introduction of backdoors to those systems. When users take corporate laptops home and use them on wireless networks, the vulnerabilities to your network increase. I have been on network assessments reviewing wireless usage and found that many a CEO, CFO, or CTO has the IT staff set up a wireless device at home for them with the same characteristics they have at work (SSID and so on). This makes it easy for them to work at home with no trouble; however, the corporate network is extremely vulnerable because an attacker can go after a corporate employee's home network and compromise his machine. When the employee goes to work, so does the attacker—now he is inside your corporate network. Common sense is needed here—and a commitment by everyone in the organization's management team to secure the network. This means not mixing corporate and home security regardless of how much fussing that C level may do; security is bigger than the individual.

Perhaps a bit more common is when other wireless devices unintentionally cause a DoS to your wireless data network—for example, that new cordless phone running on 2.4 GHz, or placement of APs near devices that generate interference and affect their operation, such as microwaves. Not all reduction in wireless connectivity is related to attackers, so remember that wireless networks are based on radio signals, and many things (walls, weather, and wickedness) can affect them.

Rogue/Unauthorized Access Points

Wireless APs can be easily deployed by anyone with access to a network connection, anywhere within a corporation or business. Most wireless deployments are in the home, so people with laptops can use them in any room in the house. The ease with which wireless technologies can be deployed should be a concern to all network administrators.

Because a simple WLAN can easily be installed by attaching a WAP (often for less than \$100) to a wired network and a wireless enabled laptop, employees are deploying unauthorized WLANs while IT departments are stuck trying to track down these rogues. Unauthorized WAPs are known more commonly as *rogue APs*.

An executive of a large technology conglomerate was recently quoted as saying something like, “The hardest network to secure against wireless threats was one that had no wireless access at all.” What this executive meant was that just because a company did not buy and install any wireless gear on its network did not mean that there wasn’t any.

The concept behind wireless technology is to give people the freedom to roam around and still be connected to their network resources. The lure of this freedom is just too tempting to some folks in corporate America, so they go out and buy wireless gear on their own and hook it up to the office network. Now, you begin to see the problem.

If you can imagine how difficult it is to prevent people from bringing software from home and installing it on their work machines, it is ten times more difficult to prevent power users from “self-adopting” wireless gear into the office LAN.

You might ask, “What is the harm in doing this?” The harm is that by installing an unauthorized AP, you have now extended an invitation to every person within its signal radius to prowl your company’s network, files, Internet access, printers, and any other devices currently connected to the private corporate network.

Your network administrators take great pains to protect the corporate network from attackers and other evildoers, and now there is a completely unprotected conduit into the company’s holiest of holies: your internal corporate network.

A well-documented company has several security policies in place that govern every type of behavior when a user connects to the network. Rogue APs subvert these policies and open the doors to all varieties of bad things happening to the network.

To be perfectly fair to the employees who might commit this wireless breach of security, it is important that the following information be made abundantly clear:

- Only authorized IT staff is allowed to connect networking equipment.
- All devices that connect to the network, especially wireless APs, must conform to established security policies.
- Any devices that have been installed by anyone other than approved IT staff will become either the property of the company or will be rendered inert (that is, smashed into a million pieces).

- Hackers install rogue APs on a company network with the intention of stealing secrets and damaging data; this means no holiday bonuses because this kind of damage can cause a company to go out of business.

Finding rogue APs has become a little easier than in the past through the use of freely available software; the section titled “NetStumbler” delves into this. This same piece of software that made life easier for hackers has now become the favored tool of network security specialists for dealing with unauthorized wireless access points.

Misconfiguration and Bad Behavior

Wireless APs are typically centrally managed in today’s enterprise networks; however, they are slow in catching up with technology. The latest version of 802.11 has evolved to include many new features that have resulted in relatively complex configuration options. Add to this the inherent capability of laptops to create ad-hoc networks via peer-to-peer technologies. The latest versions of Windows operating systems have removed the complexity involved with ad-hoc wireless networking, bypassing network security procedures automatically.

AP Deployment Guidelines

I was going to call these “the rules for attackers to deploy rogue access points,” but applying rules to those with criminal intent seemed an oxymoron. Attackers have developed some best practices that they have shared in their community because many wireless networks are relatively easy to break into. It is important that any wireless deployment use effective and efficient wireless security techniques and policies. In addition to using the best encryption and practices defined in this chapter, wireless intrusion prevention systems (WIPS) and wireless intrusion detection systems (WIDS) are commonly used to verify and protect the integrity of wireless networks. Following is a brief list of what you can do to prevent attackers from “casing the joint”:

- Know what you are trying to gain before placing the access point.
- Plan for the use of the AP; this means place it so that if you have your laptop out and working, you do not look suspicious.
- Place the AP as discretely as possible while maximizing your ability to connect to it.
- Disable SSID broadcasting, thus requiring the target’s IT staff to have a wireless sniffer to detect it.
- Disable all network management features of the AP, such as SNMP, HTTP, and Telnet.
- If possible, protect the AP’s MAC address from appearing in ARP tables.

The obvious disclaimer here is that these actions are not something you should ever do without—and I *really* stress this—*written permission*. Many companies view even the

accidental connection to their wireless network as an attack, so it is likely that you are going to be viewed as guilty until you prove your innocence.

It is also important to note that devices designed to jam radio signals have been around since before wireless ever became a standard. Because wireless is a radio frequency, it can be easily jammed with a simple transmitter purchased online.

Wireless Security

You might be wondering why someone would want to use a wireless connection with all the insecurities that seem to go along with it. All is not lost, thanks to the focus that has been placed on securing wireless networks.

From its inception, the 802.11 standard was not meant to contain a comprehensive set of enterprise-level security tools. Still, the standard includes some basic security measures that can be employed to help make a network more secure. With each security feature, the potential exists for making the network either more secure or more open to attack.

Working on the layered defense concept, the following sections look first at how a wireless device connects to an AP and how you can apply security at the first possible point.

Service Set Identifier (SSID)

By default, the AP broadcasts the SSID every few seconds in beacon frames. Although this makes it easy for authorized users to find the correct network, it also makes it easy for unauthorized users to find the network name. This feature is what enables most wireless network detection software to find networks without having the SSID upfront.

SSID settings on your network should be considered the first level of security and should be treated as such. In its standards-adherent state, SSID might not offer any protection against who gains access to your network, but configuring your SSID to something not easily guessable can make it more difficult for intruders to know what exactly they are seeing.

Finding nearby SSIDs, even if they are not broadcasting, is relatively easy. One of my favorite tools is from a wireless company known as Meraki. It offers an online web browser-based Wi-Fi Stumbler that will find nearby SSIDs, as shown in Figure 10-7. If you look, you can see that several of those listed are running WEP, which, as we have discussed, is foolish. This tool also provides helpful information such as channel, signal strength, and radio manufacturer/type. As a network administrator, this is extremely helpful and extremely convenient; I especially like the “I wish this page would” feature... now that is customer support!

A complete listing of manufacturers’ SSIDs and even other networking equipment default passwords can be found at www.cirt.net/.

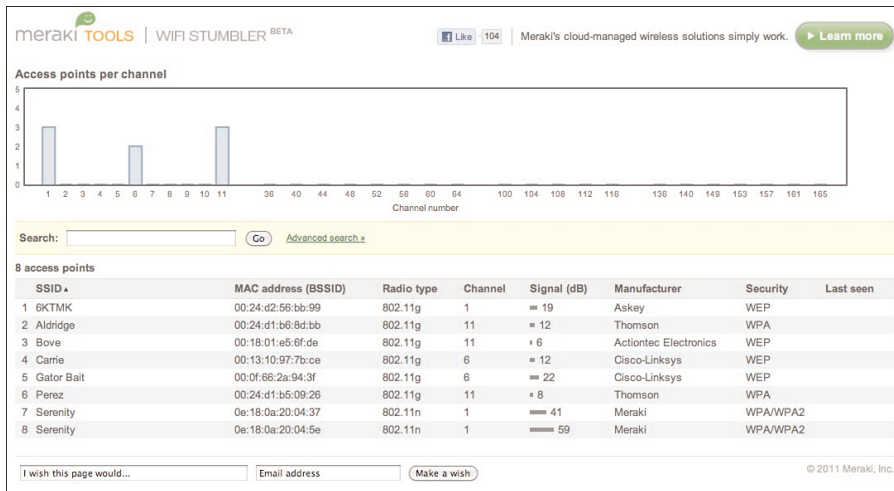


Figure 10-7 Web-Based SSID

Device and Access Point Association

Before any other communications take place between a wireless client and a wireless AP, the two must first begin a dialogue. This process is known as *associating*. When 802.11b was designed, the IEEE added a feature to enable wireless networks to require authentication immediately after a client device associates with the AP, but before the AP transmission occurs. The goal of this requirement was to add another layer of security. This authentication can be set to either *shared key authentication* or *open key authentication*.

You must use open key authentication because shared key is flawed; although that is counterintuitive, this recommendation is based on the understanding that other encryption will be used. Wireless network administrators need to be aware that accidental or malicious association is a risk that needs to be managed. A user turns on his laptop and unknowingly associates with a neighboring organization's wireless network; the user might not even be aware this has occurred. Malicious association is when a hacker uses this accidental association to gain access to your network by taking over a client and planting a tool to enable him to gain deeper access.

Wired Equivalent Privacy (WEP)

There is a lot of misconception surrounding WEP, so let's clear that up right away. WEP is not, nor was it ever meant to be, a security algorithm. WEP was never designed to protect your data. WEP is not designed to repel attackers; it simply makes sure that you do not transmit everything in clear text. The problem occurs when people see the word *encryption* and make assumptions. WEP is designed to make up for the lack of security in wireless transmission, compared to wired transmission; however, it should never be used to secure your wireless networks.

WEP Limitations and Weaknesses

WEP protects the wireless traffic by combining the “secret” WEP key with a 24-bit number (Initialization Vector, or IV), randomly generated, to provide encryption services. The 24-bit IV is combined with either the 40-bit or 104-bit WEP passphrase to give you a possible full 128 bits of encryption strength and protection—or does it? There are a few issues surrounding the flawed current implementation of WEP:

- WEP’s first weakness is the straightforward numerical limitation of the 24-bit Initialization Vector (IV), which results in 16,777,216 (2^{24}) possible values. This might seem large, but you know from discussions in Chapter 6 that this number is deceiving. The problem with this small number is that eventually the values and thus the keys start repeating themselves; this is how attackers can crack the WEP key.
- The second weakness is that of the possible 16 million values, not all of them are good. For example, the number 1 would not be very good. If an attacker can use a tool to find the weak IV values, the WEP can be cracked.
- WEP’s third weakness is the difference between the 64-bit and 128-bit encryption. Perception would indicate that the 128 bit should be twice as secure, right? Wrong. Both levels still use the same 24-bit IV, which has inherent weaknesses. Therefore, if you think going to 12 bit is more secure, in reality, you will gain absolutely no increase in the security of your network.

Of course, freely available tools can accomplish all these things and are ready for the attackers to download and use as discussed in the section “Essentials First: Wireless Hacking Tools,” later in the chapter. Using WEP is *not* advised, and if you run across a network running WEP, buy them a copy of this book, and point out this chapter to them for me!

MAC Address Filtering

MAC address filtering is another poor and unsuccessful way people have tried to secure their networks over and above the 802.11b standards. A network card’s MAC address is a 12-digit hexadecimal number that is unique to every network card in the world. Because each wireless Ethernet card has its own individual MAC address, if you limit access to the AP to only those MAC addresses of authorized devices, you can easily shut out everyone who should not be on your network.

However, MAC address filtering is not completely secure and, if you rely solely upon it, you will have a false sense of security. Consider the following:

- Someone must keep a database of the MAC address of every wireless device in your network. If there are only 10–20 devices, it is not a problem. However, if you must keep track of hundreds of MAC addresses, this quickly becomes a management nightmare.
- MAC addresses can be changed, so a determined attacker can use a wireless sniffer to figure out a MAC address that is allowed through and set his PC to match it to con-

sider it valid. Note that encryption takes place at about Layer 2, so MAC addresses will still be visible to a packet sniffer.

If you are thinking of using MAC address filtering as your sole means of security, that is a bad idea because it provides a false sense of security and prevents only unintended connections, not a directed attack. This form of wireless security should be used only with one of the methods covered in the following sections.

Extensible Authentication Protocol (EAP)

802.1X is a standard for port-level security that the IEEE ratified and updated several times. This ratification was initially intended to standardize security on wired network ports, but it was also found to be applicable to wireless networking.

Extensible Authentication Protocol (EAP) is a Layer 2 (MAC address layer) security protocol that exists at the authentication stage of the security process and, coupled with the security measures discussed thus far, provides a third and final layer of security for your wireless network. Using 802.1X, when a device requests access to the AP, the following steps occur with EAP:

1. The access point requests authentication information from the client.
2. The user then supplies the requested authentication information.
3. The AP then forwards the client supplied authentication information to a standard RADIUS server for authentication and authorization.
4. Upon authorization from the RADIUS server, the client is allowed to connect and transmit data.

Note Not everyone has a RADIUS server that is ready to use LEAP; however, Cisco APs can be configured with a feature called local AAA Authentication on a per-user basis. This enables the user database to reside in the AP instead of RADIUS and works well if you have only a limited number of users.

More than a dozen different types of EAP are available, making for a complicated set of choices. The four most commonly used EAP methods in use today follow:

- Lightweight Extensible Authentication Protocol (LEAP)
- EAP-TLS (Transport Layer Security)
- EAP-PSK (Pre-Shared Key)
- EAP-TTLS (Tunneled Transport Layer Security)

The following sections provide a quick overview of each EAP method.

LEAP

EAP-Cisco Wireless, or LEAP as it is more commonly known, is a standard developed by Cisco with the 802.1X standard and is the basis for much of the ratified version of EAP. Like EAP-MD5, LEAP accepts a username and password from the wireless device and transmits them to the RADIUS server for authentication. Cisco added additional support beyond what the standard required, resulting in several security benefits as follows:

- LEAP authenticates the client; one-time WEP keys are dynamically generated for each client connection. This means that every client on your wireless network is using a different dynamically generated WEP key that no one knows—not even the user.
- LEAP supports a RADIUS feature called *session timeouts*, which requires clients to log in again every few minutes. Fortunately, this is all handled without the user needing to do anything. Couple this feature with dynamic WEP keys, and your WEP keys change so often that attackers have a difficult time determining the key.
- LEAP conducts mutual authentication from client-to-access point and access point-to-client; this stops attackers from introducing rogue APs into your network.

There is a known limitation to running LEAP. MS-CHAPv1 is used for both the client and AP authentication and is known to have vulnerabilities; definitely look at alternatives to anything Microsoft thinks is secure. LEAP can be cracked with *asleep*, written by Joshua Wright and available at www.willhackforsushi.com, so you might want to consider stronger wireless security than LEAP.

Note Extensible Authentication Protocol (EAP) is a widely used method of authenticating; EAP is more of a format than a process. With EAP as the framework, many additional authentication methods are built upon it.

EAP-TLS

Microsoft developed EAP-TLS, which is outlined in RFC 2716. Instead of username/password combinations, EAP-TLS uses X.509 certificates to handle authentication. EAP-TLS relies on transport layer security to pass PKI information to EAP. Like LEAP, EAP-TLS offers the following:

- Dynamic one-time WEP key generation
- Mutual authentication of the client and the network

The drawbacks of EAP-TLS include the following:

- PKI is required to use EAP-TLS; however, most companies do not deploy PKI.
- Microsoft Active Directory with a certificate server can be used; however, change is difficult in this model.

- If you use Open LDAP or Novell Directory Services, you need a RADIUS server; again, not everyone has immediate access to one.
- If you have implemented PKI using VeriSign certificates, all the fields required by EAP-TLS are not present.

Unless you are ready to follow the implementation of EAP-TLS exactly as Microsoft has laid it out, you should probably look for another method.

EAP-PSK

Pre-Shared Keys (PSK) are a part of this EAP method of authentication that was designed for use in wireless networks. When using EAP-PSK, an encrypted method of communicating is used, via AES, to ensure the integrity and authentication is successful.

EAP-TTLS

Funk Software (now part of Juniper Networks) pioneered EAP-TTLS as an alternative to EAP-TLS. The wireless access point still identifies itself to the client with a server certificate, but the users now send their credentials in username/password form. EAP-TTLS then passes the credentials in any number of administrator-specified challenge-response mechanisms (PAP, CHAP, MS-CHAPv1, MS-CHAPv2, PAP/Token Card, or EAP). The only challenges to EAP-TTLS are

- They are slightly less secure than dual certificates of EAP-TLS.
- Protected EAP (PEAP) is the newer version championed by Cisco, RSA, and Microsoft.

Essential Wireless Security

As discussed, there are some possible means of securing your wireless network beyond WEP. It is unlikely, however, that anyone has a RADIUS server ready and waiting to be used; therefore, you need to identify steps you can take immediately to increase the security of your wireless network. The attention to the pitfalls of wireless LANs has inspired some organizations to ban wireless LANs altogether. However, security-conscious organizations are fortifying their wireless LANs with a layered approach to security that includes the following:

- Putting the wireless network behind its own routed interface so that you can shut off access at a single choke point if necessary.
- Regular monitoring and discovery of rogue access points and potential associated vulnerabilities of nearby APs.
- Physical and logical AP security to ensure that someone cannot walk up to an access point and alter its configuration without your knowledge.

- Changing the SSID and then picking a random SSID that gives away nothing about your company or network.
- Disabling active SSID broadcasting, making your wireless invisible to prevent accidental association is important; however, decloak can be done with KisMET and other tools. If attackers have the time (which they usually do), it won't take long to decloak your SSID. KisMET is on the BackTrack 5 distribution and was cited earlier in this chapter as an excellent tool for this.
- Configure your AP to rotate encryption keys every ten minutes or less.
- Encryption and authentication, which might include a virtual private network over wireless.
- Using 802.1X for key management and authentication.
- Looking over the available EAP protocols and deciding which is right for your environment.
- Setting the session to time out every ten minutes or less.
- Establishing and enforcing wireless network security policies.
- Implementing proactive security measures that include wireless intrusion protection and detection in a layered approach is important.
- Deploying smart cards or tokens is a strong form of security and can be used with other security needs in your network, such as VPN access. However, the drawback here is that these solutions can be expensive.
- Regular software updates that keep the AP software up to date and regularly patched.

As shown in Figure 10-8, these steps and recommendations can be illustrated as a phased approach, which enforces the concept of first knowing what the vulnerabilities are and moving forward from that point.

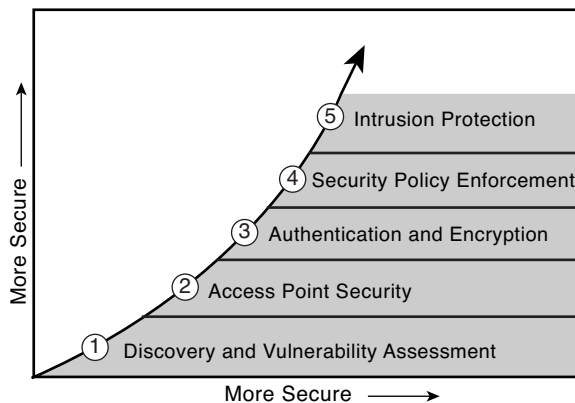


Figure 10-8 *Stages of Securing Your Wireless Network*

Essentials First: Wireless Hacking Tools

This section examines some of the tools that eliminate some of the threats discussed in the preceding sections. In theory, these tools were all designed to help network administrators take care of their networks, and they are still touted as such on each website. In reality, these are some of the same tools that attackers can and will use; thus, network administrators should also use them to ensure that their wireless networks are secure.

NetStumbler

Wireless networking is everywhere! That is not meant as hyperbole—it really is *everywhere*. Wireless technology uses radio waves to transmit data, so wireless packets are probably flowing in the air in front of you as you read this.

As everyone knows by now, where wireless packets flow, wireless APs are pumping them out. (Where there is smoke, there is fire.) If only there were a way to find out whether any WAPs were nearby. Fortunately (and unfortunately), there is a way to discover just that.

A little piece of freeware called NetStumbler is available on the Internet (www.netstumbler.com/) that provides you with such secret pieces of information as the following:

- WAP's Service Set Identification (SSID), the unique name you can assign to your WAP
- Signal strength of the discovered WAPs and whether the WAP uses WEP
- What channel the WAP transmits on, and some other sneaky bits of information

You might have even seen NetStumbler make an appearance on the local evening news under the headline, “Wireless Security Threats: You Could Be Next!” or some other scary tagline. Figure 10-9 shows the NetStumbler interface.

NetStumbler sends out a broadcast on all channels looking for a response. If your WAP is configured to respond to the broadcast (SSID broadcast “enabled” setting), NetStumbler logs that WAP and furnishes you with a “bing-bing” tone designating a target. The trick is that NetStumbler tells you all the information you need about someone else's wireless network.

Most wireless NIC configuration programs enable you to perform a *site survey*, which sniffs around for other wireless access points configured to broadcast on the same channel as your NIC. If you happen to find a WAP with the default SSID (in this case, the default SSID of a Linksys WAP is linksys) displayed, you can assume that you can connect to that WAP with little or no trouble.

One of the best features about NetStumbler is its capability to integrate laptop-based GPS units into its WAP discovery adventure. Imagine driving along with your trusty laptop on the passenger seat of your privately owned vehicle (POV) and hearing the pleasant “bing-bing” tones generated by NetStumbler as it happily sniffs out WAPs within transmitting distance. Every time your laptop makes that sound, NetStumbler queries the attached GPS

unit and records the coordinates of the WAP it found. Later, you can download the coordinates into mapping software and have a nice, little map printed out to show you where the WAPs were found. And who says technology doesn't make our lives just a wee bit more interesting?

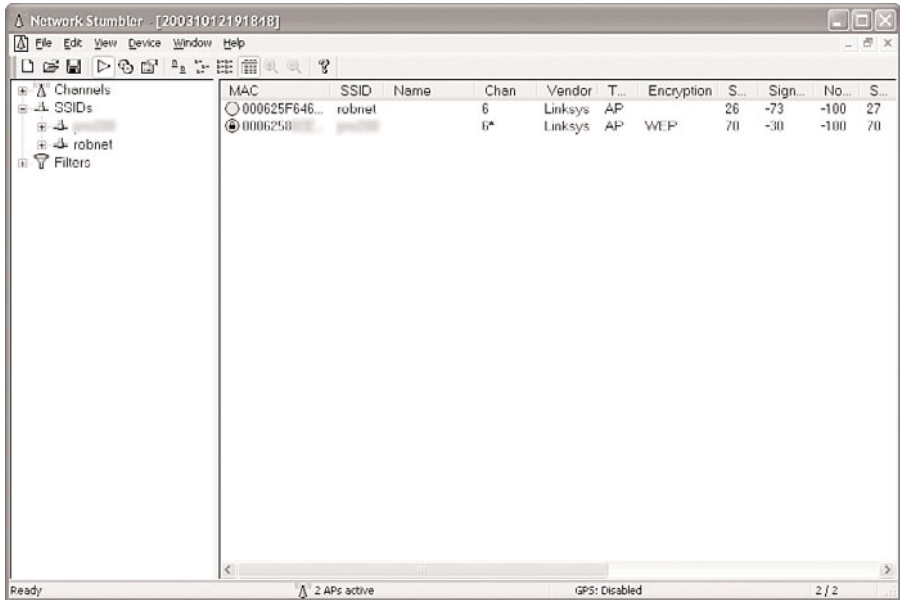


Figure 10-9 *NetStumbler Scanning*

The whole GPS issue aside, NetStumbler is not actually a hacking tool because the information it reveals is just a step above what your NIC can already help you find out. Tools such as NetStumbler are more along the lines of “reconnaissance” tools because they help you discover things that might not have been immediately obvious. NetStumbler is a chatty tool and recon is often done with passive recon tools such as KisMET.

Wireless Packet Sniffers

Sniffing packets can be both fun and profitable if you know how and what to sniff. Any network administrator can lay his hands on a packet sniffer in a matter of seconds and snag a couple hundred packets before you can even read this paragraph. The contents of these packets can reveal network secrets that have been closely guarded. *Sniffing*, or *snaarfing* in the hacker world, is the process of intercepting and recording traffic that was never supposed to be seen by anyone other than the sender or receiver.

To the layman, the idea of sniffing, capturing, or snagging packets is an often misunderstood concept; therefore, the basics of the operation deserve some brief discussion:

1. Packets travel over an Ethernet connection from source to destination.

2. A NIC set to promiscuous mode can listen in on all local traffic.
3. A packet sniffer can see and record all this traffic.
4. A packet sniffer can also decode the packet and display neat things such as the source MAC address, the destination MAC address, and the data payload contained in the packet.
5. Packets contain things such as unencrypted Windows passwords, logins or password combinations sent in clear text, account numbers, and other tasty things relished by hackers.

Now that you know about wired packet sniffers, you also need to meet their wireless cousins. How is this possible, you ask? Can I actually capture wireless packet traffic? Could it be that easy? Do hackers know about this? The answers are, yes, yes, and *yes*. Capturing packets in a wireless network is actually *much* easier than in a wired network because wireless packets are all around in the air, and you don't need to be physically sitting on a wired segment.

Yes, hackers know about sniffing wireless connections, and they have made the most of it. Have you turned on a MAC filter on your WAP? Packet captures rat you out by telling the hacker the MAC address's source. It is easy to spoof a MAC address on your wireless NIC, especially with a program called SMAC, lovingly created by a group of guys at KLC Consulting. If hackers "sniff" your wireless packets, they can decode the packets, read the MAC address of a machine listed in the WAP's MAC filter, plug that number in SMAC, and impersonate a machine authorized to use the WAP. It can do all this in less than 1 minute. That is correct—60 seconds. In the time it takes to dip a chip in salsa and eat it, a hacker can intrude on your network.

Aircrack-ng

Aircrack-ng is an 802.11 WEP and WPA-PSK keys cracking program that can recover keys after enough data packets have been captured. It implements the standard FMS attack along with some optimizations such as KoreK attacks and the all-new PTW attack, thus making the attack much faster compared to other WEP-cracking tools. Aircrack-ng is a set of tools for auditing wireless networks.

You can learn more about this product and company online at www.aircrack-ng.org and in the BackTrack 5 distribution.

OmniPeek

Wireless networks require the same kinds of analytical and diagnostic tools as any other LAN to maintain, optimize, and secure network functions, with one notable exception. In a LAN environment, all signals are conducted over fixed, well-defined, and "electrically stable" network of cables. This is in stark contrast to wireless networks, where signals transmit using radio frequency (RF) technology. Radio frequency waves propagate outward in all directions from their source and are sensitive to disruption or interference. The

quality of the transmitted signal varies over time and space, even if the source and destination remain fixed. The path between the source and destination also has a significant impact on the quality of the resulting communication. Open propagation of data means that anyone can receive the data, even those not “connected” to the network, making security a far bigger issue for WLANs. The use of unlicensed spectrum by 802.11 also increases its vulnerability to interference because it must share its available bandwidth with non-802.11 devices, including Bluetooth, cordless telephones, and microwave ovens.

Fortunately, the 802.11 WLAN standard offers even more data to packet analysis than any of the other members of the 802 family of protocols. WildPackets products enable the creation of highly flexible, cost-effective wireless network analysis solutions. OmniPeek is a comprehensive wired and wireless network analyzer with complete support for IEEE 802.11 wireless LAN protocols. Real-time expert analysis provides an advanced set of expert troubleshooting and diagnostic capabilities.

Features include the following:

- Full 802.11 WLAN protocol decodes
- Multi-NIC support
- Distributed operation with wireless probes or AP capture adapters
- Display of data rate, channel, and signal strength for each packet
- SSID tree of nodes
- Expert analysis of network performance in real time, including VoIP expert diagnoses and wireless problem events
- Designation of nodes as Trusted, Known, and Unknown identifies rogue APs easily
- Expert ProblemFinder settings that include description, possible causes, and possible remedies
- Peer Map, which is a continuously updated graphical view of traffic between pairs of network nodes, showing volume, protocol, node address, and node type
- Alarms, triggers, and notifications, all user-definable
- Security audit template with predefined security audit filters
- Scan/surf by channels, ESSID or BSSID
- VoIP analysis tools
- Application performance tools
- Forensics analysis

You can learn more about this product and company online at www.wildpackets.com/products/network_analysis_and_monitoring/omnippeek_network_analyzer.

Wireshark

Wireshark is the world's foremost network protocol analyzer. It enables you to capture and interactively browse the traffic running on a computer network. It is the de facto (and often de jure) standard across many industries and educational institutions.

Wireshark development thrives because of the contributions of networking experts across the globe. It is the continuation of a project that started in 1998.

Wireshark has a rich feature set that includes the following:

- Deep inspection of hundreds of protocols, with more being added all the time
- Live capture and offline analysis
- Standard three-pane packet browser
- Multi-platform: runs on Windows, Linux, OS X, Solaris, FreeBSD, NetBSD, and many others
- Captured network data can be browsed via a GUI or via the TTY-mode TShark utility
- The most powerful display filters in the industry
- Rich VoIP analysis
- Decryption support for many protocols, including IPsec, ISAKMP, Kerberos, SNMPv3, SSL/TLS, WEP, and WPA/WPA2
- Coloring rules can be applied to the packet list for quick, intuitive analysis

You can learn more about this product and company online at www.wireshark.org/.

Note You can use other wireless tools, such as KisMET and KisMAC, which are wireless AP locators and include support for GPS location and positioning, to create maps of all known, open wireless APs in a city, building, or your neighborhood.

Chapter Summary

This chapter has hopefully shed some light on the technology that drives wireless and the first steps for beginning to secure a wireless network. You should be concerned about a variety of areas surrounding wireless; however, you can apply clear, layered steps to secure a wireless network with minimal impact to users. Of utmost importance are the steps you take today to increase security that will not hamper or affect the security of your wireless network. Ultimately, though, it is the responsibility of the IT department to keep users up to date on dangers and techniques to keep themselves and the network safe and secure. User awareness training and current policy is critical in this area of information security.

The chapter concluded with a discussion of the *freely* available tools relating to attacking and securing wireless networks. Attackers commonly use these tools; more important,

however, those who want to find flaws in their wireless network security *should* use them to patch them up and prevent easy attacks.

Chapter Review Questions

1. How are the terms 802.11 and Wi-Fi used? In what ways are they different or similar?
2. What are the five benefits to organizations that would provide reasons for them to implement a wireless network?
3. Wardriving is the most common means of searching for wireless networks. What is needed to conduct a wardrive, and why is it so useful for attackers?
4. What is one type of freely available wireless packet sniffer?
5. Are wireless networks vulnerable to the same types of denial of service attacks as wired networks? Are they vulnerable to any additional attacks that wired networks are not?
6. What are the four most common types of EAP available for use?

Intrusion Detection and Honeypots

...I think computer viruses should count as life. I think it says something about human nature that the only form of life we have created so far is purely destructive. We've created life in our own image....—Stephen Hawking

By the end of this chapter, you should know and be able to explain the following:

- The essentials of an intrusion detection system (IDS) and why is it necessary even if you have a firewall
- The difference between an IDS and intrusion prevention system (IPS)
- The difference between a network intrusion detection system (NIDS) and a host intrusion detection system (HIDS)
- How an IDS detects intrusions (that is, attacks) and the potential ways a response can occur
- Some of the potential IDS solutions available today

Answering these key questions will enable you to understand the characteristics and importance of intrusion detection systems in your network's overall security. By the time you finish this chapter, you should have a solid appreciation for network security, its issues, how it works, and why it is important.

The machines have taken over the world. Check into it if you want, but the truth is that they have taken over, and you simply provide the power to run them. You exist as some kind of power cell and nothing more. (Pretty weird so far, huh?)

Does this sound like some kind of nightmare, or perhaps the plot of a high-end science fiction movie? Take a moment to decide whether the guy in the trench coat and sunglasses is telling you the truth. Are you ready to cross through the looking glass to actually see what's going on? Are you ready to give up 24 hours of cable TV, media propaganda, chocolate milk, and video games? You decide whether you want the truth. *You can't handle the truth!* Or can you?

You wake up and find yourself surrounded by a glass cocoon filled with sticky viscous fluid and discover that you have probes plugged in to your spinal cord. Could this story get any worse?

It can, and it does. You start unplugging the probes one by one. Before you completely realize where you are, creepy spider machines start hovering around you and checking you out (don't you hate it when that happens?) and smacking you around.

And then...and then.... Go ahead and turn the lights back on; yes, you over there by the light switch—flip the switch. You need to stop for a moment to discuss what in the world, if anything; all this has to do with a chapter on intrusion detection systems (IDS) and honeypots.

Although this scenario is “borrowed” from a popular movie produced in 1999, this story gives you a sneak peek at the basic premise of how an IDS works. IDSs function on three basic premises:

- Where to watch
- What to watch for
- How to react

The first premise, “where to watch,” tells the IDS the logical location it will be monitoring for something to happen. The little story has you as the “where to watch” portion. The evil machine empire has instructed the creepy spider machines to monitor you and make sure that you do not wake up.

The second premise, “what to watch for,” tells the IDS conditions for which it is supposed to be looking for to raise an alarm or some other kind of action. In this case, the creepy spiders were programmed to look for you to wake up and unplug the probes. Back in the old days, all it took was someone waking up to get the creepy spiders going. Things have changed, haven't they?

The third premise, “how to react,” is the action the IDS has been told to take when a situation meets certain parameters. The creepy spiders were programmed to fly up to your pod and smack you around if you happen to wake up and start monkeying around with your sleep chamber.

Now put all the spiders and sci-fi stuff aside for a minute and take a look at a real-world example of an IDS in action:

1. You install an IDS to watch the Internet connection and those trying to get into your network through your firewall.
2. You tell the IDS what types of hacks and attacks to look for based on their packet and connection type and what activities these might generate.
3. You tell the IDS to page you and send you an email when one of these attacks occurs.
4. A malicious hacker attempts to initiate a port scan that scans the first 1000 TCP ports.

5. The IDS sees the sequential connection attempts to all these ports, checks its database, and sees that this behavior matches the profile you entered that tells it how to recognize a port scan.
6. The IDS reacts to the port scan and based on the responses you've set up, it attempts to email you and page you.
7. Suddenly, the port scans increase, and they also come from another source.
8. The IDS also notifies you of this attempt.

Now, assuming, that you have properly configured your IDS, it sits and watches your network 24 hours a day, ready to alert you at the first sign of any funny business.

Sounds pretty cool so far, doesn't it? IDSs have two major flaws:

- They are voyeuristic appliances; in other words, they just watch.
- False positives and complacency can occur.

First, the IDS can watch only one interface at a time and while it is watching that single interface, the IDS watches only for conditions you tell it to monitor for. If it has not been programmed to watch for the port-scan attack, it will not notify you when one does occur.

Finally, an IDS can actually become an ally to hackers. Impossible, you say? How many times do you still run right out of the house and check your car when you hear the factory-installed alarm go off in the middle of the night? The same "crying wolf" situation can occur with an IDS. If your pager starts filling up with messages sent by the IDS, you start filtering out what you believe to be false positives; this could lead to you missing the pages that could mean something.

The secret to successfully configuring and deploying an IDS is *tuning*. You must deploy the IDS in a lab first, see what normal traffic causes the IDS to alert, and then start "turning down the squelch"—that is, decreasing the IDS's sensitivity to these conditions. You can also resist the urge to alert on everything that occurs. Most people want to be notified of every little burp that takes place, but this is not realistic. IDSs are not perfect, and they generate false positives from time to time. Now take a real-world look at the essentials behind intrusion detection.

Essentials First: Intrusion Detection

Networks of all sizes are designed to enable the sharing of information, and only rarely is security a part of that design. Many businesses are leveraging IP-based networks, such as the Internet, to bring remote offices, mobile workers, and business partners into their trusted internal network environments. The Internet is continuously growing and connecting more and more places; as it becomes increasingly reliable, companies can redefine how corporate applications function. The clearest example is how almost everything is becoming based on HTML. Although this enables businesses to have broader interaction with customers, streamline operations, reduce costs, and increase revenues, it also comes at a price and with risks.

The reach and openness that make the Internet such a powerful business tool also makes it a tremendous liability. Simply put, the Internet was designed to connect and share, not to secure and protect. This bears repeating: The Internet was not designed to secure and protect—period—the Internet is a web connecting the world together; it is not a super highway with law enforcement.

The websites and portals that welcome remote sites, mobile users, customers, and business partners into the trusted internal network might also be welcoming attackers who would misappropriate network resources for personal gain. As discussed in Chapter 10, “Wireless Security,” the growth of wireless networks is compounding this problem.

The question becomes this: How are these mission-critical communications protected from an inherently insecure medium such as the Internet? This book covers various means of increasing the security of these resources by adding layers of protection. The most common layers of security in a network are an Internet router prescreening packet and a stateful firewall. However, your organization has both a web and email server that must be accessible from the Internet to function. You cannot block this traffic because your business depends on it. You also know that, as the Internet has grown, so have the sophistication of the attacks; however, the knowledge level required to conduct these attacks has decreased, as shown in Figure 11-1.

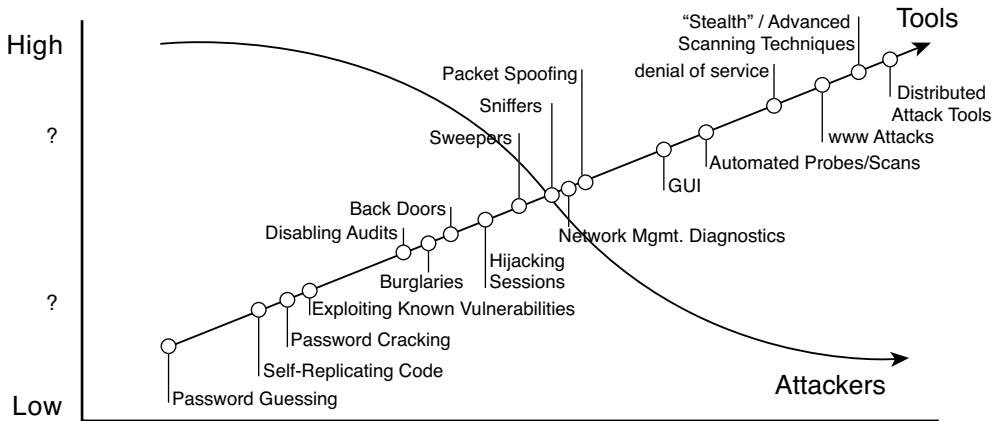


Figure 11-1 *Attack Sophistication and Attacker Skills*

Neither the router nor the firewall can tell you whether that WWW packet actually contains an attack or a customer request; unfortunately, many people have placed their trust in these devices, which can fall short in the detection arena. Perhaps your organization has a talented system administrator who is trusted to secure and lock down business-critical servers or implement thorough security policies and procedures. None of the security solutions discussed so far address the need to *detect attacks or intrusion attempts!*

In Internet terms, IDSs are rather young; research began in the 1980s with the efforts and writing of Anderson and Denning. In the 1980s, the government first began using basic

IDS functionality on what was then still the ARPANET. Late in the 1980s, members of the Haystack Project formed Haystack Labs as a commercial venture into developing host-based intrusion detection. Network-based intrusion detection followed in the 1990s with Todd Heberlein leading the charge. By then, several organizations were developing IDS tools, Haystack Labs, SAIC; in 1993 the United States Air Force implemented Automated Security Incident Measurement Systems (ASIM), and the team that developed this solution formed the Wheel Group in 1994, as shown in Figure 11-2.

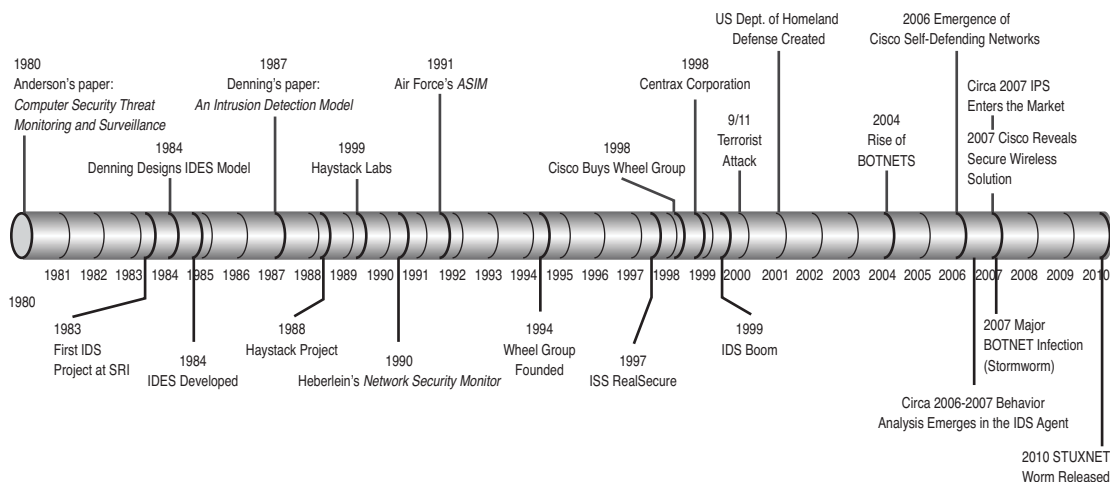


Figure 11-2 *IDS Development Timeline*

This is relevant to the discussion because in 1994, Cisco purchased the Wheel Group; this acquisition formed the core of the IDS and security services.

Note If you want to obtain a more detailed look at the history of IDS, check out the following article at <http://www.symantec.com/connect/articles/evolution-intrusion-detection-systems>.

IDS Functional Overview

Whether an attacker's motive is intellectual challenge, espionage, political, financial, or even just to make trouble; your network will face an attack. Not only is it common sense to monitor these attacks, but in many cases, it is also a business imperative. Starting in the early 1990s, new products began to appear to deal with this aspect of network security: intrusion detection systems (IDS). An IDS is like an alarm system for your network. The network is protected, but without the IDS (alarm), you would never know whether an attacker was trying to get entry. The goal of intrusion detection is to monitor network assets to detect unusual behavior, inappropriate activity, and attacks, or stop the attack intrusion and even provide information to prosecute the attacker.

IDSs that are available on the market today promise a plethora of feature sets and capabilities. In evaluating an IDS for your organization, the following capabilities should generally be the focus, beyond traditional event logging:

- **Event correlation:** When an IDS is deployed in a busy network with multiple IDSs, the ability to correlate events (attacks) is crucial to ensure that your network is secure. Consider that an attack could span multiple segments as one host is compromised and then used to attack another, and so on. Without proper event correlation, this attack could cause great confusion and lead to many hours of wasted resources attempting to isolate the cause of the outage. Event correlation enables the IDS administrator to quickly track down and relate events that occur across multiple sensors deployed in different subnets, or perhaps in different geographical locations and over extended periods of time.
- **Centralized sensor management:** Having an IDS correlate events is important, and having all the IDS managed via centralized management is just as critical. In the real world, every device (server, router, and firewall) creates logs; however, they are rarely checked, let alone reviewed. Therefore, having a centralized management platform that enables event correlation and response control over multiple sensors and the ability to run detailed reports on your network's security is crucial for success.
- **Customizable signatures and thresholds:** Company or business-specific applications, software upgrades, new operating systems, viruses, and intelligent hackers are always looking for and discovering new vulnerabilities. There is always a delay from the time a new vulnerability is discovered and when IDS developers release a new signature that detects the attack used to exploit the vulnerability. Therefore, an IDS must provide administrators with the ability to create attack signatures to deal with any eventuality.
- **Elimination of false positives:** Just like every operating system (such as Windows) that comes with all the features enabled, so do IDS devices. In other words, they are overly sensitive out of the box and provide a lot of false positives, thereby resulting in fear, uncertainty, and doubt (FUD) about your network security. You can understand, then, that every good IDS must have the capability to eliminate false positives. Having too many alarms sounding and gathering too much information can be a hazard in and of itself. An actual attack can be overlooked by getting drowned in a log-file. Take caution; however, to eliminate only a rule or feature set if you are sure of its impact. If you are sure, wait 24 hours after implementation and review. Double-check that you are capturing the data you want to capture.
- **Standards-based implementation:** An important aspect of deploying any technology is choosing a standards-based implementation. Many vendors create products that perform wonderful security services, but few are interoperable or provide the framework for future implementations. An IDS is no exception to this rule, and few standards currently exist. Because the most important aspect of integrating an IDS and managing it are its reporting capabilities; a standard has emerged based on the Common Vulnerabilities and Exposures (CVE) database. The CVE database both

classifies and groups vulnerabilities into an easily referenced system. CVE compatibility is important for IDS because it provides reporting capabilities that far surpass the typical cryptic reporting historically found in IDS. By integrating CVE-compatible IDSs, organizations can use other CVE-compatible tools, such as vulnerability assessment (VA) tools, to further enhance the accuracy and criticality of event reporting. CVE has become widely adopted and will continue to be a standard method of reporting and classifying network security events (<http://cve.mitre.org/cve/>).

- **Intrusion prevention functionality:** Intrusion prevention is essentially the ability to actively respond to and prevent intrusions and unwanted traffic. The term *intrusion prevention* has recently been the subject of much confusion and is often marketed as a competing technology to intrusion detection; however, the reverse is true. In today's market, an IDS must support the capability to actively respond to suspected threats.
- **Signature matching:** Monitors all traffic traversing a network and then matches each packet or series of packets with known attack patterns (signatures). The IDS then responds either passively or actively to that event. The response can vary from generating an SNMP alarm, crafting an email alert, or actively stopping the attacker from completing the attack (also referred to as *intrusion prevention*).
- **Anomaly detection:** Enables an IDS to establish a baseline of normal traffic patterns and information flows, and then respond whenever the normal thresholds are exceeded (for example, if a new protocol is detected on a network). Anomaly detection becomes most effective when it's coupled with protocol decoding, whereby the IDS knows what normal behavior is expected within certain protocols and responds if abnormal commands or requests are detected.

Despite a common misconception, an IDS cannot monitor everything. Having an IDS as a layer in your overall security plan is a good idea; however, depending on it as an end-all, be-all security solution is a bad idea. It is part of your network security model, one layer among what should be several layers.

You can deploy an IDS in a variety of locations within a network to further increase an organization's security and protection. There are four flavors of IDSs:

- **Host-based:** Monitors the characteristics of a single host and the events occurring within that host for suspicious activity (for example, network traffic, system logs, running processes, application activity, file access and modification, and system or application configuration changes). These systems are most often deployed on critical hosts such as publicly accessible servers and servers containing sensitive information. Use specialized software applications, called *agents*, that are installed on a computer (typically a server) to watch all inbound and outbound communication traffic to and from that server and to monitor the file system to identify unauthorized, illicit, and anomalous behavior. The installed agent uses a combination of signatures, rules, and heuristics to identify unauthorized activity. Host-based intrusion detection systems (HIDS) are extremely effective on mission-critical, Internet-accessible application servers, such as web or email servers, because they can watch the applications at

its source to protect them. Just like the network-based IDS (NIDS), the role of a HIDS is passive.

- **Network-based:** Monitors network traffic for particular network segments or devices and analyzes the network and application protocol activity to identify suspicious activity. It can identify different types of events of interest. It is most commonly deployed at a boundary between networks (that is, border firewalls, routers, VPN servers, remote access servers, and wireless networks). NIDS deal with information passing on the wire between hosts. They are typically referred to as *packet sniffers* and reside directly on the network and watch all the traffic that traverses the wire. NIDS are effective at both watching for inbound or outbound traffic flows and traffic between hosts on or between local network segments. NIDSs are typically deployed in front of and behind firewalls and VPN gateways to measure the effectiveness of those security devices and interact with them to add more depth to the network's security. NIDS use a network tap, port span, or hub to collect packets traveling over a given network. As stated before, an IDS does not actively block network traffic; its role is passive: gathering, identifying, logging, and alerting.

Note I've used the term *passive system* in describing NIDS and HIDS, but you need to understand the major difference between a passive and reactive system. In a passive system, the IDS sensor detects a potential security breach, logs the information, and signals an alert on the console and the system administrator. In a reactive system, also known as an *intrusion prevention system (IPS)*, the IPS auto-responds to the suspicious activity by resetting the connections or by reprogramming the firewall to block network traffic from the suspected malicious source.

Note Cisco NID supports 802.1Q trunking and can thus be set up to monitor multiple VLANs per single interface. Do not overburden the sensor. This means that if it is a Cisco NID, a NID can monitor more than one interface at a time.

- **Wireless:** Monitors wireless network traffic and analyzes it to identify suspicious activity involving the wireless networking protocols. It cannot identify suspicious activity in the application or higher-layer network protocols (TCP and UDP) that the wireless network traffic is transferring. It is most commonly deployed within range of an organization's wireless network to monitor it.
- **Network Behavior Analysis (NBA):** Examines network traffic to identify threats that generate unusual traffic flows, such as distributed denial-of-service (DDoS) attacks, certain forms of malware, and policy violations. These systems are most often deployed to monitor flows on an organization's internal network and are also sometimes deployed where they can monitor flows between an organization's network and the external network.

All four classifications of IDS implementation offer different techniques for detecting and deferring malicious activity, and you can deploy combinations of these four to provide the most effective enhancement to a layered defense strategy. Table 11-1 lists some examples of the various types of IDS by type, what they are commonly used for, their strengths, and some suggestions as to available products.

Tip The authors do not have a “dog in this fight.” There are many products on the market; just a few are listed. We recommend you take the time to do some research to make sure the product you decide to purchase meets all your organizations needs.

Table 11-1 *Types of Intrusion Detection Systems*

IDS Classification	Used For	Strengths	Products Available
Host-Based	Monitoring host application and operating system activity; focusing on the network, transport, and application layers of the OSI model	This is the only IDS classification that can monitor end-to-end encrypted communications.	Cisco Security Agent* SNORT HBSS Deep Security
Network-Based	Monitoring network, transport, and application layer activity	Can analyze the widest range of application protocols. The only IDS that can accurately analyze many of them.	Cisco IPS Tripwire Prelude Hybrid IDS
Wireless	Monitoring wireless protocol activity and unauthorized wireless LANs in use (that is, a rogue access point)	The only IDS that can monitor wireless protocol activity.	AirDefense/Motorola Guard Isomair Wireless Sentry
Network Behavior Analysis (NBA)	Monitoring network, transport, and application layer activity that causes anomalous network flows	Best for use when combating denial-of-service (DoS) attacks and identifying malware infections within your organization.	Cisco Traffic Anomaly Detector DefensePro

*Cisco announced December 10, 2010 as the end-of-sale and end-of-life date for the Cisco Security Agent.

These various types of IDS solutions should be deployed together to provide a truly effective layered defense with visibility into, and control of, an organization’s communications.

IDSs also provide organizations a check and balance on the effectiveness of their security systems and the overall effectiveness of their security dollars. The next section discusses the overall capabilities of an IDS.

Host Intrusion Detection System

Host Intrusion Detection Systems (HIDS) monitor, detect, and respond to user and system activity and attacks on a given host. In contrast to NIDSs, HIDSs are installed on the host (for example, the web or email server) to be monitored. HIDS monitors the host's audit and event logs, whereas a NIDS monitors packets. Rather than trying to identify packets' contents versus attack signatures, the HIDS approach attempts to identify known patterns of local or remote users doing things they should not be doing.

Note NIDSs deal with TCP/IP packets transmitted from host to host over a network, whereas HIDSs are concerned with what occurs on the hosts themselves by monitoring usage and log activity. A NIDS is like a parking lot attendant who watches all the cars coming and going out of the garage, whereas a HIDS is more like an attendant who watches the one space in which you park inside the garage.

HIDSs act much like antivirus software (however, they are not a replacement for it) with extended capabilities that greatly increase the level of security that can be provided. HIDSs are best suited to combat security threats against hosts because of their capability to monitor and respond to specific user actions and file accesses on the server. The majority of computer threats come from within organizations, from many different sources such as disgruntled employees or corporate spies. HIDSs monitor servers by providing information about the following:

- Intrusion attempts or successes and suspicious behavior by authorized users.
- Scans of the host to ensure that they conform to accepted security practices such as having all the latest patches and not having unnecessary services running.
- Audit policy management and centralization, supply of host-based forensics, statistical analysis and evidentiary support, and, in certain instances, some measure of access control. More robust tools typically provide these functions.

The deployment of HIDS is fairly straightforward; it is an application that resides on a server that watches for file system changes, registry changes, open ports, running applications, and all traffic originating to and from the host on which it resides. Server farms are often placed on their own network, and application servers are strong candidates for HIDSs.

Where multiple hosts are concerned, HIDSs should be configured to report to a centralized management console to provide event correlation and enterprisewide reporting. Typical candidates for HIDSs deployments are web servers, file servers, or any application server that provides network communication resources to the public Internet.

To get complete coverage at your site using a host-based intrusion detection system, you need to load the IDS software on every computer. Following are two primary classes of host-based intrusion detection software:

- **TCP, or host, wrappers/personal firewalls:** A TCP wrapper is an access control list (ACL) system used in host-based networking. It is used to filter network access to Internet protocol servers on Linux or BSD operating systems. It enables host, or sub-network, IP addresses, names, and indent query replies to be used as tokens to filter for access control purposes. Host wrappers or personal firewalls can be configured to look at all network packets, connection attempts, or login attempts to the monitored machine. This can include, but is not limited to, dial-in attempts or other nonnetwork-related communications ports.
- **Agent-based software:** Host-based software agents can monitor accesses and changes to critical files and changes in user privilege.

Either approach is more effective in detecting trusted-insider attacks than a network-based IDS, and both are more effective for detecting attacks from the outside. Table 11-2 lists the more popular wrapper packages and agent-based software.

Table 11-2 *HIDS Detection Software*

Class	Examples
Host Wrappers	TCPwrappers (UNIX) Nuke Nabber (Windows)
Personal Firewalls	WRQ's AtGuard (www.atguard.com)
Host-Based Agents	Cisco Security Agent* Cybersafe (www.cybersafe.com) ISS (www.iss.net) Tripwire (www.tripwiresecurity.com)

*Cisco announced December 10, 2010 as the end-of-sale and end-of-life date for the Cisco Security Agent.

Network Intrusion Detection System

Network intrusion detection systems (NIDS) sit and “capture” all the packets on the network segment to which they are connected. This reading is similar to a packet sniffer; however, the differences appear after the packets are captured or sniffed. NIDSs are built on the wiretap concept and can be implemented in a couple different ways. These methods have been developed to deal with the prevalence of LAN switches and how they operate to isolate traffic. An IDS must see as much of the network traffic as possible to be effective. The different NIDS implementation methods are as follows:

- **Inline wiretap:** This method of capturing packets places a physical *tap* in between (that is, inline between) two network devices. The NIDS would be plugged in to this tap.

- **Port mirroring:** Depending on the switch you use, port mirroring, also known as *port spanning*, is perhaps a more flexible solution. This technique tells the switch to send to another port copies of every packet that, for example, is to be sent to the port your firewall is plugged into. The NIDS connects to this mirrored port.

Some NIDSs look for a fingerprint match by comparing the packet to the attack signatures it has in its database, whereas others look for unusual packet signatures indicating an attack is in progress. The NIDS inspects the packets as they pass through a sensor. The sensor can see only the packets that happen to be carried on the network segment it's attached to. Packets are considered to be of interest if they match a certain signature. Depending on the NID you implement, the packets are analyzed against a variety of signatures. Overall, there are three primary types of signatures:

- **String signatures:** Look for a text string that indicates a possible attack. This has a tendency to produce false positives. You can refine the string signature to reduce the number of false positives by using a compound string signature. A compound string signature for a common UNIX-based web server attack might be “cgi-bin” AND “aglimpse” AND “IFS”.
- **Port signatures:** Watch for connection attempts to well-known, frequently attacked ports. Examples would be telnet (port 23), ftp (port 21/20), SUNRPC (port 111), and IMAP (port 143). If any of these ports aren't used by your site, incoming packets to these ports should be considered suspicious activity.
- **Header condition signatures:** Watch for dangerous or illogical combinations in packet headers. The most famous example is Winnuke, where a packet is destined for a NetBIOS port and the Urgent pointer or Out-of-Band pointer is set. This results in the dreaded Blue Screen of Death (BSOD) for Windows-based systems. Another well-known header signature is a TCP packet with both the SYN and FIN flags set, signifying that the requestor wants to start and stop a connection at the same time.

Some issues relate to scalability and timeliness that the IDS industry is still trying to overcome. NIDS have had some trouble scaling as network speeds have increased, and with Gigabit Ethernet making inroads to networks of all sizes, it will not be long before 10-Gigabit speeds will be used. Of course, NIDS want to capture every packet and analyze its contents; this makes these new speeds a bottleneck that has not yet been completely solved. In addition, the updating of attack signatures is not yet close to being where it should be to detect the latest attacks. It is clear that IDS vendors and how they update signatures are still a far cry from the timeliness the antivirus community has achieved.

Note Cisco has incorporated the various types of intrusion detection or prevention into many of its products through additions to its operating system and dedicated modules or devices for various components. For example, the ASA 5500 series and Catalyst 6500 series both have modules that incorporate network intrusion detection directly into them that enables increased accuracy when capturing packets and defending your network.

NIDS deployment is entirely based on the existing network design and architecture in place at each location. The more network segments a network has usually determines the number and placement of NIDSs.

Traditional NIDS placement enables them to be the most effective on the network perimeter, such as on both sides of the firewall (internal and external), near the VPN appliance server, and on links to business partner networks. This placement enables an organization to measure the real effectiveness of its prescreening routers and firewalls. These links tend to be low bandwidth (T1 speeds) such that an IDS can keep up with the traffic. This provides a good measure of checks and balances and is ideal for the security-aware organization, where application servers behind the firewall are accessible to the public Internet. Another high-value point is the corporate WAN backbone. A frequent problem is hacking from “remote” areas of the network into the main corporate network. Because WAN links tend to be low bandwidth, NIDS can be extremely beneficial.

Security best practice says that, when considering an IDS solution, both internal and external NIDSs should be used. This enables the NIDSs to monitor attacks from the Internet and internal threats. It might seem a bit odd to have two NIDSs; however, remember that statistically, the majority of attacks come from internal sources. Neglecting either location reduces the effectiveness of the IDS solution and greatly decreases your network’s security.

Wireless IDS

Like the dragon Smaug, you’re vulnerable in your underbelly. And unless you protect your weakest spot, you, too, will be felled with a carefully aimed black arrow. With the advent of wireless technologies and the saturation of wireless devices on the commercial market, it is more important than ever before to secure your wireless network and to ensure no one can break through this oh-so-big chink in your armor. A wireless IDS helps protect this spot by monitoring wireless network traffic and analyzing its wireless networking protocols to identify suspicious activity involving the protocols.

A wireless IDS is similar to a network-based IDS in many ways; they share the same components and those components have essentially the same functionality; the difference being the sensors. Wireless IDS sensors perform the same basic role as network IDS sensors, but they function differently because of the complexities of wireless communications. A major difference, however, is what the wireless IDS can monitor. A wireless IDS works by sampling traffic within two frequency ranges (2.4 GHz and 5 GHz) and each band is separated into channels. (802.11b and g support 14 channels and 802.11a supports 12 channels.) Currently a sensor cannot simultaneously monitor all traffic on a band; a sensor must monitor a single channel at a time.

Wireless sensors are available in three different ways:

- **Dedicated:** A device that performs wireless IDS functions but does not pass network traffic from source to destination. A dedicated sensor is often completely passive, functioning in a radio frequency (RF) monitoring state to “listen” to, or sniff, wireless traffic. Dedicated sensors can focus on detection and do not need to carry wireless

traffic; they offer stronger detection capabilities than wireless sensors bundled with APs or wireless switches. That being said, the dedicated sensors may be cost-prohibitive to your organization; purchasing, installing, and maintaining a dedicated switch is much more than implementing a set of bundled sensors that can be installed on existing hardware.

- **Bundled with an access point (AP):** Typically provide less rigorous detection capability than a dedicated sensor. This is because the AP must split its time between doing its job and monitoring traffic on multiple channels or bands.
- **Bundled with a wireless switch:** Typically used to help administrators manage and monitor their wireless devices. Those wireless switches that do offer an IDS function typically do not offer detection capabilities as strong as bundled APs or dedicated sensors.

Wireless IDSs have several security capabilities that traditional network-based IDSs provide, such as information gathering, logging, detection, and prevention. Because wireless IDS technology is relatively new, these capabilities vary greatly between vendors.

Network Behavior Analysis

A network behavior analysis (NBA) tool examines network traffic to identify threats that generate unusual traffic flows, such as distributed denial-of-service (DDoS) attacks, certain forms of malware, and policy violations. These systems are most often deployed to monitor flows on an organization's internal network and are also sometimes deployed where they can monitor flows between an organization's network and the external network. An NBA solution typically uses both sensors and consoles; some products on the market also offer an NBA solution with a management server, called an *analyzer*. NBA sensors are typically only available as hardware appliances. These sensors are similar to network-based IDS sensors in that they "listen," or sniff, packets to monitor network activity on a single segment. Other NBA sensors have the capability to monitor the entire network, but they rely on network flow information gathered from routers and other networking devices. Flow consists of the following components:

- Source and destination address
- Source and destination ports (TCP or UDP)
- Number of packets and bytes transmitted
- Session time stamps

NBAs have many of the same security capabilities as the other IDS classifications:

- Information gathering
- Logging
- Detection

The major difference lies in what the NBA detects. NBA technologies have the capability to detect malicious activity using, primarily, anomaly-based detection with some stateful protocol analysis. Using this unique, nonsignature-based detection array, an NBA can detect DoS attacks, scanning, worms, tunneled protocols, backdoors, and policy violations.

NBAs also provide limited intrusion prevention capabilities depending on sensor type. A passive sensor can attempt to end an existing TCP session by sending a TCP reset (RST) flag to both the source and the destination address. Inline NBA sensors offer firewall capabilities used to drop, or reject, suspicious network activity. The passive and inline NBA sensors share a couple of IPS characteristics. The first characteristic is the capability to instruct network security devices (firewalls and routers) to reconfigure themselves to block certain types of suspicious activity or route that suspicious activity elsewhere. The second such characteristic is the shared ability to run an administrator-specified script or program when certain suspicious activity is detected.

How Are Intrusions Detected?

An IDS has a special implementation of TCP/IP that enables it to gather the packets and then reassemble them for analysis. It is not enough to simply sniff the packets; an IDS must examine them. An IDS can use one of three methods to detect intrusion:

- Pattern matching or signature-based
- Statistical anomaly-based
- Stateful protocol analysis

A pattern matching or signature-based model uses a set of rules, or signature, to detect an attack in progress. A device used for intrusion detection is loaded with a set of signatures. Each signature contains information about the kind of activity to look for in traffic passing through the network to detect whether an attack is under way. When the traffic passing through matches the pattern contained in a signature, an alarm is generated, notifying the network administrator of the intrusion.

Statistical anomaly-based IDS rely on establishing thresholds for various types of activity on the network. This type of activity can lead to detecting an intrusion taking place. The drawback to this method is it is difficult to rely on solely; as the sophistication of the means of detecting an intrusion got better, so did the people attacking your organizations. Many attacks do not lend themselves to easily being detected based on threshold limits, which might require more manual intervention in the tuning process.

Finally, there is the stateful protocol analysis, or deep packet inspection, process. This is a process of comparing predetermined profiles of industry-accepted definitions of non-harmful protocol activity for each protocol (FTP, SSH, TFTP, and so on) state against observed behavior to located anomalies coming through your IDS. The thought process is that when a user begins a TFTP session, IDS should expect to see certain commands, and if it sees anything other than those basic commands (view, list, and so forth) it flags the activity as malicious.

Every IDS vendor (of which there are several) has buzzwords of every type to confuse the buyer on the explanation of how an IDS performs its job. This seems counter-productive because each vendor wants to sell, but alas, the world is a fickle place and security is full of snake oil! This section takes a high-level look at the methods any good IDS should use.

Signature or Pattern Detection

Signature or pattern-based detection compares known threat signatures to observed events to identify incidents. Signature/pattern matching is the most common method of detecting attacks, and it means that the IDS must recognize every attack technique to be effective. An IDS has large databases with thousands of signatures that enable the IDS to match attack signatures or patterns. This particular methodology is effective at detecting known threats, but largely ineffective at detecting unknown threats and many variants of known threats. Signature-based detection cannot track and understand the state of complex communications; therefore, it cannot detect most attacks that compromise multiple events.

For example, many IDSs are used to monitor abuse, such as a user visiting pornography or gambling websites while at work. Detecting that kind of misuse is based on a keyword; however, consider a different scenario in which someone uses ICMP to scan and map out your network.

This type of attack detection takes place at a more granular level than protocol analysis or anomaly detection. As a result, specific events are identified that, for example, indicate that a compromise has occurred. One of the most frequently matched patterns is when an attacker ensures that he has achieved root permissions on a host. The host replies that root access was achieved in a packet that will be sent to the attacker and that can be analyzed for the word *root*. This is a greatly simplified example, but it demonstrates what an IDS looks for (that is, matches).

Anomaly-Based Detection

Anomaly-based detection is used to compare definitions of what activity is considered normal against observed events to identify significant deviations. This method uses profiles developed by monitoring the characteristics of typical activity over a period of time. The IDS then compares the characteristics of current activity to thresholds related to the profile. Anomaly-based detection methods can be effective at detecting previously unknown threats. Common problems with anomaly-based detection are inadvertently including malicious activity within a profile, establishing profiles that are not sufficiently complex to reflect real-world computing activity, and generating many false positives that take a fair amount of time and effort to tune out.

Anomaly detection is similar to the training of spam filters because a period of learning by IDS allows it to determine normal baseline levels of activity. Of course, normal is different for every network. The thought behind this approach is to measure a baseline of statistics such as file activity, user logins, CPU utilization, disk activity, and so on. After the baseline is established, IDS is used to detect statistical anomalies.

For example, assume that you are monitoring activity and your IDS begins to note that, early every morning, many of the hosts on your network become active. You might not immediately know what is going on, but you are alerted that you should investigate.

Stateful Protocol Analysis

Stateful protocol analysis compares predetermined profiles of generally accepted definitions of benign protocol activity for each protocol state against observed events to identify deviations. Unlike anomaly-based detection, stateful protocol analysis relies on vendor-developed universal profiles that specify how particular protocols should and should not be used. It can understand and track the state of protocols that have a notion of state, which enables it to detect many attacks that other methods cannot.

Attacks use methods of altering the underlying protocol information to be successful. For example, the Ping of Death is successful because it alters the packet size and, through protocol verification, this would be detected. An IDS has a verification system that can flag invalid packets; this can include valid packets that are severely fragmented, which again proves that communication stream reassembly is important.

An important aspect of protocol verification is that of application verification, where the IDS detects inappropriate application protocol behavior. For example, the WinNuke attack uses NetBIOS (a valid protocol) but adds out-of-band information, which is valid but is used only to attack a host.

Problems with stateful protocol analysis include the following:

- It is often difficult or impossible to develop completely accurate models of protocols.
- It is resource intensive.
- It cannot detect attacks that do not violate the characteristics of generally acceptable protocol behavior.

Combining Methods

Attackers continually modify and improve their abilities, thereby making them increasingly difficult to detect. To combat this, IDS continues to evolve, becoming smarter and better at detection by combining the methods they use to detect intrusions.

For example, an IDS might have the capability to combine the methods of signature-based pattern matching, protocol analysis, and anomaly detection. This capability to use multi-method attack detection is another example of the ever-evolving way in which IDSs continue to grow.

Intrusion Prevention

An IPS picks up where an IDS leaves off by providing the capability to prevent an attack from being successful at the earliest possible moment by blocking, or rejecting, packets that match a particular signature or behavior. To make this effective, the IPS sits inline as

opposed to using a network tap or port span. The most effective IPS, ideally, works with an IDS; many vendors have seen this need and combined the two technologies to make an IPS-capable IDS.

There are many types of IPS technologies, which are differentiated primarily by the types of events that they can recognize and the methodologies that they use to identify incidents. In addition to monitoring and analyzing events to identify undesirable activity, all types of IPS technologies typically perform the following functions:

- **Stop the attack:** The IPS can stop an attack by one of three means:
 - Terminate the network connection or user session being used for the attack.
 - Block access to the target from the offending user account source address (IP address).
 - Block all access to the targeted host, service application, or other resource.
- **Change the environment:** The IPS could change the configuration of other security controls to disrupt an attack. Common examples are reconfiguring a network device to block access from the attacker or to the target, and altering a host-based firewall on a target to block incoming attacks.
- **Modify the attack's content:** Other IPS technologies can remove or replace malicious portions of an attack to render it benign. A simple example would be an IPS removing an infected file attachment from an email and then permitting the cleaned email to reach its recipient.

IDS Products

Many IDS/IPS systems exist, and a lot of confusion surrounds them because there is little in the way of standards for how they operate. It is difficult to provide a direct comparison between products because terminology, meanings, features, and functionality have not matured to a level at which an effective comparison can occur. However, many products are based on the work done by the open source community efforts in the IDS arena. The foremost of these products is Snort (www.snort.org).

Snort!

Snort! Snort! Snort!

Don't worry—the snorting that you are hearing is not coming from some sort of weird beast; it is coming from an open source IDS developed by Sourcefire. Following is a description of Snort, quoted from the website (www.snort.org):

...Snort is an open source network intrusion prevention and detection system (IDS/IPS) developed by Sourcefire. Combining the benefits of signature, protocol, and anomaly-based inspection, Snort is the most widely deployed IDS/IPS technology worldwide. With millions of downloads and approximately 300,000 registered users, Snort has become the de facto standard for IPS...

...Snort is capable of performing real-time traffic analysis and packet logging on IP networks. It can perform protocol analysis, content searching/matching, and can be used to detect a variety of attacks and probes, such as buffer overflows, stealth port scans, CGI attacks, SMB probes, OS fingerprinting attempts, and much more....

Snort uses a flexible rules language to describe traffic that it should collect or pass, and a detection engine that uses a modular plug-in architecture. Snort has a real-time alerting capability as well, incorporating alerting mechanisms for syslog, a user-specified file, a UNIX socket, or WinPopup messages to Windows clients using Samba's smbclient.

Snort has three primary uses. You can use it as a straight packet sniffer such as tcpdump(1), a packet logger (useful for network traffic debugging, and so on), or as a full-blown network IDS.

The Snort application is an extremely well-written command-line application, albeit sometimes difficult to configure and monitor, unless you've cut your IT teeth on the command-line interface. To make it more click-friendly, several companies have made GUI interfaces to work hand-in-hand with Snort. The following figures show screen captures of a front-end third-party GUI application called IDScener (version 1.1 RC4 is the most recent version); it is a GUI interface developed by a group of people at Engage Security (www.engagesecurity.com).

As shown in Figure 11-3, the IDScener GUI front end enables users to have a graphical configuration and monitoring interface.

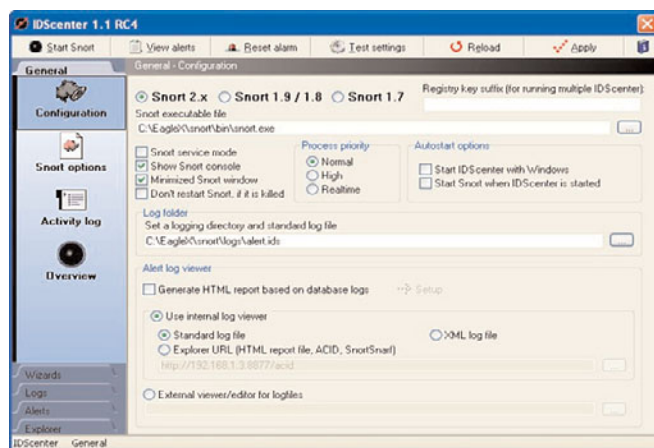


Figure 11-3 IDScener Main Snort Configuration Screen

Figure 11-4 gives you an idea of the basic configuration options that can be set for a Snort operation. This screen in IDScener is crucial for beginning the inevitable tweaking that must occur, and Figure 11-5 shows the method users have for selecting the intrusion or attack profiles that Snort will be required to look for and how to notify customers. In

addition, Figure 11-5 shows the alerting options that can be set when situations that require administrative attention occur.

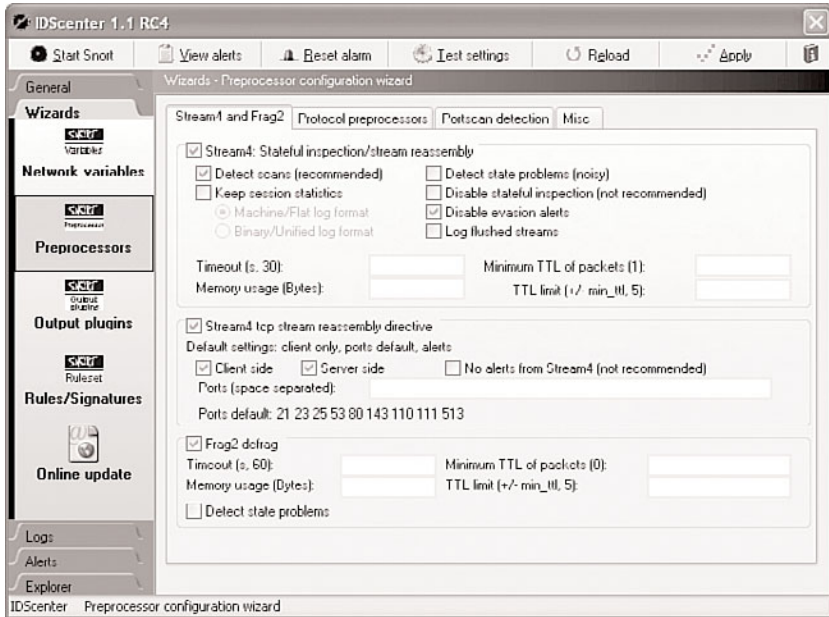


Figure 11-4 IDScenter Snort Rule Configuration Page

You might already know about Snort if you are familiar with Linux or other *nix operating systems, but what you might not know is that the IDScenter software is Win32-based, as is the version of Snort that it manages.

Created by people who truly want Snort to work under Windows, a handful of Win32 Snort installation packages are available on the Internet. The majority of these packages work well, but a few of them need additional development cycles. Also keep in mind that after you configure and deploy a Snort machine, you cannot use it for anything else after you engage the monitoring functions.

Limitations of IDS

Still an evolving technology, IDS has some manageable limitations given its overriding benefits. An IDS should always be deployed in addition to prescreening routers and firewalls. Primary systems, such as firewalls, encryption, and authentication, are rock solid. Bugs or misconfiguration often lead to problems in these devices, but the underlying concepts are proven and accurate. Regardless, some of the limitations are as follows:

- **HIDS versus NIDS debate:** This should never be a debate; both are needed and should work together in a unified approach to increase your network's security because they play different roles. Even though both tools are needed, each has its draw-

backs and limitations. Network-based IDSs have two major limitations: switched networks and false alarms. In a switched network, network-based IDS must be able to view and analyze all network traffic of the network it is protecting to be effective. Because most networks use a switch of some form or fashion, a sniffer cannot see all the network traffic—only the traffic on the segment to which it is attached. This means you'd have to deploy a NIDS at the perimeter, similar to a honeypot (see Figure 11-6 in the next section); this protects you somewhat from external attacks but does nothing to protect against an internal attack. Furthermore, depending on how your NIDS is configured, you might end up with a high rate of false positives. Modern-day enterprise networks amplify this disadvantage because of the massive amounts of data that need to be analyzed. Host-based IDSs have a couple of disadvantages: complexity and failure due to compromise. First, the implementation of HIDS can get complex in large networking environments. With several thousand possible endpoints in a large network, collecting and auditing the cumbersome log files generated by each node can be overwhelming. Second, if the HIDS system is compromised, the host could cease to function, resulting in a stop on all logging activity. Furthermore, if the IDS system is compromised and the logging still continues to function, the trust of such log data is severely diminished.

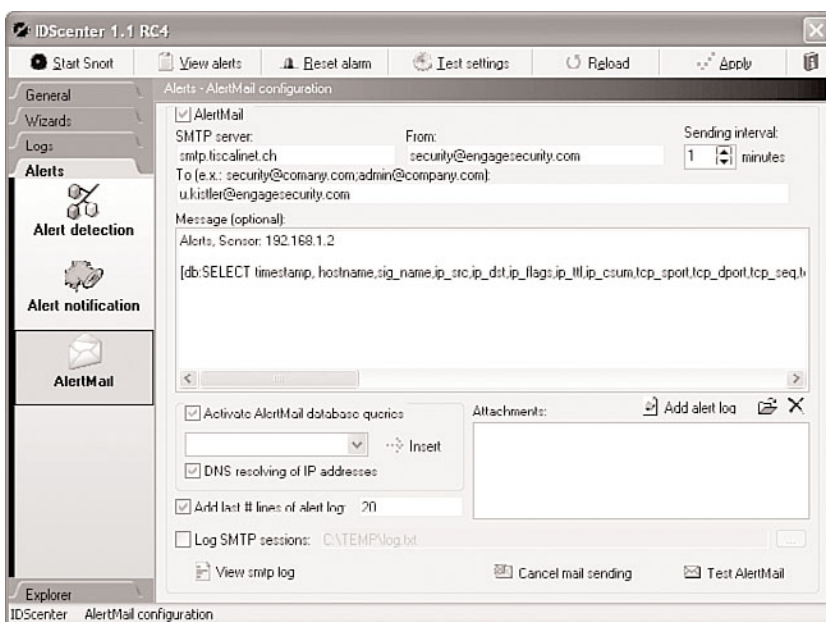


Figure 11-5 *IDScenter Notification*

- **Attack patterns and signatures:** Similar to antivirus software that uses virus definition files (VDF) to remain current, IDS products use attack pattern signatures. These signatures range from the simple—checking the value of a header field—to the highly complex signatures that actually track the state of a connection or perform extensive protocol analysis. Signature capabilities vary greatly among IDS products and are not

always updated with the latest attack signatures. Some network IDS products provide little ability to customize existing signatures or write your own, whereas other IDS products give you the ability to customize all their signatures and write almost any signature you can think of. Another important factor to consider is that some IDS products can check only certain header or payload values, whereas other products can give you the data from any portion of any packet.

- **False positives and false negatives:** A *false positive* is any normal or expected behavior identified as anomalous or malicious. The major problem that false positives create is that they can easily drown out legitimate IDS alerts. A *false negative* is any alert that should have happened but didn't. False negatives produce two problems:
 - There are missed attacks that will not be mitigated.
 - False negatives give a false sense of security.
- **Resource limitations:** NIDS sit at centralized locations on the network. They must keep up with, analyze, and store information generated by potentially thousands of machines. NIDS must emulate the combined entity of all the machines sending traffic through its segment. This means if you have a large network segmented into 30 VLANs and you want to monitor traffic on all segments, you need a NID for each segment. You can purchase one NID and place it on the edge of your network, but then you capture traffic only coming in from potential external threats, and according to the NSA 93 percent of all attacks come from within your own network boundaries.

Note When buying an IDS, ask the vendor how many packets per second the system can handle. Many vendors try to tell you how many bits per second, but per-packet is the real performance bottleneck. It is now increasingly common to have IDS signature-matching techniques in various IDS/IDP products with throughput up to 10 Gbps for Ethernet speeds on your enterprise networks and up to 40 Gbps speeds at your core. Make sure the IDS/IPS you implement does not become a bottleneck to productivity.

- **Long-term state:** A classic problem is slow scans, in which the attacker slowly scans the system. The IDS cannot store that much information over that long of a period, so it cannot match the data together.
- **Sensor blindness:** IDSs are built on regular computers that do not have any special capabilities; thus, it is possible to saturate the link to which they connect and blind them, thereby causing them to drop packets they should have been recording. For example, the open source port-scanning tool nmap includes a feature known as *decoy scans*, which causes nmap to send hundreds of scans using spoofed IP addresses. It therefore becomes an improbable task for the administrator to discover which of the IP addresses are real and which are decoy addresses. These two scenarios retain forensics data, however. If the attacker is suspected, the data is still there to find. Another attack is to fill up event storage.

- **Storage limitations:** When attackers try to blind the IDS sensor, they might have the dual purpose of filling up the sensor database or hard drive. This causes the sensor to delete events or stop recording events.
- **Denial of service:** An IDS is extremely complicated because it has an entire TCP/IP implementation running. As a result, IDS can be susceptible to attacks. Attackers can often download the same IDS their targets use free of charge, and then experiment to find packets that will disable the IDS. During the attack, the intruder then disables the IDS and continues the attack undetected.
- **Fragmentation:** The act of breaking up large packets into multiple smaller packets. The receiving TCP/IP stack then reassembles the packets and their data. Most IDSs do not have the capability to reassemble IP packets. Simple tools exist that can automatically fragment attacks to evade IDS.

Note Fragmenting the IP packets in the middle of the TCP header has long been used to evade firewall port filtering. Some industrial-grade NIDSs can reassemble traffic. Also, some firewalls can normalize traffic by forcing reassembly before passing the traffic through to the other end. IDS sensors have monitor ports (the sniffing port) that have no IP address assigned to them and are therefore not susceptible to DoS attacks. A management interface that has an IP address assigned to it should be placed in a VLAN and isolated from normal network access.

- **Pattern evasion/change:** Many simple NIDSs rely on *pattern matching*. Attack scripts have well-known patterns, so simply compiling a database of the known attack scripts provides pretty good detection; however, it can easily be evaded by simply changing the attack script and thus rendering the IDS pattern match useless.
- **IDS evaluation tools:** Many tools are freely available to test the accuracy and usefulness of an IDS. The two most commonly used are snot and Stick. These tools create thousands of attacks to see whether the IDS can sense them. Attackers can use these tools to hide their attacks or to potentially blind the IDS.

These limitations do not mean that the uses of IDS are invalid or somehow lessened. Hacking is so pervasive and attack tools so readily available that it is astounding what an IDS can detect. Properly maintained and managed, IDS dramatically improves the security of any network. However, one of the key fundamental points of properly using an IDS has been saved for the end of this section.

A security policy is crucial to the successful use of an IDS. You might ask, “But why do I need another policy and process to follow? They are such a burden!” It cannot be emphasized enough that the assets your security measures are designed to protect have value, and not ensuring that *everyone* involved in protecting them follows the same standards would be a grave mistake. One cowboy can ruin it for everyone.

Essentials First: Honey pots

“...When having a smackerel of something with a friend, don't eat so much that you get stuck in the doorway trying to get out....”—Winnie the Pooh, *Pooh's Little Instruction Book*

You are probably wondering what Winnie the Pooh and his predilection with honey are doing in a book about network security. It's less about Pooh and more about the analogy. Pooh was always getting himself in trouble because he would always be attracted to the honey and then eat too much. If my memory serves, the preceding quote is from Pooh after he snuck into Rabbit's house and found all the honey. He then helped himself and upon trying to leave found himself stuck in the hole (yes, I have three children). That's essentially the purpose of a honeypot: It brings in the hacker/black-hat types and traps them there so that you can log their activities.

This portion of the chapter covers honeypots to demonstrate that, just as an active device such as an IDS has a role in securing your network, so does a passive device such as a honeypot, which does not have the same limitations as an IDS.

Honey pot Overview

Until this point, this book has not discussed taking the fight to the attackers, so let's refocus. First, a definition: Honey pots are highly flexible computer system security tools with different customizable applications used to expressly lure and “trap” people who attempt to penetrate your organization's computer systems through probes, scans, and intrusions. This target audience includes the hacker, cracker, and script kiddie, regardless of their location in the world. When I first heard of the honeypot concept, I was confused. Why in the world would you want such a device on your network? It seems to me that having a computer designed to let attackers hack into would not serve much of a purpose. As it turns out, if correctly implemented and closely monitored, these network decoys serve a threefold purpose:

- Honey pots distract attackers from more valuable resources on your network, thus allowing the protection of your resources by distracting attackers to devices that they presume are real.
- Honey pots provide early warning about new attacks and intrusion attempts. IDS can generate false positives, whereas those who are likely to intend harm access only a honeypot because it is nonproductive.
- Honey pots allow for an in-depth examination of an attacker's activities during and after the exploitation of the honeypot. This might seem like something only someone involved in research might do, but think about what you can learn. You can use this education to ensure that the real security resources on your network are correctly configured or patched.

Note Lance Spitzner, an expert on honeypot systems, documents them in a series of articles titled “Know Your Enemy” as a part of the Honeypot Project (www.honeypot.net). He describes how to track attackers through the system to gain sufficient information about how they operate in these articles.

Clearly, the problem of false positives discussed in the IDS section is not a real issue with honeypots. Specifically, if an attack happens to a honeypot, a passive but monitored device, you will know it. This actually means that detection of attacks is no longer much of an issue, is it? In the real world, you often see honeypots deployed on a demilitarized zone (DMZ); however, the honeypot is not listed in DNS, WINS, or registered, nor is it linked to a production machine in any way. If the honeypot begins to get scanned from hosts within the DMZ, that tells you something. What if the honeypot is inside the network and it gets attacked? These placements of honeypots are passive in that they are waiting for someone to attack them.

The design and intent of honeypots fall into two categories:

- **Production honeypots:** Used by organizations concerned with the security of their networks; we focus on these. A production honeypot is typically deployed with a certain goal or intent in mind. They are easy to use, capture only limited information, and are used primarily by companies or corporations.
- **Research honeypots:** Just the opposite. They are complex to deploy and maintain, capture extensive information, and are primarily used by research firms, military, and government organizations.

Note There is currently some question as to the legality of honeypots and whether they fall under the banner of wire-tapping devices. As silly as that sounds, the FBI and other law-enforcement agencies are still battling over this question.

In addition, honeypots can be classified by their function, as follows:

- **Port monitors:** A rather straightforward type of device, these honeypots listen on ports targeted by attackers. By design, they respond to port scans, thus letting the attacker attempt to connect. These types of honeypots log connection attempts on a port.
- **Deception systems:** Take the next step from just monitoring a port and deceive attackers by interacting with them as a real system would. This means that, instead of just replying on TCP port 110 such as an email server configured for POP3, a deceptive honeypot responds as if it were a real mail server. Deception systems do not implement every aspect of a mail server; rather, they implement just enough to make it sweet as honey to an attacker.

- Multideception systems:** Increasing yet another level are the more advanced honeypots that not only enable multiple services that can be emulated, but can also simulate different operating systems. One of the most commonly used tools for this purpose is Specter, which you can find at www.specter.com/.

Note You can explore additional aspects of honeypots where there are entire systems dedicated as honeypots. Then, the detection is taken a step further through the use of an IDS when honeypots are in use. You can find one of the best resources for honeypots at www.honeypots.net/honeypots/links. You can also download a freeware honeypot for Win32 machines called, of all things, “Honey Potter” from <http://honeypott4.tripod.com/>. It is a basic piece of software (it is free, after all) that provides an introduction to honeypots.

Honeypot Design Strategies

Perhaps the clearest and most present danger is that when your honeypot works correctly, it detects attackers coming after your network and its resources. In practice, this means that you already have a criminal in some part of your network. As a result, you must take care of a few items to ensure the security of the network.

Use a firewall! Yes, a firewall—even though the honeypot is designed to let attackers in, still use a firewall to ensure that they do not get too suspicious. Create a rule set that allows basic Internet functionality out from the honeypot back to the Internet. Experts recommend that you should allow all inbound traffic to reach the honeypot, but allow only FTP (ports 20/21), ICMP, and DNS (port 53) outbound.

Figure 11-6 includes basic/simplistic honeypot architecture showing you potential locations for honeypots within your organization’s network.

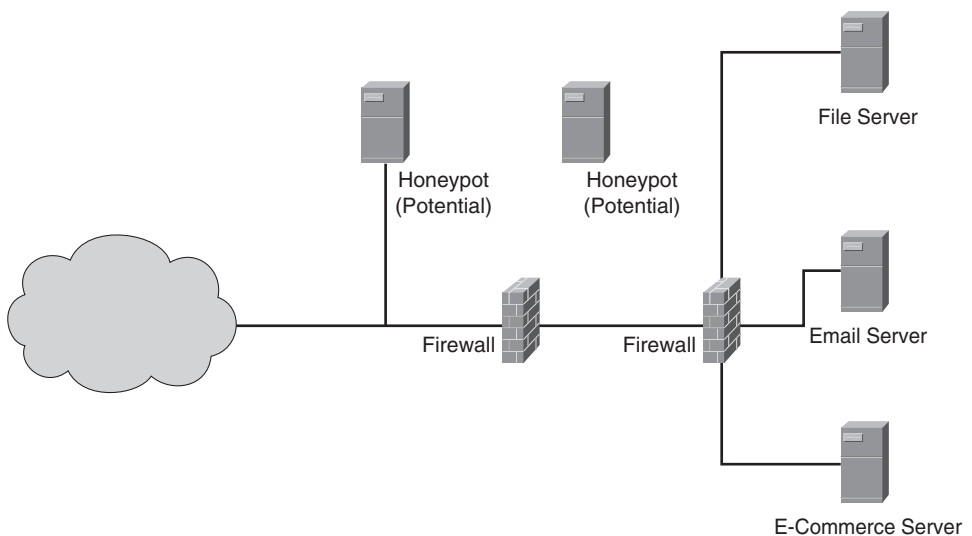


Figure 11-6 *Simplistic Honeypot Architecture*

The way you can see an attacker's activities is through various logs and through the actual honeypot logs. Failure to ensure that these are working will make your life difficult and basically nullify your entire motivation for setting up a honeypot.

Note Some people feel that capturing criminals in this manner is something that should be considered a form of entrapment. This is a misconception because honeypots are not active lures—they do not advertise themselves. A honeypot is not stumbled into by any legitimate user, and a good user would never “root kit” you.

Honeypot Limitations

Even with all their benefits, honeypots do not fix a single security problem. Instead, they are used for misdirection, prevention, detection, and information gathering by being closely monitored and designed to look like something they are not for the attackers to hack into. Conceptually, this means that a honeypot should not be used for production because its value lies in being probed, attacked, or compromised. A honeypot has many benefits; it also has the following limitations:

- **Open-Door:** If the system does indeed get hacked, it can be used as a stepping-stone to further compromise the network.
- **Complexity:** Honeypots add complexity. In security, complexity is bad because it leads to increased exposure to exploits.
- **Maintenance:** Honeypots must be maintained, just like any other networking equipment/services.

Chapter Summary

This chapter introduced two of the newest available security-related technologies: intrusion detection and intrusion prevention. This chapter began by exploring the two fundamental types of IDS: host-based that run on servers and network-based IDS that run on a network. This chapter also covered the basic operation of an IDS and concluded by covering honeypots.

Chapter Review Questions

1. When was the first commercial IDS developed and by whom?
2. What are the two types of IDS and should they be deployed together or separately?
3. Define and discuss NIDSs. How and where are they effective in a network?
4. Define and discuss HIDSs. How and where are they effective in a network?
5. When is anomaly detection the most effective and why?

- 6.** Which intrusion detection methodology also verifies application behavior?
- 7.** List and define each of the two techniques an IDS can employ to prevent an attack.
- 8.** List the three most important IDS limitations, in your opinion, and explain why you choose them.
- 9.** True or false: Honeypots distract attackers from more valuable resources.

Tools of the Trade

“The happy people are those who are producing something; the bored people are those who are consuming much and producing nothing.”—William Ralph Inge

By the end of this chapter, you should know and be able to explain the following:

- The fundamental types of attacks that your network might experience
- How to conduct or contract a security assessment of your network’s security
- How to use the results from a security scan and vulnerability assessment to better secure your network
- How to conduct or contract a penetration test of your network’s security

Answering these key questions will enable you to understand the overall characteristics and importance of network security. By the time you finish this book, you should have a solid appreciation for network security, its issues, how it works, and why it is important.

The HaXor that stole Christmas...what a great way to start this tools-of-the-trade chapter.

Every holiday season all over the world, people experience an OOBIE, otherwise known as an out-of-box experience. OOBIE is an acronym that refers to the excitement and wonderment that many people enjoy when they open the box their new computer comes in. The smell of the new plastic, the tactile sensation of the new keyboard, the sound that a new computer makes when you boot it up for the first time—all the sights and sounds that come with getting your brand new, shiny SuperComp 2000 laptop with 4 gigabytes of super-duper speedy RAM, 500-gigabyte Serial ATA hard drive, otherwise known as a MacBook Pro—or so I have heard.

As you might imagine, this scene plays out in many households throughout the world every time the Christmas season rolls around. The previous year’s PC is relegated to being the de facto “family” computer—the one that never gets its hard drive defragged or patched and consequently takes three days to boot up.

This year's personal computer is going to little Johnnie (or Joanne, to be politically correct) because he is a freshman in high school now and is required to turn in top-notch reports for biology and chemistry and whatever other classes require report writing on a computing platform 150 times more powerful than the computers on the Space Shuttle and NORAD combined.

Besides, little Johnnie/Joanne needs something pretty powerful for playing all those online games available via the brand spanking new broadband connection you got last month when you were planning ahead for the big box under the tree! And, of course, the computer will be up in his/her room to make doing homework less of a chore and more of an individual accomplishment, achieved in the combination bedroom and office. Perhaps Johnnie/Joanne might even have received a laptop; that broadband modem has wireless, so you can use it when you need to work, too.

Take a moment to get a few things crystal clear. First, teenage children do not need a computer capable of breaking encryption in less than two days; several government agencies are capable of doing just that, and they actually do not like the additional competition. Second, the words "You've got mail" will not be heard through the speakers of the super computer in question when it connects to an unsupervised broadband connection that has huge download speeds rivaling a DS3. You can expect to hear the sounds of heavy metal, rap, and whatever other kinds of music they can download via MP3s.

Regardless of whatever story your little high-school sophomore tells you, if they are in any kind of computer science course at school (as more than 80 percent of them are), they are striving for one goal: to be crowned "Uber Haxor" (pronounced oober hacksor) by their little felonious classmates. That's right; your little baby that used to eat peas and carrots with their toes is but a few mouse-clicks away from being brought up on charges under the U.S.A. Patriot Act, and the shiny new PC you bought for Christmas is the high-tech hotrod that might end up getting them an extended stay at the "gray bar motel." How many of the attacks, techniques, and tools discussed in this book cost money? Not many; most are free and those that do have cracks are available on the Internet.

The combination of intelligence and a burgeoning contempt for authority in any form (teenager) can make a state-of-the-art computing device a dangerous thing if it ends up in the wrong hands. Now, you might be saying to yourself, "My children would never do anything like that. I've brought them up to respect authority and have taught them the difference between right and wrong." All this might be 100 percent correct, but in educating the little tykes, you might have forgotten that the Internet is still as wild and wooly as the west was in the 1800s. Sometimes, it is not the victim's fault—malicious software and hacks are easy to miss.

Surfing the Internet is a common occurrence for children who have grown up in the past 15 years, and the amorality of the Internet lends itself to bad decisions. At last count (and some people have actually counted), thousands of websites are dedicated to hacking, cracking, and computer crime. Finding information on how to write viruses is easier and more fun than locating a recipe for double fudge brownies (but not as tasty).

Broadband Internet access has created a culture of anonymity that has never existed before for children seeking ways to rebel and embarrass their parents for grounding them

or taking away the car keys. Email, YouTube, MySpace, Facebook, and video chat rooms have empowered children to explore the boundaries of society in an instant and exploit the weakness of that society on a whim when they determine that society has treated them badly.

Even at this point, you might still be convinced of your child's enduring innocence and good intentions when it comes to behaving responsibly with regard to Internet usage. You might be correct in maintaining your belief; but then again, if you had asked the attacker's mother who was recently in the news if her son was capable of these kinds of acts, she probably would have denied that her son was capable of executing the attacks, and you know how that story turned out.

This chapter discusses the security tools that attackers use so that readers can understand what they are up against. The chapter then examines the tools available to identify weaknesses in your network and the anatomy of a security audit, which is a crucial piece to ensure that your network is secure.

Essentials First: Vulnerability Analysis

This section looks at some of the tools that are freely available to attackers. The fundamental truth this section teaches is that *the bad guys have good tools*. Previous chapters touched on many of the specific attack tools; however, attackers have a broad toolset with which they can launch multilevel attacks against your network. When an attacker gains a foothold in your network, it is leveraged to exploit another aspect of your network.

Attackers, for example, can, do, and will take advantage of weak authentication and authorization, improper allocations, poor security implementation, shared privileges among users or applications, or even poor employee security habits to gain unauthorized access to critical network resources.

Throughout this book, you have seen many ways to allow even the best security procedures and technologies to be circumvented. To understand how this is done, you must spend some time understanding the exact methods and tools the attackers use. Perhaps it is exciting to see an attacker's tools; however, if *any* network resource is your responsibility, you must ensure that you use these tools to assess your network's security. It is extremely important that you use these tools on behalf of your network so that vulnerabilities can be detected and found before an attacker uses them against you. You can count on these tools being used in your network—the decision you must make is *who is going to use them first?*

Fundamental Attacks

Leading-edge security technologies, policies, and procedures can quickly have their effectiveness nullified if those who are responsible for network security do not understand the methodology and tools that will be used against your network. This chapter discusses some of the methodologies and various tools that attackers use, how they operate, and

the tools and techniques you can use to protect your network resources against these hacker tools.

Even the best security technologies and procedures can be rapidly nullified unless you know the precise methods and tools being employed against you. Therefore, it is crucial to be able to identify the various tools of the hacker trade, how they operate, and what kinds of protections thwart these attacks. This includes a thorough knowledge of the common tools and techniques discussed in the sections that follow.

IP Spoofing/Session Hijacking

This type of attack occurs when an attacker creates a packet with a different IP address to gain entry to a system. This attack exploits trust relationships by allowing the attacker to assume a trusted host's identity. The header for each IP packet contains the source and destination IP address of the packet. By forging the header so that it contains a different IP address, an attacker can make it appear that the packet was sent from a different machine. The machine that responds will respond back to the forged source IP address. For this attack to be successful, the attacker must determine the "patterns of trust" for the target host—that is, for example, the range of IP addresses that the host trusts. After the attackers determine the pattern of trust, they can move on to the next step of the attack by either compromising the host or disabling it in some manner. These types of attacks are often used as the first step in the overall attack strategy. IP address spoofing is most frequently used in denial-of-service (DoS) attacks—the types of DoS attacks are covered in more detail in a few pages.

Note Because attackers spoofed an IP address (that is, made it up so the target trusts it—the address could be a local LAN address, whereas the attackers are not local), the attackers might not see the response from the target. This means that the attackers are blind to their success, which is why this is usually a first step. It is common to blindly exploit vulnerabilities in this matter and, after the host is compromised, move to the next step.

IP Spoofing/Session Hijacking Tools

A variety of tools accomplish exploitation through IP spoofing/session hijacking. A quick Google search on "IP Spoofing Tools" returned more than 139,000 hits. The list that follows describes just a few:

- **Dsniff:** A collection of tools for network auditing and penetration testing specifically known as Dsniff, filesnarf, mailsnarf, msgsnarf, urlsnarf, and webspay passively monitor a network for interesting data (passwords, email, files, and so on). arp spoof, dns spoof, and macof facilitate the interception of network traffic.
- **Hunt:** A program for intruding into a connection, watching it, and resetting it. Hunt was an outgrowth of similar products such as Juggernaut, but it has several features that cannot be found in these products.

- **Ettercap:** A powerful Apple OS X, Windows, and UNIX-based program employing a text-mode GUI, easy enough to be used by script kiddies. All operations are automated, and the target computers are chosen from a scrollable list of hosts detected on the LAN. Ettercap can perform four methods of sniffing: IP, MAC, ARP, and public ARP. It also automates a variety of other tools.
- **TotalSpooF:** A free and useful utility that enables you to spoof websites. This spoofing enables the hacker to fool a website that has security into thinking that the hackers are already in the site so that the site logic will not check the authentication credentials again and enable entry into the secure site.

Prevention

Preventing these kinds of attacks is as important as understanding them and the tools that are used. Virtual private networks (VPN) are effective against IP spoofing because a VPN encrypts the original IP addresses as they are transmitted across the network. If either the data or the source address proves to be tampered with, the packet is deleted. This prevents an attacker from penetrating a system without access to the VPN encryption keys.

Packet Analyzers

A packet analyzer (also known as a *network analyzer* or *protocol sniffer*) is a tool that intercepts and logs traffic passing over a digital network or part of a network. As data streams flow across the network, the sniffer captures each packet and decodes and analyzes its content according to the appropriate RFC or other specifications. Sniffers generally come either software-based (for PCs/PDAs) or hardware-based (on a dedicated computer). Using sophisticated network sniffers that can decode data from packets across all layers of the OSI model, attackers can steal usernames and passwords and use that information to launch further attacks. In general, attackers can use sniffers by compromising the corporation's physical security—say, walking into the office and plugging a laptop into the network. Many sniffers have built in “expert systems” that determine critical network data without the user having to need any skills beyond clicking a button.

With the growing use of wireless networks, someone in the parking lot with a wireless device can access the network. By using sniffers, attackers can obtain valuable information about usernames and passwords across public or private networks—in particular, from applications such as FTP, Telnet, and others that send passwords in the clear. Protocols such as SMTP, IMAP, and POP3 are used for remote access to email applications via simple username and password authentication techniques and are especially susceptible to sniffer attacks. Because users tend to reuse passwords across multiple applications and platforms, attackers can potentially use the acquired information to obtain access to other resources on the network, where their confidentiality could be compromised.

Denial of Service (DoS) Attacks

DoS attack methods overload networks by making so many requests that regular traffic is slowed or completely interrupted. These attacks methods have existed since the Internet became open to the public—not only in theory, but also in practice. Although DoS attacks have been around for decades, the ugly offspring Distributed Denial-of-Service attacks are much newer, first seen in late June/early July 1999.

A “standard” DoS attack does not involve breaking into the target; instead, the attacker’s goal is to simply overload the target (router or web server) with so much fake traffic that it cannot cope. When the target cannot cope, genuine users cannot connect and are therefore denied service.

Smurf Attacks

A smurf attack is a type of DoS attack that exploits the use of the Internet Control Message Protocol (ICMP, aka ping) and the IP’s network and broadcast addresses. A smurf attack’s purpose is to disable a target host or network by consuming all its resources; aside from this, it causes no permanent damage. Every IP subnet has two special addresses:

- The network address, which is the first address in the subnet
- The network broadcast address, which is the last address in the subnet range

The IP network address serves as the identity address of a given subnet in the IP routing table. The IP broadcast address was devised as a method for sending information to all the hosts in a given subnet. Most IP implementations respond to messages with the network or broadcast address as the source address. This support is known as *directed broadcast*. This feature is also data used for legitimate purposes.

Fraggle Attack

Fraggle is an attack similar to a smurf attack, but instead of using ICMP, it uses UDP. The attack broadcasts a spoofed UDP packet to the network, which in turn replies to the victim’s system. The larger the network, the larger the amount of traffic redirected to the victim’s system.

SYN Flood Attack

In TCP/IP, a three-way handshake occurs whenever a client attempts to connect to a service, such as FTP or HTTP. The three-way handshake is defined as follows:

1. The client sends a packet with the SYN (synchronization) flag in the TCP header set to the service.
2. The service responds with a SYN-ACK (synchronization-acknowledgment).
3. The client sends a handshake to the service (a SYN-ACK transaction) and the session is considered established so data begins to flow.

Note Richard Stevens (*TCP/IP Illustrated*, p. 231) explains the format of the SYN packets and is an excellent resource for those wanting to understand the details of TCP/IP.

A SYN flood attack is when the client does not respond to the service’s SYN-ACK, thereby tying up the service until it times out. In this type of attack, the client never responds because the client’s source address is forged (spoofed). The attacker’s goal is to send SYN packets to the service faster than it takes for the service to timeout waiting for the client’s SYN-ACK response. This causes the service to become so busy acknowledging the SYNs

and waiting for the client that it cannot answer requests for service from legitimate users and therefore denies them service.

Teardrop Attack

A teardrop attack uses fragmented UDP packets. The first fragment is fine, but the second packet overwrites part of the first fragmented packet. This results in a memory error, and the system crashes.

Distributed Denial-of-Service

A Distributed Denial-of-Service (DDoS) attack generates false traffic from multiple hosts across the Internet. A DDoS attack uses multiple computers throughout the network that it has previously infected with a *DDoS daemon* (program); these computers are then known as *zombie computers*.

Note A DDoS daemon is a specialized computer program designed for use in controlling and coordinating a DDoS attack. As of this writing, there are five known programs: Tribal Village (TFN); Tribe Flood Network 2K (TFN2K); Trinoo; Stacheldraht, which is German for “barbwire”; and Trinity. You can learn more about these programs by visiting the following URLs:

<http://staff.washington.edu/dittrich/misc/trinoo.analysis>
<http://staff.washington.edu/dittrich/misc/tfn.analysis>
<http://staff.washington.edu/dittrich/misc/stacheldraht>.
http://packetstormsecurity.org/distributed/TFN2k_Analysis-1.3.txt

These zombie computers all work together as a *zombie network* to send out bogus connection messages, thereby increasing the amount of open connections with which the target must deal. This prevents legitimate users from accessing the services being offered. A DoS attack can be conducted using various bogus connection techniques. Table 12-1 provides a listing of the various DDoS tools and the ports used to communicate between the various components.

Table 12-1 *DDoS Tools Communication Matrix*

DDoS Tools	Attacker to Master Communication	Master-to-Daemon Communication	Daemon-to-Master Communication
Trinoo	Port 27665/tcp	Port 27444/udp	Port 31335/udp
Tribal Village (TFN)	ICMP Echo/Echo Reply	ICMP Echo reply	ICMP Echo/Echo Reply
Stacheldraht	Port 16660/tcp	Port 65000/tcp	ICMP Echo reply
Trinity	Port 6667/tcp	Port 6667/tcp & port 33270/tcp	
Shaft	Port 20432/tcp	Port 18753/udp	Port 20433/udp

DoS attacks are easy to implement and can cause significant damage, thereby disrupting a server, website, or network's operation and effectively disconnecting them from the Internet. DoS attacks attack the work differently based on the type of attack and which section of the network architecture are being targeted. Following are types of attacks:

- Bandwidth
- Logic
- Protocol

For example, a SYN flood is a protocol attack that uses fictitious, half-open TCP connection requests that exhaust the resources of the targeted system.

Preventing DoS Attacks

You might be wondering how you can defend against DoS attacks. This is perhaps one of the most difficult attacks to defend against because many of the attacks come in the form of traffic that would be considered a normal occurrence on your network. Perhaps one of the most common defenses is to rate limit certain types of traffic. For example, you might want to allow ping (ICMP); however, too much of it could be considered a DoS attack, so you would rate limit ICMP. In contrast, you must carefully watch other types of traffic. Perhaps limiting HTTP (web) traffic to your Internet e-commerce site would be a mistake!

There is no surefire method by which you can protect yourself 100 percent against DoS attacks. They (the attackers) continually take advantage of bugs and exploits present in the operating system, JavaScript, and Internet browser (Mozilla, Internet Explorer, and so on); even the Adobe Flash Player is not immune to vulnerabilities. The best method I have found is to keep an organization's computing environment up to date with the latest security patches, allow only necessary traffic into, and out of, your network by using defense in depth, (honeypots, ACLs, firewalls, and DMZs), and monitor the traffic coming in from outside by using an IDS or IPS. Having to effectively defend against every DoS attack type strikes fear into your IT and security staff. Carefully consider these restrictions, their impact on your business, and the risk mitigation should an attacker successfully launch a DoS attack. I've said it before and I will continue to say it: You must educate your staff on the risks and how to identify and protect your assets.

Other Types of Attacks

When planning an attack on someone, be it a nation or corporation, you must use all weapons you have to obtain your objective. It is the same way with hackers and malcontents who want nothing more than to disrupt your corporation's day-to-day business for political gain, personal vendetta, or boredom. The following sections describe a few other attack types or vectors that you need to be aware of.

Ping of Death

A ping of death attack uses the characteristics of ICMP to the attacker's benefit. The compromised hosts are directed to attack the designated target via a continuous stream of ping packets. This causes an incredible number of ping requests coming from thousands of compromised hosts to begin impacting the host. This is certainly an unwelcome scenario, but the attacker has also altered the ICMP packet (ICMP echo request). Each packet does not contain the compromised host's source address; instead, each packet's source address is the target's address. (That is, the source and destination address are the targets.) The target system then transmits a response (ICMP echo reply) to each packet, which is destined to itself, thus causing traffic to increase exponentially until the target crashes because of its inability to handle such a high volume of traffic. To take this further, imagine if a broadcast were allowed onto a LAN with the target's source address. The real caveat implicit in this attack is that the traffic appears completely normal and is typically allowed into any network and through firewalls, and so on. You can see this attack in action with a properly positioned sniffer or probe.

Man-in-the-Middle Attacks

In a Man-in-the-Middle (MitM) assault, the attacker places himself in the middle of a communication flow between two hosts: usually a server and a client. The attacker then intercepts messages transmitted between the two hosts. The attacker can look for a variety of things. Perhaps he wants to see how much money is in your bank account, your password to the fantasy football website, or even block the connection. The interesting part and point of concern here is that if the attacker does not alter anything, you will not know that the packets are being intercepted by an attacker in the middle!

Network sniffers, such as Ettercap (<http://ettercap.sourceforge.net/>), are often used to accomplish this type of attack. You can also find that MitM attacks can be used to reconstruct public cryptic keys. This discussion is beyond the scope of this text; however, protecting passwords and keys is always a good idea in case you need another reason for a password policy.

ARP Spoofing (aka ARP Poisoning)

ARP spoofing is one way in which a MitM attack can be successful if executed on either a wired or wireless LAN. The process of updating a target computer's ARP cache with forged entries is referred to as *ARP poisoning*. This technique involves the attacker constructing a forged ARP request and reply packets to change the Layer 2 Ethernet MAC address to one of his choosing. By the attacker sending forged ARP replies, a target computer could be convinced to send frames destined for Host B to instead go to the attacker's computer first so that they can be recorded and read. When done properly, Host A has no idea that this ARP redirection took place, as shown in Figure 12-1.

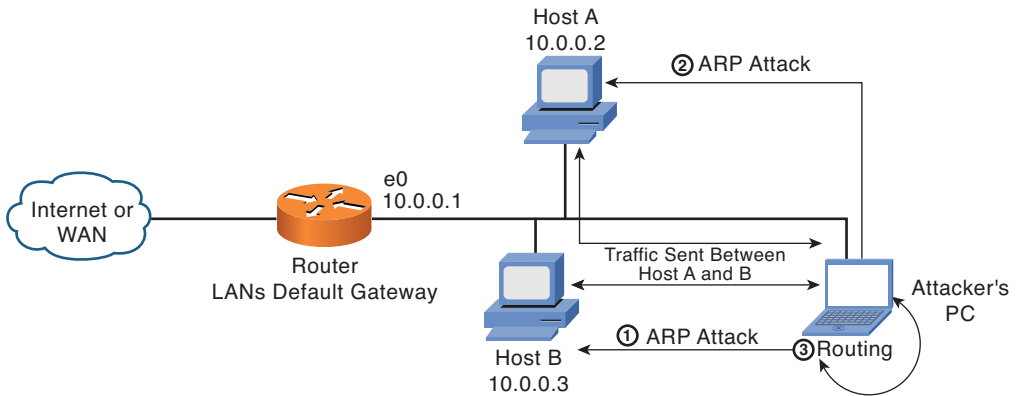


Figure 12-1 *Man in the Middle: ARP Spoofing*

IP spoofing plays an important role in MitM attacks. In IP spoofing, an attacker intercepts a legitimate communication between two parties. The attacker then controls the flow of communication and can eliminate or alter the information sent by one of the original parties without either the sender or receiver being aware. By doing so, an attacker can fool a victim into disclosing confidential information by spoofing the identity of the original sender.

Back Doors

A back door, or trapdoor, is a secret way of gaining access to a program, operating system, BIOS, or network service. You often see these types of back doors in computer games where a certain phrase or key combination provides you unlimited money, power, health, and so on. Back-door entry to resources can be accidentally or intentionally opened by users or by design; consider the following examples:

- Deliberately placed by system developers to allow quick access during development and not turned off before release.
- Placed by employees to facilitate performance of their duties because the “proper procedure” made them think it made their jobs harder, so there must be a smarter and easier way. Users might not be as technical as your IT staff, and often they find back doors because they do not have a preconceived notion of how something should work.
- Normal part of standard “default” operating system installs that have not been eliminated by OS hardening, such as retaining default user logon ID and password combinations. Again, vendors do not want technical support calls, so they make it as easy and open as possible. This means that your IT staff must review and harden every server!
- Placed by disgruntled employees to allow access after termination. In many cases, an employee suspects that the loss of his job is coming, which makes him angry and feel unappreciated, so he wants to ensure that he can strike back as necessary when the time comes.

- Created by the execution of malicious code, such as viruses or a Trojan horse that takes advantage of a vulnerability in an operating system or application.

As discussed previously, these attacks and tools are the most common types of vulnerabilities used by attackers. Understanding them will better allow you to understand the tools and techniques discussed in the following section.

LAND (Local Area Network Denial) Attack

A LAND attack is a DoS that consists of using a type of IP spoof-based attack where the source and destination address are the same. This attack crashes some TCP/IP implementations that do not know how to handle the packet. This is a rarity in terms of appearance in the real world but is a standard signature on ISS, NFR, Dragon, and Cisco Net Ranger and IDS-IOS.

Xmas Tree Attack

In an Xmas tree attack, a TCP packet sent to any known service port sets all the code flags (URG, ACK, PSH, RST, SYN, and FIN). An alternative version of this attack is a TCP packet without any flags set. Both cases are the result of packet craft and do not exist in the wild.

Ping Pong Attack

Following are two variations on the ping pong attack:

- A flood of spoofed packets to the echo service (UDP/TCP port 7), which is a simple service that echoes back any data sent to it.
- Sending a spoofed UDP message that appears to be from the chargen service port to the echo (UDP port 7) service on another system. The chargen service responds to any packet sent to the service port with a 72-byte random character string. After the spoofed connection is established, the echo port sends traffic to the chargen port and a loop develops.

Both variations consume CPU resources; enough attacks cause the system to become CPU-bound and crash.

Firewalking

Many people consider firewalls immune to attacks or standard techniques that enable attackers to figure out their rule sets to bypass them. That belief was true for a while; however, new techniques are always being developed and, in this context, *firewalking* is a concept that enables the attacker to send specially crafted packets *through a firewall* to determine what ports and services are permitted through the firewall. Attackers with this knowledge can make their port scans *bidden* and thus map your network through your firewall.

Firewalking works because IP packets contain a field that prevents them from being sent around a network forever. This field is known as Time-To-Live (TTL). When this field reaches zero, the packet is discarded. In firewalking, this field is set to a value that enables the packet to get beyond the firewall and then be dropped by a host or device after the

firewall. What enables this to happen is that the TTL value is one of the first things checked and, if its value is zero, a device sends back a packet acknowledging that it is being dropped without the original packet ever actually being processed.

Security Assessments and Penetration Testing

Companies with security offerings these days often have a *security assessment* as their first step in assisting a client in securing their network. A security assessment is an excellent first step for an organization concerned with understanding the extent of the security on their network (and its effectiveness). A *strongly recommended* practice is that individuals outside your organization perform security assessments on a yearly basis. This provides an objective and honest evaluation of your security, and because vulnerabilities are always being discovered, your network would be evaluated often enough to understand its effectiveness. A variety of available types of security assessments exist:

- Internal vulnerability and penetration
- External vulnerability and penetration
- Physical security

Before arranging a security assessment of any sort, you should learn more about the processes and procedures that the vendor is going to use. Too many security service companies exist to risk your company's security without some due diligence; you should review the following paragraph:

Understand the plan for the security assessment. If not planned and understood, assessing the actual network vulnerabilities can cause havoc in your network. There must be a legal agreement on the scope of the testing and the extent to which it will go; this protects both parties. Finally, it is important to define the success criteria of an assessment so that both parties understand what is to be accomplished.

The following sections examine the recommended approach that you should take and the benefits to the security of your network for each type of assessment.

Internal Vulnerability and Penetration Assessment

According to a recent study by the FBI, internal users and processes account for more than 60 percent of network security threats in today's enterprises. These threats are a result of improper configuration of network devices, lack of effective security procedures, and outdated and unpatched software. Security consultants should be able to identify these threats to determine your network's level of risk to intentional or accidental threats.

Today's organizations find it difficult to stay up to date on the numerous new vulnerabilities found each day in operating systems and applications. Security consultants should be aware of the latest vulnerabilities and help you assess the state of your internal network security mechanisms. They should also be able to recommend corrective steps for moving forward with your organization's security goals.

Assessment Methodology

Internal network security assessments must be performed onsite at your location and focus on internal security risks associated with policies, procedures, and networked hosts and applications. At a minimum, a security consultant should perform the following work:

- Gather customer-provided network information, if applicable.
- Gather and document publicly available network information for your review so that you can understand what an attacker would know.
- Perform network mapping techniques to determine the topology and physical design of your network.
- Perform network application probing and scanning.
- Consider OS fingerprinting and vulnerability detection to expose vulnerable hosts.
- Identify traffic patterns and flows to compare with expected normal business expectations.
- Detect any potentially weak user authentication systems, such as users who never change passwords or insecure wireless networks.
- Vulnerability analysis using public, private, and custom tools.
- Manually verify all detected vulnerabilities to ensure that false positives are not reported.
- Observe internal security practices and policies throughout your network.
- Analyze findings and report analysis along with specific recommendations for moving forward.

The end result of an internal risk assessment should be a document that contains the assessment methodology, work performed, and details gathered on every system, including the high-risk systems found vulnerable to attack and detailed lists of vulnerabilities. The assessment results document provides a clearer picture of your network architecture and security risks. The document should also contain the results of all work performed and conclusions from each test phase about the remediation required and the relative priority of these recommendations. Of course, this document must also include recommendations for mitigating detected network security risks in a cost-effective manner.

External Penetration and Vulnerability Assessment

As traditional business systems become more distributed among an organization's geographically disperse locations, the risk of external attacks increases. These risks are further exaggerated by improper router and firewall configuration and insecure, outdated, or improperly configured web-based applications.

Today's small and medium-sized businesses find it difficult to stay up to date on the numerous new vulnerabilities found each day in operating systems and applications. There

are numerous security firms that can help you assess the state of your current perimeter defense mechanisms and recommend steps for moving forward with your organization's security awareness.

Assessment Methodology

External penetration and vulnerability assessments are performed against your network at places where it interacts with the outside world. This could be through connections to the Internet, wireless, phone systems, and other remote access locations. The intent of this type of security assessment is to determine where and how your network is vulnerable to external attacks.

In many cases, an external assessment and an internal security assessment look at the same types of things. The difference is the point of view, and in this case it is from the outside trying to look in to see what can be discovered. The following list examines the work that should be done for an external penetration and vulnerability assessment:

- Gather customer-provided network information, if applicable.
- Gather and document publicly available network information for your review so that you can understand what an attacker would know.
- Perform stealthy network mapping techniques to determine your network's topology and physical design and to see whether these simulated attacks can be detected.
- Perform network application probing and scanning.
- Look for firewalking, wardialing, and wardriving, as needed. Firewalking has already been discussed. Wardialing is a technique of using a modem to automatically scan a list of telephone numbers, usually dialing every number in a local area code, to search for computers, bulletin board systems, and fax machines. Wardriving is the act of searching for Wi-Fi hotspots or wireless networks by a person in a moving vehicle, using a laptop, PDA, or smartphone (there is an "app" for that), and even game consoles (Nintendo DS and Sony PSP); *Treasure World* for the DS is a commercial game in which gameplay completely revolves around wardriving.
- Use OS fingerprinting and vulnerability detection to expose vulnerable hosts.
- Identify traffic patterns and flows to compare with expected normal business expectations.
- Detect any potentially weak user authentication systems, such as users who never change passwords or unsecure wireless networks.
- Perform vulnerability analysis using public, private, and custom tools.
- Manually verify all detected vulnerabilities to ensure that false positives are not reported.
- Analyze findings and report analysis along with specific recommendations for moving forward.

The end result of an external penetration and vulnerability assessment is a document that contains the same level and type of information as an internal assessment, except from an external point of view. Although this chapter separately examines internal and external assessments, these assessments are best performed together in the real world. They then provide a clearer picture of your network's security, end-to-end.

Physical Security Assessment

This book focuses on the logical security of networks, which is only part of the coverage that this type of assessment provides. Many assets are physical in nature and can be harmed through cruder and perhaps simpler methods than have been discussed. For example, are your IT resources kept in a room with overhead water-based sprinklers? If so, that is not physically secure because microchips and water do not mix. A simple DoS would be to trigger the fire alarm in your building and let the water do the rest. I hope your tape backups are protected from water damage and that they are current.

Although this is a digital age, today's IT systems still depend on physical hardware and reside in physical locations. Without the use of proper physical security mechanisms, all other security measures in place can be defeated. As the sensitivity of an organization's information increases, physical security takes a more important role. What good is it to have the latest firewall, IDS, and VPNs if you leave the door open to your equipment?

Physical security controls can be either deterrent or detective in nature and are designed to limit your organization's exposure to physical threats. A physical security risk assessment can help your organization design and implement cost-effective physical security measures to deter would-be attackers, monitor suspicious activities, and ultimately protect your valuable corporate resources from tampering, compromise, or destruction.

Assessment Methodology

A physical security assessment must be performed onsite at your location and focus on physical security measures and internal practices of a physical nature that are in place to protect your network resources. A physical security assessment should entail the following:

- Observe external building access points and safeguards in place.
- Observe physical safeguards in place, such as closed-circuit cameras, badge access, and visitor sign-in practices.
- Review physical protection mechanisms for IT resources, but also paper records.
- Determine physical safeguards in place for securing IT equipment, such as restricted access to computing environment, floppy-drive locks, redundant power sources, and protected data communication channels.
- Observe employee habits as related to physical security.
- Observe the physical disposal methods of critical data; do you recall dumpster diving?

- Make recommendations for securing your IT resources from physical security breaches.
- Understand the backup procedures and storage of critical data.
- Examine vendor and visitor access policies (if they exist) to determine how unknown individuals are handled.

The end result of a physical risk assessment is a document that contains the methodology followed, the work performed, the results of the work performed, and recommendations for mitigating detecting physical security risks in a cost-effective manner.

Many of these assessments cannot be automated to any great degree, so you must open your network and its resources to a trusted outside organization. When selecting this organization, you should request the following:

- Review of industry standard certifications to ensure that there is at least a measurable level of competence associated with those who are assessing your network.
- Contact several references of the company you are thinking about using and make sure that the references are relevant to the services you need performed.
- Ask for and review sample assessments. This can be difficult to do because assessments usually contain sensitive customer data, but any company committed to providing security services should have the capability to show you a sanitized version.
- Set expectations and deliverables clearly in the agreement to proceed or contract and so forth, thereby protecting yourself and the vendor's employees. Clear communication can solve 99 percent of the world's problems.
- Ask the security company to walk you through the assessment process before it comes to your location. If it cannot recite the process from memory, chances are it has either not been in business very long or the person you are speaking with is not a field technician.

Miscellaneous Assessments

Following are other types of assessments related to security in some ways that you should consider:

- **Procedural risk assessment:** This assessment enables security professionals to review your security policies and procedures to ensure that they conform to best practices. Chapter 2, "Security Policies," discusses policies and procedures of this nature.
- **Disaster recovery:** If your organization is based in an area of the world that is susceptible to tornados, hurricanes, earthquakes, lightning strikes, floods, fire, or some combination of these, the need for a plan to recover your network infrastructure and critical data becomes more important with every passing day. The influence and persuasiveness of IT is ever-increasing.

- **Information handling security assessment for banks and medical offices:** With new legislation for the security of financial and medical records (HIPPA for medical and Gramm-Leach-Bliley Act for financial) coming out each year, professions tasked with maintaining these types of records must meet increasingly higher data security standards or face jail time.

Assessment Providers

A simple Google search on security assessments reveals more than 400,000 hits, and this number will continue to grow. Some companies, such as the following, are worth mentioning as excellent providers of assessments services:

- **Cisco Secure Consulting Services:** Provides enterprise customers with comprehensive security analysis of large-scale, distributed client networks externally from the perspective of an outside hacker and internally from the perspective of a disgruntled employee or contractor, according to its website. You can learn more at www.cisco.com/go/securityconsulting.
- **Qoncert:** Provides customized security solutions and assessments for customers of all sizes with a specialization in ensuring that business focus drives the security solution versus the more common occurrence of IT driving business. You can learn more at www.qoncert.com.

Security Scanners

When hackers want to breach your systems, they typically look for well-known security flaws and bugs to attack and exploit, some of the most common of which have been discussed in earlier chapters of this book. A true attack, which has the end goal of penetrating and controlling the target system, relies on the attacker gathering the most accurate and comprehensive view of an organization's security. As attackers evaluate the network, they exploit vulnerabilities to determine precisely how to get control of valuable information assets.

Attacking vulnerabilities used to be a time-intensive procedure that required a lot of knowledge on the part of the attacker. Today, however, automated tools have changed all that. Gone are the days of having to figure out the publicly available exploit codes and maintaining them all to be effective. Now a Google search provides anyone with enough information to be dangerous.

You might be wondering whether we are talking about the role of attackers and commiserating with them about how hard it is to control the data. Or perhaps we are talking about how network administrators are faced with such a daunting task. Actually, we are talking about both, and the point is that this is no longer the biggest concern. Today, attacks have been scripted and published, and companies have formed to automate the detection of attacks and exploit vulnerabilities. This section looks at the most comprehensive of these tools that, included with the individual tools covered elsewhere in this book, should provide you with excellent resources to detect vulnerabilities and begin correcting them.

Features and Benefits of Vulnerability Scanners

Applications that perform security scan and vulnerability assessments do the scanning and calculations in the background. Frankly, the focus here is not in *how* vulnerabilities are detected, but on *what* is vulnerable. The following four categories summarize the value of these scanning tools:

- **Scan and detection accuracy:** Scans and reported vulnerabilities must be accurate with minimal false positives—defined as normal activity or configuration that the system *mistakenly* reports as malicious. The opposite also holds true then: There can be no false negatives—defined as malicious activity that is not detected.
- **Documentation and support:** Must be clear, concise, well written, and easy to understand (like this book). This includes reporting documentation and application operation so that users can figure out how to make the application work and see the documented findings in the report.
- **Reporting:** The most important aspect of a vulnerability scanner is when you need to know the next steps after a vulnerability has been detected (that is, what was detected and how to fix it); thus, a report must be customizable, useful, and accurate.
- **Vulnerability updates:** New vulnerabilities are constantly being released and, with today's technology, every system should have a way to automatically update.

Freeware Security Scanners

Your corporation needs to have the capability to scan for attacks, identify them, and notify you that they are occurring. To help bridge the gap between nothing and a robust multifeature solution, several different freeware security scanners are available. You can pull one off the Internet and implement it.

Metasploit

The Metasploit Framework is a penetration testing toolkit, exploit development platform, and research tool that provides useful information and tools for penetration testers, security researchers, and IDS signature developers. The framework includes hundreds of working remote exploits for a variety of platforms. It is written in the Ruby scripting language and is provided on the BSD license. The latest version is 3.5.0. You can find more information about Metasploit on www.metasploit.com. It is also part of BT5, as mentioned in Chapter 1, “There Be Hackers Here.”

NMAP

Network Mapper (NMAP) is a free and open source (license) utility for network exploration or security auditing. Many systems and network administrators find it useful for tasks such as network inventory, managing service upgrade schedules, and monitoring host or service uptime.

NMAP uses raw IP packets to determine what hosts are available on the network, what services (application name and version) those hosts offer, what operating systems (and OS versions) they run, what type of packet filters or firewalls are in use, and dozens of other characteristics. It was designed to rapidly scan large networks but works fine against single hosts.

NMAP runs on all major computer operating systems, and official binary packages are available for Linux, Windows, and Mac OS X. In addition to the classic command-line NMAP executable, the NMAP suite includes an advanced GUI and results viewer (Zenmap); a flexible data transfer, redirection, and debugging tool (Ncat); and a utility for comparing scan results (Ndiff).

The most recent version is 5.35DC1. You can read more about NMAP and its functionality at <http://nmap.org>.

SAINT

SAINT is a product suite that offers a complete solution to evaluate the threats and vulnerabilities that affect your network. It consists of SAINTscanner, SAINTexploit, SAINTmanager, and SAINTwriter.

SAINTscanner scans your network to detect anything that could allow an attacker to gain a foothold, create a DoS, or obtain access to sensitive information. SAINT offers heterogeneous scanning that identifies vulnerabilities across operating systems, desktop applications, network devices, web applications, and databases.

SAINTexploit is the penetration testing component that is integrated with SAINTscanner. SAINTexploit automates the penetration testing process, examines vulnerabilities discovered by the scanner, exposes where the attacker could breach the network, and exploits the vulnerability.

SAINTmanager is a remote management console for organizations that want to centrally manage multiple scanners and help manage the vulnerability life cycle.

SAINTwriter is the report writer. It is built in to all SAINT's products for reporting on vulnerability assessment, penetration testing, trouble tickers, and vulnerability management.

You can find more information on SAINT and its entire suite of products at www.saintcorporation.com.

Nessus

Nessus is to vulnerability detection what Snort is to IDS: an open source solution supported by a community of Internet volunteers. You can learn more about Nessus at www.nessus.org.

In Their Own Words

The following section is a direct quote from the Nessus web page on how it describes its product:

The “Nessus” Project aims to provide to the Internet community a free, powerful, up-to-date and easy to use remote security scanner.

A security scanner is software, which will audit remotely a given network and determine whether bad guys (a.k.a. “crackers”) may break into it, or misuse it in some way.

Unlike many other security scanners, Nessus does not take anything for granted.

That is, it will not consider that a given service is running on a fixed port—that is, if you run your web server on port 1234, Nessus will detect it and test its security. It will not make its security tests regarding the version number of the remote services, but will really attempt to exploit the vulnerability.

Nessus is very fast, reliable and has a modular architecture that allows you to fit it to your needs.

Scan and Detection Accuracy

Scans and reported vulnerabilities must be accurate, with minimal false positives—defined as normal activity or a configuration that the system *mistakenly* reports as malicious. The opposite also then holds true: There can be no false negatives—defined as malicious activity that is not detected.

Nessus has good capabilities to detect vulnerabilities and is accurate in the vulnerabilities it detects and finds. Being an open source project, Nessus is constantly being watched, tested, studied, and improved upon. Technical support can be purchased through Tenable Network security. With so much visibility, this product has become configurable for those with the knowledge to understand its underpinnings. Figure 12-2 shows the Nessus setup screen and its flexibility, strength, and many possible options.

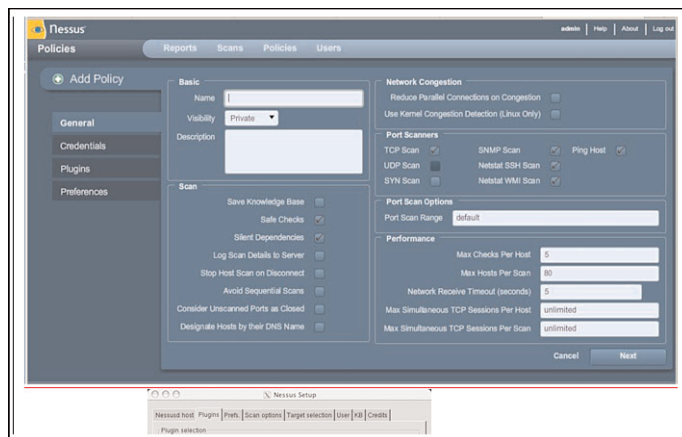


Figure 12-2 Nessus 4.2.2 Vulnerability Policy Screen

Documentation and Support

Documentation must be clear, concise, well written, and easy to understand. This includes reporting documentation and an application operation so that users can figure out how to make the application work and see the documented findings in the report.

Nessus documentation is excellent. The installation and configuration guides are online and easily downloaded for all operating system platforms. It took me 10 minutes to download, obtain a key, and register the software for my Microsoft XP Pro laptop and then, another 30 minutes to download and install the plug-ins. There was no need to add or remove any third-party software. No MAN pages to wade through, and no fumbling around a command-line interface.

Although the Nessus Project began as a free remote security scanner, in 2005, Tenable Network Security (the company that was co-founded by the Nessus developer) changed Nessus 3 to a proprietary license. In 2008, Tenable revised the feed license, which enables home users full access to the plug-in, and began charging for a professional license. What does all that mean? It means Nessus 4.2.2 is free of charge for personal or home use; however, in an enterprise environment this software license must be purchased and registered before you can use it. Technical support can be purchased for both the home and professional versions; however, there is a mailing list and forum that has many of the core Nessus programmers, and they can be helpful.

Reporting

The most important aspect of a vulnerability scanner is when you need to know the next steps after a vulnerability has been detected (that is, what was detected and how to fix it); thus a report must be customizable, useful, and accurate.

Nessus creates reports in a variety of formats; the most useful is HTML. These reports are fully hyperlinked with complete analysis of the vulnerabilities detected, their level of risk to your network, and great, pretty pictures that can visualize the vulnerabilities. On a downside, the reports are UNIX-centric and full of contextual and grammatical errors, not something that you can share without serious editing. The information is accurate, just not as polished as a commercial product.

Vulnerability Updates

New vulnerabilities are constantly being released and, with today's technology, every system should have a way to automatically update.

Nessus is kept up to date via scripting that can be automated to ensure that it has the latest signatures. Nessus runs on the following platforms: Microsoft Windows, Mac OS X, Linux, FreeBSD, Solaris, and iPhone.

Nessus is a good vulnerability scanner that has exceptional functionality as a result of its open source status. There are some excellent resources available for Nessus. You can find other related links at the following URLs:

www.securityprojects.org/nessuswx/

www.nessus.org/

Nessus 4.2.2 Home Edition is freely available for download and requires no purchase because it is open source software.

Retina Version 5.11.10

Retina is eEye's premiere security scanner that leads its suite of security products. eEye offers several products that focus on securing the Microsoft product line. You can learn more about Retina at www.eeye.com.

In Their Own Words

The following section is a direct quote from the eEye corporate web page on how it describes its vulnerability scanner, Retina:

Retina, founded from over a decade of technology innovation by eEye's world renowned security research team, is an integrated end-to-end vulnerability and compliance solution designed to help organizations with protection and compliancy by defining and monitoring relevant IT controls.

Retina monitors both patch and configuration vulnerabilities and compliance to pre-defined configuration baselines and provides automated notification of violations. The environment is assessed, capturing established security controls along with any vulnerabilities or configuration violations that impact the network. Detailed reports providing prescriptive guidance and recommendations are then forwarded and response is initiated to ensure that corrective action can be taken in a timely fashion.

Retina's management console is a fully integrated and rich Internet-enabled application for security and compliance management. Now you can simplify the management of distributed, complex infrastructures while protecting your mission critical assets from evolving threats with a single unified management system.

Scan and Detection Accuracy

Scans and reported vulnerabilities must be accurate with minimal false positives—defined as normal activity or configuration that the system *mistakenly* reports as malicious. The opposite also holds true, then: There can be no false negatives—defined as malicious activity that is not detected.

Retina has an excellent presentation interface for the execution of scans; it is intuitive and comes with a variety of other tools, so it can be used for more than just a vulnerability scanner. One of the best features that Retina provides is the easy customization, scheduling, and penetration audit customization. What sets Retina apart is the capability to create scanning policies with different scans for different devices. For example, you can scan Internet servers differently than employee PCs. Figure 12-3 shows an example of targeting a specific device in Retina.

Documentation and Support

Documentation must be clear, concise, well written, and easy to understand. This includes reporting documentation and application operation so that users can figure out how to make the application work and see the documented findings in the report.

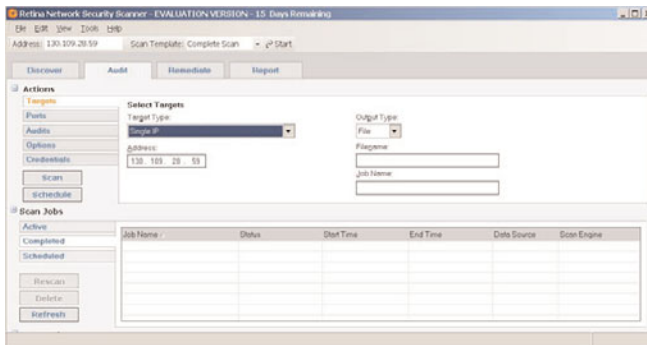


Figure 12-3 *Selecting a Target Range in Retina*

Retina documentation is included in the Windows help file and appears to be complete, answering many of the questions a typical user would have. It does not contain many in-depth how-to's, but it provides enough examples that their lack is not a hindrance. A web-based form submitted to the eEye technical support team provides only support options for users.

Reporting

The most important aspect of a vulnerability scanner is when you need to know the next steps after a vulnerability has been detected (that is, what was detected and how to fix it); thus, a report must be customizable, useful, and accurate. Figure 12-4 shows the summary of vulnerabilities after Retina completes its scan of your network.

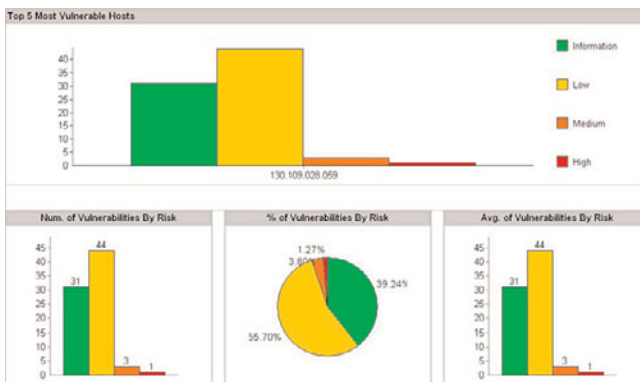


Figure 12-4 *Vulnerability Summary by Risk Level*

Like Nessus, Retina provides an overview of the detected vulnerability along with links to additional information and corrective actions, such as a Microsoft hotfix. Figure 12-5 shows specific detected vulnerability information, what the risk is, and where to find the manufacturer's fix.

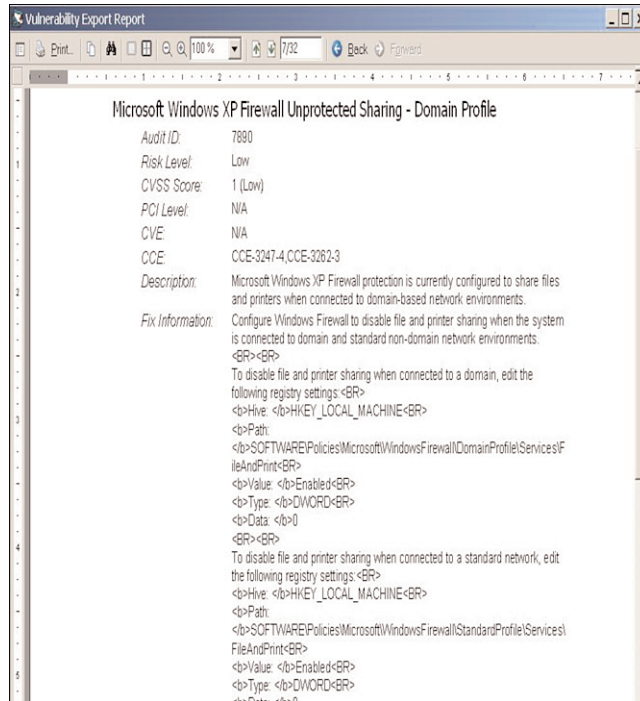


Figure 12-5 Vulnerabilities Details

Vulnerability Updates

New vulnerabilities are constantly being released and, with today's technology, every system should have a means of updating itself automatically.

Retina is exceptional in this regard; it can be configured not only to update its list of vulnerabilities, but also the application itself. Curiously, the open source movement has missed the boat on that feature. Retina takes a bit of getting used to, and it is an effective vulnerability scanner.

CORE IMPACT Pro (a Professional Penetration Testing Product)

Vulnerability identification, detection, and prioritization are all assessment functions. You can classify a product as penetration testing only if it actually exploits a given vulnerability. Vulnerability assessment and penetration testing complement each other. They do different things, so they must be two separate categories. This confusion is often encountered during the sales process. The penetration testing product picks up where the vulnerability scans leave off.

Vulnerability assessment does an adequate job of providing the tester with a snapshot of the current network configuration. Unfortunately, this snapshot does not address the impli-

cation of a successful intrusion to organizational assets. It relates only what the vulnerabilities are; it does not probe deeper to reveal what happens when the vulnerabilities are exploited. The following details the limitations of vulnerability assessments and scanners:

- Provides just partial information assurance
- Identifies only vulnerabilities; does not provide meaningful weighting of vulnerabilities or prioritization of remedies
- Produces a long list of potential weaknesses, often including numerous false positives
- Does not demonstrate what information assets can be compromised
- Cannot simulate real-world attacks
- Does not exploit trust relationships between network components, nor demonstrate the implications of a successful attack

CORE Security is roaring into the security penetration testing marketplace with its exploit product, CORE IMPACT. Yes, *exploit* product. (It runs applications in the product.) CORE IMPACT actually does not detect vulnerabilities; instead, it exploits a vulnerability and installs an agent on the targeted server. This agent then enables you to escalate attacks and own the target machine. The CORE IMPACT product eliminates the annoying and embarrassing occurrence of false positives. Although the following section discusses at length how CORE IMPACT achieves this, you can learn more about CORE IMPACT at www.coresecurity.com.

In Their Own Words

The following section is a direct quote from the CORE Security web page describing its product:

CORE IMPACT Pro enables you to perform frequent, realistic and effective penetration testing throughout your enterprise. After first identifying and validating any vulnerabilities that provide unauthorized access to your network, IMPACT Pro takes the testing process a step further by emulating multi-staged attacks that pivot between network systems, endpoints, web applications and wireless networks to access your organization's most valuable information and resources.

CORE IMPACT enables you to safely assess an organization's security posture against the top four attack methods that jeopardize data today.

The product's unified interface provides a consistent methodology for replicating data breach attempts that spread among these attack vectors. For instance, IMPACT can replicate an attack that initially compromises a web server or end-user workstation and then propagates to backend network systems. Only IMPACT allows you to utilize penetration testing to assess your information security in such an integrated, comprehensive, in-depth and seamless fashion.

Scan and Detection Accuracy

Scans and reported vulnerabilities must be accurate, with minimal false positives—defined as normal activity or configuration that the system *mistakenly* reports as malicious. The opposite also holds true, then: There can be no false negatives—defined as malicious activity that is not detected. IMPACT provides integrated Rapid Penetration Testing (RPT) capabilities across four attack categories:

- Network
- Client-side
- Web application
- Wireless

The four test approaches differ in the Information Gathering and Attack and Penetration stages. This is not a scanner; only limited scanning is possible during each of the RPT attack categories. For instance, during the information gathering phase using the Network RPT, information is gathered about the target network using network discovery, simple port scanning, and target operating system and service identification modules.

Documentation

Documentation must be clear, concise, well written, and easy to understand. This includes reporting documentation and application operation so that users can figure out how to make the application work and see the documented findings in the report.

CORE IMPACT Pro generates clear, informative reports that provide data about targeted systems and applications, results of end-user penetration testing tests, audits of all exploits performed, and details about proven vulnerabilities. These reports can be produced in HTML, PDF, or Microsoft Word formats. IMPACT Pro provides the framework to generate the following reports:

- **Activity:** Provides a detailed log of all testing activity that is being carried out, including the relevant data that organizations might need to share with auditors reviewing its security programs.
- **Attack Path:** A powerful visual representation of the manner in which tests can exploit individual vulnerabilities and achieve subsequent access to other systems and applications.
- **Client-Side Penetration Test:** Provides detailed results of assessments performed on endpoints and end users, including information about any social engineering tactics used to trigger tests.
- **Client-Side User:** Helps organizations understand exactly how well their end users stand up to social engineering attacks involving both email and web-based delivery models, including spear phishing assessments.

- **Delta:** Gives your organization an integrated view into vulnerabilities resident across a range of different assets, including network systems and client systems.
- **Executive Summary:** Offers a high-level view of penetration tests performed and understanding of how ubiquitous vulnerabilities are, where they reside, how they can be exploited, and where to begin remediation efforts.
- **FISMA Vulnerability Validation:** Provides results of penetration testing performed by government entities and other organizations working to remain compliant with the Federal Information Security Management Act of 2002 (FISMA).
- **Host:** Provides IMPACT Pro users with precise details about how their systems and applications can be compromised via real-world hacking or malware attempts.
- **PCI Vulnerability Validation:** Provides results of penetration testing performed with the goal of remaining compliant with the Payment Card Industry (PCI) Data Security Standard.
- **Trend:** Enables users to track data from up to 52 penetration tests over time, graphically representing changes in an organization's security posture as exploitable vulnerabilities are identified, remediated, and retested.
- **Vulnerabilities:** Provides IMPACT Pro users with specific details about all the weaknesses successfully exploited during penetration testing and how those flaws can be used by attackers to obtain control of a tested system and establish a beachhead for subsequent activity.
- **Web Application Executive:** Provides summarized information of every vulnerable web page found during testing and how those problems can be exploited by real-world attackers.
- **Web Application Vulnerability:** Provides comprehensive information about every security flaw that can be exploited during penetration testing, including those available to SQL injection, cross-site scripting, and remote file inclusion attacks.
- **Wireless Penetration Test:** Details wireless networks discovered, client-to-access point relationships, and access point profile information. This report also includes information about which networks were tested against attacks, which were successfully compromised, and which weaknesses allowed the compromise.

Normally, these reports are standard type reports that you would expect. What makes them unique, however, is that IMPACT enables them to be customized and printed according to the level of detail you want to present. For example, the report given to an organization's executive team should differ greatly from the report presented to the IT staff. IMPACT enables this level of customization.

Documentation and Support

The most important aspect of a vulnerability scanner is when you need to know the next steps after a vulnerability has been detected (that is, what was detected and how to fix it). Therefore, a report must be customizable, useful, and accurate.

When learning new software or applications, I find that it is important that the product has good documentation and support. This enables users to learn on their time versus other methods, such as training or scheduled web seminars (which I'm not a big fan of).

Vulnerability Updates

New vulnerabilities are constantly being released, and with today's technology, every system should have a way of updating itself automatically.

CORE IMPACT Pro provide real-time updates including new penetration testing exploits and tests for additional platforms as they become available. The support team from IMPACT advises you if and when new modules are published and provides a link enabling you to download them the same day, directly from within the IMPACT Pro software, which enables easy updating of the attack modules through a single click of a button. CORE Security is committed to making the product grow and evolve so it has an aggressive development schedule. You cannot find every possible vulnerability within CORE IMPACT; however, there are also continual updates in this regard. It is a challenge to determine exactly which vulnerabilities become modules and, so far, observations have shown that good choices and options are rather limited; however, they are quickly growing.

Chapter Summary

This final chapter covered several additional new vulnerabilities and described how they are used to attack systems. Understanding these common attacks is crucial for understanding what the rest of the chapter explained. Security assessments and penetration testing are effective tools that, if used correctly, enable your network to be evaluated by qualified engineers who deploy the proper security analysis tools to find the vulnerabilities. A good security assessment, however, covers more than just the logical vulnerabilities in your network. The remainder of this chapter was dedicated to the various security scanning tools that are available, some of which are free open source solutions.

Chapter Review Questions

1. What is the difference between a Man-in-the-Middle attack and a denial-of-service attack?
2. Define what a DDoS attack is and how it functions. How is it different from a standard DoS attack?
3. Name some common denial-of-service attacks.
4. Identify and explain three reasons that can result in a back door exploit being present on a system.
5. Define the concept of firewalking.
6. Where should an external penetration and vulnerability assessment be performed in your network?
7. When considering vulnerability scanners, why are a program's capability to conduct an accurate scan crucial?

This page intentionally left blank

Answers to Review Questions

Chapter 1

1. What is a target of opportunity?

Answer: A target of opportunity is one in which a vulnerability has been detected by an attacker, who decides to try an exploit because the target has enabled him to find it.

2. What is a target of choice?

Answer: A target of choice occurs when attackers choose you as a target. Their reason is irrelevant because this is a mental commitment on the part of the attackers.

3. What is the purpose of footprinting?

Answer: Footprinting is the process attackers take to understand a target's network and associated systems. This is a continuous process used throughout all planned attacks, and in which attackers want to gain as much information about the target as possible.

4. Which of the following are ways by which an attacker can gain access?

- a. Operating system attacks
- b. Application attacks
- c. Misconfiguration attacks
- d. Script attacks
- e. All the above

Answer: E. All the above

5. List four network security organizations.

Answer: CERT

SANS

SCORE

Security Focus

ICAT

Center for Internet Security

6. Briefly explain why it is important for attackers to cover their tracks.

Answer: Presuming that attackers have compromised a system, the ability to remove the forensic evidence of their actions (in other words, cover their tracks) enables the attackers to use the compromised system at their leisure if the system administrators never know they have been compromised.

7. Social engineering can be damaging without an overt attack happening. Explain why.

Answer: The purpose of social engineering is to trick a person into believing that the attacker is someone else and thereby allowing that person to believe that the attacker is entitled to sensitive information.

8. What kind of information might be found if an attacker dumpster dives at your place of work?

Answer: Perhaps there might be financial reports, customer lists, human resource information, or other sensitive data. The point here is to never simply throw out information that might have value.

9. DNS information gained through WHOIS is used for what kind of reconnaissance?

Answer: WHOIS information is used for passive reconnaissance.

10. What two free reconnaissance tools are available with most versions of the Windows operating system?

Answer: Nbtstat and net view.

Chapter 2

1. How important is it to involve other departments and employees in the crafting of security policies?

Answer: Involving your fellow employees is crucial to a policy's success. Their involvement allows everyone to understand and support the company's commitment to security.

2. True or false: It is a well-known fact that users circumvent security policies that are too restrictive. Explain your answer.

Answer: Absolutely true. The tighter you create your security policies, the harder it is for users to function effectively. Therefore, you must balance security and productivity.

3. What are three things you should keep in mind when writing or reviewing a security policy?

Answer:

Determine who gets access to each area of your network.

Determine what they can access and how.

Balance trust between people and resources.

Allow access based on the level of trust for users and resources.

Use resources to ensure that trust is not violated.

4. Why is it important to include an enforcement section in every security policy?

Answer: The enforcement section defines the penalty for failure to follow the policy. Dismissal is typically the most severe penalty, but in a few cases, criminal prosecution should be listed as an option.

5. An Acceptable Use Policy defines what kind of expectations for users?

Answer: An AUP defines the systems to be used for business purposes that serve the interests of the company, your clients, and your customers.

6. When and under what circumstances should you reveal your password to someone?

Answer: No one in a company should ever ask for your password; if a technical difficulty occurs, the password will be reset. Never reveal your password to anyone and, if asked, immediately report the request to corporate security.

7. Which of the following sample passwords would be considered effective when checked against the corporate password policy?

- a. wolfpack
- b. thomas67
- c. simonisnot4
- d. sJ8Dtt&efs
- e. Missing\$4u

Answer: D is clearly the correct answer because it has all the proper characteristics of a secure password as outlined in the password policy.

8. Define VPN and the role it can play within a company's network infrastructure.

Answer: A network is constructed using a public network such as the Internet to connect systems to a main site, typically the headquarters. VPNs use encryption mechanisms to protect data transmitted across the Internet. Additional protections are put in place to ensure that only authorized users or devices can connect via a VPN.

9. VPNs support a technology called *split-tunneling*. Define this technology and explain whether it should be used in a network.

Answer: Split-tunneling is a method of configuring a VPN, and it is either on or off. Essentially, if split-tunneling is on, users can connect to the corporate network and the Internet simultaneously. This presents a danger to the corporate network's security because if an attacker were to take control of the computer creating a VPN to the corporate network, the attacker can also gain access to the company's network via the VPN.

10. How frequently should security policies be updated or reviewed?

Answer: Ensure that your policies are updated annually, if not sooner, to reflect the changes of the past year.

Chapter 3

In lieu of review questions, Chapter 3, "Processes and Procedures," provides a list of references including checklists, best practice links, security websites, and the like that are useful for those implementing network security. Refer to the end of Chapter 3 for more information.

Chapter 4

In lieu of review questions, Chapter 4, "Network Security Standards and Guidelines," provides a comprehensive list of websites from Cisco, the NSA, and Microsoft related to network security standards and guidelines. Refer to the end of Chapter 4 for more information.

Chapter 5

1. What are the six security design concepts you should consider when looking at the security technologies for securing your network?

Answer: Layered security, controlling access, role-specific security, user awareness, monitoring, and keeping systems patched.

2. What rule is always implicitly present at the end of every packet filter?

Answer: Deny all packets.

3. When a device performs a stateful packet inspection, what characteristics in a packet's header are inspected, and why are they important?

Answer: Firewalls perform a stateful packet inspection and monitor the IP header information to track the status of a connection.

4. What are some limitations of a stateful packet inspection?

Answer: SPI cannot inspect or track every type of packet; for example, ICMP and UDP are not stateful.

5. Define the differences between public and private IP addresses.

Answer: Private addresses are for internal, non-Internet use. Public addresses are those used on the Internet.

6. Compare and contrast the three different versions of NAT, and identify which of them is the most commonly used.

Answer: Static, dynamic, and overloading. Refer to the bulleted list in the section "Network Address Translation (NAT)" in Chapter 5 for a full comparison. Overloading is the most commonly used form of NAT.

7. What are the two types of proxy firewalls?

Answer: Standard and dynamic firewalls.

8. Why is content filtering so important to networking?

Answer: Content filtering protects a company by restricting harmful websites.

9. What is the potential value of PKI to securing a network and e-commerce?

Answer: Seamless global security.

10. AAA provides security for what aspect of a network?

Answer: Network devices.

11. Search the Internet and find three potential vendors that can offer an effective RADIUS solution. Describe what features about each are beneficial.

Answer: Cisco ACS and Funk Steel belted RADIUS are two vendor-specific RADIUS solutions.

Chapter 6

1. How long, in bits, is the DES key?

Answer: 56 bits.

2. True or false: In 3DES, the same key is used to encrypt at each of the three stages.

Answer: True.

3. Define a hash in your own words.

Answer: By way of an analogy, a hash is a grinder that takes something recognizable, such as beef or pork, hashes it, and ends up with something unique that is based on the original. In this case, it is hamburger or sausage.

4. What creates a digital signature?

Answer: A hash.

5. Define authentication and provide an example.

Answer: Authentication is the process of identifying an individual or device based on the correct username/password combination.

6. Define authorization and provide an example.

Answer: Authorization defines what individuals are allowed to access. An example is the question “Have they been authenticated?”

7. A hash check occurs at what point in the operation of MD5?

Answer: When using a one-way hash operation such as MD5, you can compare a calculated message digest against the received message digest to verify that the message has not been tampered with. This comparison is called a *hash check*.

8. Of the security protocols covered in this chapter, which of them use generic routing encapsulation (GRE)?

Answer: PPTP and L2TP.

9. Describe several security benefits of L2TP.

Answer: Refer to the bulleted list in Chapter 4 in the “Benefits of L2TP” section.

10. What are the three core SSH capabilities?

Answer: Secure command shell, secure file transfer, and secure port forwarding.

Chapter 7

1. Who needs a firewall?

Answer: Everyone connected to the Internet or with IT resources to protect needs a firewall. Depending on a router and ACLs is an incomplete solution in layering your network’s defense.

2. Why do I need a firewall?

Answer: A firewall provides protection for your network resources through technologies such as SPI, which is not possible with any other device.

3. Do I need a firewall?

Answer: Yes, yes, yes; you need a firewall!

4. How is a firewall an extension of a security policy?

Answer: A firewall's rules reflect the network security policy that your organization has expressed in a written security policy.

5. What is the name of the table in a firewall that tracks connections?

Answer: State table.

6. What fundamental does a DMZ fulfill?

Answer: The DMZ protects Internet-accessible servers and services.

7. What are four benefits of a DMZ?

Answer: Auditing of DMZ traffic, locating an intrusion detection system (IDS) on the DMZ, limiting routing updates between three interfaces, and locating DNS on the DMZ.

8. Can firewalls enforce password policies or prevent misuse of passwords by users?

Answer: No, they cannot.

9. Do firewalls guarantee that your network will be protected?

Answer: Firewalls do not provide any sort of guarantee that your network will be protected; they are a tool for your use in building the layers of defense and protection needed.

10. Are all firewalls created equal?

Answer: No, not all firewalls are created equal; they are created different. It behooves you to understand the role and responsibility of the firewall prior to making a purchasing decision.

Chapter 8

1. Because every company that connects to the Internet has a router, should you deploy security on those routers?

Answer: Definitely! You have the router and this book, and you need to protect your network; use the knowledge presented here to go out and start some packet screening at the router. Layered security is best!

2. What is the value of edge routers being used as choke points, and how effective can they be in increasing your network's security?

Answer: The value of edge routers being configured as choke points is that they can prevent access to specific devices and applications in a performance-friendly way. This increase in security is typically provided through the use of standard and extended access control lists that can address traffic concerns at Layers 2, 3, and 4 of the OSI reference model.

3. Which four features from classic IOS Firewall features have been implemented in the Zone Based Policy Firewall?

Answer: Stateful packet inspection

VRF-aware Cisco IOS Firewall

URL filtering

Denial-of-service (DoS) mitigation

4. What are the two major changes to the way you configure IOS Firewall Inspection, compared to the Cisco IOS Class Firewall?

Answer: Introduction of the zone-based configuration or architecture and a new configuration policy language referred to as Cisco Policy Language (CPL).

5. Can the Cisco IOS IDS have multiple points of packet inspection?

Answer: Of course, you can have multiple points of packet inspection in the form of ACLs. The only requirement of the FFS and CBAC is that the filtering must occur after the inspection. Having the FFS determine access based on conversation direction maintains the capability for the router to still function primarily as a router.

6. Temporary access control lists have timers associated with them. Define how they function based on protocol (ICMP, UDP, and TCP).

Answer: ICMP and UDP sessions are removed based on configurable inactivity timers. TCP sessions are removed 5 seconds after the exchange of FIN packets. If an RST (reset) packet appears, the session is terminated and corresponding ACL entries are immediately removed.

7. What is the difference between atomic and compound signatures?

Answer: Atomic signatures are concerned with attacks directed to single hosts, whereas compound signatures look at attacks directed to groups of machines.

8. What happens when an attacker uses chargen and echo together? How would you stop this from occurring in a Cisco router?

Answer: Pointing the chargen service at the echo service creates a loop that causes an enormous amount of traffic to be generated and eventually overwhelms the router's CPU and RAM resources; therefore, this provides the makings of a serious denial-of-service attack (DoS). The easiest way to prevent this kind of attack is to disable these services on the router.

The commands to do so are **no tcp-small-servers**, which disables echo, chargen, discard, and daytime, and **no udp-small-servers**, which disables echo, chargen, and discard.

Chapter 9

1. Can you have unencrypted VPNs?

Answer: Yes; in that case, other protocols are used to handle the encryption.

2. What are the three types of VPNs?

Answer: Site-to-site, extranet, and remote.

3. Select three VPN features and benefits, and explain how your organization can directly benefit from each.

Answer: VPNs are secure, encrypted traffic and can link sites securely over the Internet.

4. VPN concentrators are designed for many users—explain how many and when you should use them.

Answer: VPN concentrators are built to handle the requirements of VPNs and are available in models suitable for everything from small businesses with up to 100 remote-access users to large organizations with up to 10,000 simultaneous remote users.

5. Does the VPN Client Software for PCs support Apple's powerful new operating system, Mac OS X?

Answer: Yes.

6. When does split-tunneling occur?

Answer: Split-tunneling occurs when remote VPN users or sites are allowed to access a public network (the Internet) at the same time that they access the private VPN network, without placing the public network traffic inside the tunnel first.

7. In relation to a data stream, what role does authentication play in securing it?

Answer: Authentication establishes the integrity of the data stream and ensures that it is not tampered with in transit. It also provides confirmation about the data stream origin.

8. When tunneling data in IPsec, what are the three protocols that play a role in process?

Answer: GRE, IPsec, and ISAKMP.

9. In site-to-site VPNs, what are the two different encapsulating protocols and what are the differences between them?

Answer: In site-to-site VPNs, the encapsulating protocol is usually IPsec or generic routing encapsulation (GRE). GRE includes information about what type of packet you encapsulate and about the connection between the client and server. The difference depends on the level of security needed for the connection, with IPsec being more secure and GRE having greater functionality. IPsec can tunnel and encrypt IP packets, whereas GRE can tunnel IP and non-IP packets. When you need to send non-IP packets (such as IPX) over the tunnel, use IPsec and GRE together.

10. Name three of the benefits of IKE.

Answer: Eliminates the need to manually specify all the IPsec security parameters at both peers.

Enables you specify a lifetime for the IPsec SAs.

Enables encryption keys to change during IPsec sessions.

Enables IPsec to provide antireplay services.

Enables CA support for a manageable, scalable IPsec implementation.

Enables dynamic authentication of peers.

11. What is one important difference between SSL and AnyConnect VPNs?

Answer: AnyConnect is a client that lives on the ASA and downloads to your Mac or PC, whereas SSL is a certificate-based VPN hosted by the ASA

Chapter 10

1. How are the terms 802.11 and Wi-Fi used? In what ways are they different or similar?

Answer: These terms describe the IEEE wireless standard and are used interchangeably. Wi-Fi is the buzzword associated with the 802.11 standard.

2. What are the five benefits to organizations that would provide reasons for them to implement a wireless network?

Answer: Attractive price: Deploying a wireless LAN can be cheaper than a wired LAN because you do not need wires; just hook up an access point and it can provide service to multiple computers.

Mobility: Boost user productivity with the convenience of allowing them to wirelessly connect to the network from any point within range of an access point.

Rapid and flexible deployment: Quickly extend a wired network with the ease of attaching an access point to a high-speed network connection.

Application-agnostic: As an extension of the wired network, wireless LANs work with all existing applications.

Performance: Wireless LAN offers a high-speed connection that, although equal to Ethernet, is quickly passing it in speed.

3. Wardriving is the most common means of searching for wireless networks. What is needed to conduct a wardrive, and why is it so useful for attackers?

Answer: Ideally, attackers conducting a wardrive need a program to detect wireless networks such as Net or Mac Stumbler installed on a laptop. They can gain additional information through the use of a GPS device and an antenna.

4. What is one type of freely available wireless packet sniffer?

Answer: Ethereal.

5. Are wireless networks vulnerable to the same types of denial-of-service attacks as wired networks? Are they vulnerable to any additional attacks that wired networks are not?

Answer: Yes, and they are also susceptible to attacks that interfere with radio signals, such as jamming, because wireless networks are based on radio signals.

6. What are the four types of EAP available for use?

Answer: Following are the four commonly used EAP methods in use today: EAP-MD5, EAP-Cisco Wireless (also known as LEAP), EAP-TLS, and EAP-TTLS.

Chapter 11

1. When and who were the first to develop a commercial IDS?

Answer: Late in the 1980s, members of the Haystack Project formed Haystack Labs as a commercial venture into developing host-based intrusion detection.

2. What are the two types of IDSs, and should they be deployed together or separately?

Answer: In general, two basic forms of IDSs are in use today: network-based and host-based IDSs. Both types of sensors offer different techniques for detecting and deferring malicious activity, and both should be deployed in correlation to provide the most effective enhancement to a layered defense strategy.

3. Define and discuss NIDS and how and where they are effective in a network.

Answer: Network-based intrusion detection sensors, or NIDSs, reside directly on the network and watch all traffic traversing the network. NIDSs are effective at both watching for inbound or outbound traffic flows and traffic between hosts on or between local network segments. NIDSs are typically deployed in front of and behind firewalls and VPN gateways to measure the effectiveness of those security devices, and to interact with them to add more depth to the security of your network.

4. Define and discuss HIDSs and how and where they are effective in a network.

Answer: Host-based intrusion detection sensors, or HIDSs, are specialized software applications installed on a computer (typically a server) to watch all inbound and outbound communication traffic to and from that server and monitor the file system for changes. HIDSs are extremely effective on mission-critical, Internet-accessible application servers such as web or email servers because they can watch the applications at the source to protect them.

5. When is anomaly detection the most effective, and why?

Answer: Anomaly detection becomes most effective when coupled with protocol decoding, whereby the IDS knows what normal behavior is expected within certain protocols and responds if abnormal commands or requests are detected.

6. Which intrusion detection methodology also verifies application behavior?

Answer: Protocol analysis.

7. List and define each of the two techniques an IDS can employ to prevent an attack.

Answer: Sniping: Enables the IDS to terminate a suspected attack through the use of a TCP reset packet or ICMP unreachable message.

Shunning: Enables the IDS to automatically configure your prescreening router or firewall to deny traffic based on what it has detected, thus shunning the connection.

8. List the three most important IDS limitations, in your opinion, and explain why you choose them.

Answer: Answer will spur classroom discussion. Some items are

1. Complexity of implementation (HIDS versus NIDS)
 2. Attack patterns and signature updates
 3. False positives
9. True or false: Honey pots distract attackers from more valuable resources.

Answer: True.

Chapter 12

1. What is the difference between a Man-in-the-Middle attack and a denial-of-service attack?

Answer: Essentially, these attacks differ in two ways: maliciousness and results. A *denial-of-service (DoS) attack* occurs when an attacker sends multiple service requests to the victim's computer until they eventually overwhelm the system, causing it to freeze, reboot, and ultimately not be able to carry out regular tasks. A *Man-in-the-Middle (MitM) attack* occurs when intruders inject themselves into an ongoing dialog between two computers so that they can intercept and read messages being passed back and forth.

2. Define what a DDoS is and how it functions. How is that different from a standard DoS attack?

Answer: Quantity of devices sending the attack. Both are DoS attacks and use the same weapons against you (ICMP flood, SYN flood, teardrop attacks, and so on) but the Distributed Denial-of Service (DDoS) attack uses multiple systems to flood the bandwidth or resources of a targeted system, typically focused on one or more web servers. These systems are compromised by attackers using a variety of methods. Typically, a DoS/DDoS attack starts with someone downloading a Trojan onto one system in your network; that Trojan installs an agent that then replicates and installs agents on multiple machines within your network, effectively giving the attacker a botnet within your organization.

3. Name some common DoS attacks.

Answer: ICMP flood (smurf attack, ping flood, and ping of death), SYN flood, and teardrop attacks.

4. Identify and explain three reasons that can result in a back door exploit being present on a system.

Answer:

1. Deliberately placed by system developers to allow quick access during development and not turned off before release.
2. Placed by employees to facilitate performance of their duties because the “proper procedure” made them think that it made their jobs more difficult, so there must be a smarter and easier way. Users might not be as technical as your IT staff, and often they find back doors because they do not have a preconceived notion of how something should work.
3. Normal part of standard default operating system installs that have not been eliminated by OS hardening, such as retaining default user logon ID and password combinations. Again, here you see that vendors do not want technical support calls, so they make it as easy and open as possible. This means that your IT staff must review and harden every server.
4. Placed by disgruntled employees to allow access after termination. In many cases, an employee suspects that he is going to lose his job. This makes him feel angry and unappreciated, so he wants to ensure that he can strike back as needed when the time comes.
5. Created by the execution of malicious code, such as viruses or a Trojan horse that takes advantage of an operating system or application’s vulnerability.

5. Define the concept of firewalking.

Answer: *Firewalking* is a concept and tool that enables the attacker to send specially crafted packets through a firewall to determine what ports and services are permitted through the firewall. Attackers with this knowledge can make their port scans hidden and thus map your network through your firewall.

6. Where should an external penetration and vulnerability assessment be performed in your network?

Answer: External penetration and vulnerability assessments are performed against your network at places where it interacts with the outside world.

7. When considering vulnerability scanners, why are a program’s capability to conduct an accurate scan crucial?

Answer: Scan and detection accuracy. Scans and reported vulnerabilities must be accurate with minimal false positives, defined as normal activity, or a configuration that the system mistakenly reports as malicious. The opposite also holds true, then: There can be no false negatives, defined as malicious activity that is not detected.

This page intentionally left blank

Index

A

- AAA (authentication, authorization, accounting), 156-158
 - accounting, 157-158
 - authentication, 156-157
 - authorization, 157
 - RADIUS (Remote Authentication Dial-In User Service), 158-159
 - TACACS (Terminal Access Control Access Control System), 159-160
- acceptable encryption security policy, 46
- Acceptable Use Policy, 46, 57-64
 - Conclusion section, 63-64
 - Enforcement section, 63
 - General Use and Ownership section, 58-59
 - Overview section, 57-58
 - Purpose section, 58
 - Scope section, 58
 - Security and Proprietary Ownership Information section, 59-60
 - Unacceptable Use section, 60-63
- access
 - controlling, 128
 - hackers, 26-30
 - RBAC (role based access control), 128
- access points, wireless networking
 - association, 319
 - rogue/unauthorized, 316-317
- accounting, 157-158
- ACLs
 - packet filtering, 131-136
 - grocery list analogy*, 132-136
 - static, creating, 224
- acquisition assessment policy, 47
- Acrobat (Adobe), 34
- Active Port Scan Results example, 18-19
- Active reports (CORE IMPACT Pro), 384
- Active X, attacks, 37
- address filtering (MAC), wireless networking, 320-321
- ad-hoc wireless networking, 306
- administrative access, limiting, 111
- Adobe software, attacks, 34
- advisories (security), 86-98
 - Apple, 89
 - awareness, 88
 - Cisco, 89
 - incidents, responses, 90-91

- Microsoft, 89-90
- NIST security documents, 90
- responding to, 87-98
- roles, 91
- AES (Advanced Encryption Standard), 172-173
- aggressive mode (IKE), 274
- Aircrack-ng, 327
- alerts (security), 86-98
- all-in-one firewalls, 204
- analog/ISDN line security policy, 46
- Analogy as a Standard Access List example, 134
- anomaly detection, IDSs (intrusion detection systems), 337, 346-347
- anti-establishment hacking, 3
- antivirus process security policy, 47
- antivirus software, attacks, 33
- AnyConnect VPN Secure Mobility Solution, 295
- AP deployment guidelines, wireless networking, 317-318
- Apple, NSA (National Security Agency) Security Configuration Guides, 121
- Apple security advisories, 89
- application service providers (ASP) standards, 47
- application-level protection, 144-147
- applications, attacks, 27-28
- ARP spoofing, 367-368
- ASAs (Adaptive Security Appliances), VPNs (virtual private networks), 264
- Attack Path reports (CORE IMPACT Pro), 384
- attack patterns, IDSs (intrusion detection systems), 351
- attack signatures, IDSs (intrusion detection systems), 351
- attacks
 - Active X, 37
 - applications, 27-28
 - ARP spoofing, 367-368
 - automated, 27
 - back doors, 368-369
 - botnets, 36
 - brute force, 37
 - compressed files, 37
 - DDoS (Distributed Denial of Service), 36
 - vulnerability analysis*, 365
 - DoS (Denial of Service), 36
 - preventing*, 366-367
 - vulnerability analysis*, 363
 - wireless networks*, 315
 - firewalking, 369-370
 - fraggle, vulnerability analysis, 364
 - Heartland Payment Systems, 50
 - ICMP flood, 38
 - IP spoofing, 37
 - Java, 37
 - Land (C), 37
 - LAND (Local Area Network Denial), 369
 - misconfiguration, 28
 - MitM (man-in-the-middle), 367
 - operating systems, 27
 - origins, 32-33
 - packet analyzers, 363
 - packet sniffing, wireless networks, 313-314
 - phishing, 35
 - ping of death, 367
 - ping pong, 369
 - ping scans, 37
 - port scan, 36
 - process, 9-32
 - covering tracks*, 31-32
 - enumeration*, 23-26
 - escalating privilege*, 30
 - footprinting*, 11-17
 - gaining access*, 26-30

- reconnaissance*, 9-11
- scanning*, 18-23
- rogue access points, wireless networks, 316-317
- scripted, 29-30
- session hijacking, vulnerability analysis, 362-363
- Smurf, 37
 - vulnerability analysis*, 364
- sniffing packets, 39
- source routing, 37-38
- SYN flood, 36
 - vulnerability analysis*, 364-365
- targeted, 27
- teardrop, 37, 365
- TJX Companies, 52
- UDP flood, 36
- unauthorized access points, wireless networking, 316-317
- vulnerability analysis, 361-370
- Xmas tree, 369
- zero day, 36
- attorneys, Internet, 51
- audit vulnerability scanning, 47
- authentication, 156-157
 - EAP (Extensible Authentication Protocol), 321-323
 - IPsec VPNs (virtual private networks), 268-269
 - multi-factor, 161-167
 - OSPF (Open Shortest Path First), 251-254
 - RADIUS (Remote Authentication Dial-In User Service), 158-159
 - two-factor, 161-167
- authorization, 157
- automated attacks, 27
- automatically forwarded email, 47
- awareness
 - security advisories, 88
 - users, 128

B

- back doors, 368-369
- backup software, 100
 - attacks, 34
- bandwidth, as hacking target, 5
- bandwidth availability, wireless networking, 307
- best practices, 98-102
 - change control processes, 98
 - Cisco, 110-118
 - IOS*, 110-111
 - passwords*, 110-111
 - hotfixes, 101
 - security updates, 101-102
 - service packs, 101
- Blade Runner, 85
- Bluetooth device security, 47
- botnets, 5, 36
 - Zeus, 31
- branch design zone guides, 107
- Broderick, Matthew, 307
- browsers, attacks, 36
- brute force attacks, 37
- brute force guess passwords, 29
- buffer memory overflows, 29
- Business Case, Extranet Connection Policy, 75

C

- campus design zone guides, 107-108
- capturing passwords, 29
- Carlin, George, 299
- casing the joint. *See* footprinting
- centralized sensor management, IDSs (intrusion detection systems), 336
- CERT coordination center, 40
- change control processes, 98
- Childs, Terry, 53

choke points, edge routers, 220-224

choke routers, 221-224

**CIS (Center for Internet Security),
40-41**

Cisco

best practices, 110-118

IOS, 110-111

passwords, 110-111

ISE (Identity Services Engine), 166

NSA (National Security Agency)

Security Configuration Guides, 119

**Cisco AnyConnect VPN Secure
Mobility Solution, 295**

Cisco IOS Firewall IDS

FFS IDS, 230-234

intrusion detection, 229-234

Cisco SAFE 2.0, 106

Cisco Secure Consulting Services, 375

Cisco security advisories, 89

Cisco TrustSec, 164-167

**Cisco Validated Design (CVD)
program, 107-110**

Cisco Web Reputation Filters, 155

**client software, VPNs (virtual private
networks), 264**

client-based filtering, 149

clients (email), attacks, 34

**Client-Side Penetration Test reports
(CORE IMPACT Pro), 384**

**Client-Side User reports (CORE
IMPACT Pro), 384**

code listings

Active Port Scan Results (1-2), 18-19

Analogy as a Standard Access List
(5-1), 134

Firewall with Self-Hosted Internal
Web Server (7-2), 209-214

Query Via nstat (1-5), 25

RADIUS Configuration (5-3), 159

Sample Cisco ASA Firewall Rules
(7-1), 199

Secure IOS Template (8-1), 235-250

Standard Access List Filtering Packets
(5-2), 135

TACACS Configuration (5-4), 160

Telnet to Mail Server, Doing Some
Reconnaissance (1-3), 20-21

Using DNS for Passive

Reconnaissance via dig Command
(1-1), 13-14

Using nbtstat -c to Display NetBIOS
Names (1-6), 25

Using Windows Net View (1-4), 24

commands, dig, 13-14

common security policies, 48-49

**common vulnerabilities and exposures
(CVE), 39-40**

compressed files, attacks, 37

compromised confidential data, 196

**Computer Crime and Intellectual
Property website, 50**

Conclusion section

Acceptable Use Policy, 63-64

Extranet Connection Policy, 76-77

Password Policy, 68-69

Virtual Private Network (VPN)
Security Policy, 71

confidential data, compromised, 196

configuration

IPsec, 284-286

ISAKMP (Internet Security
Association Key Management
Protocol), 281-283

perimeter routers, 220

routers, as VPN peers, 281-286

VPNs (virtual private networks),
286-289

content filtering, 147-150

limitations, 150

controlling access, 128

CORE IMPACT, 30

CORE IMPACT Pro, 382-386

documentation, 386

reports, 384-385

vulnerability updates, 386

corporate policies, 53-57coverage, wireless networking,
306-307**curiosity, hacking, 3**CVD (Cisco Validated Design)
program, 107-110CVE (common vulnerabilities and
exposures), 39-40**cyberwarfare, 4****D**Data Center Design Center guides,
108-109data integrity, IPsec VPNs (virtual
private networks), 268-269

database credentials coding, 47

database software, attacks, 34

DDoS (Distributed Denial of Service)
attacks, 36

vulnerability analysis, 365

deception systems, honeypots, 355

Definitions section, Wireless
Communication Policy, 73Definitions section (security policy),
56

delivering, security policies, 77-78

Delta reports (CORE IMPACT Pro),
385Demilitarized Zone (DMZ), firewalls,
206-214Denial of Service (DoS). *See* DoS
(Denial of Service) attacksdeployment, VPNs (virtual private
networks), 270-271**design concepts**

controlling access, 128

incident response teams, 130-131

layered security, 128

monitoring, 129

RBAC (role based access control), 128

user awareness, 128

design strategies, honeypots, 356-357

detailed packet flow, SPI (Stateful
Packet Inspection), 138-139detection software, HIDS (host-based
intrusion detection systems), 341

dial-in access policies, 47

Diffie-Hellman algorithm, IPsec,
279-280dig command, DNS passive reconnais-
sance, 13-14

disaster recovery, 374

Distributed Denial of Service (DDoS)
attacks. *See* DDoS (Distributed
Denial of Service) attacksDMZ (Demilitarized Zone), firewalls,
47, 206-214DNS (Domain Name System) attacks,
35passive reconnaissance, dig command,
13-14**documentation**

CORE IMPACT Pro, 384-386

security scanners, 379

DoS (Denial of Service) attacks, 6, 36

IDSs (intrusion detection systems),
353

preventing, 366-367

vulnerability analysis, 363

wireless networking, 315

downstream liability, 195-196

downtime

backups, 100

networks, 196

dynamic NAT, 142
dynamic proxy firewalls, 145

E

EAP (Extensible Authentication Protocol), 321-323
EAP-PSK, 323
EAP-TLS, 322-323
EAP-TTLS, 323
eavesdropping, wireless networks, 313-314
echo reply (ICMP) attacks, 38
economic motivations, hacking, 4
edge routers
 as a choke point, 220-224
 configuring, 220
 as a packet inspector, 220
Email and Communications Activities subsection, Acceptable Use Policy, 62-63
email clients, attacks, 34
E-mail Retention policy, 47
employee information, 6
encryption
 AES (Advanced Encryption Standard), 172-173
 Triple DES, 171-172
encryption modes, IPsec, 271-272
Enforcement section, 56
 Acceptable Use Policy, 63
 Password Policy, 68
 Wireless Communication Policy, 73
enterprise firewalls, 204
enumeration, 23-26
escalating privilege, 30
Establishing Connectivity section, Extranet Connection Policy, 75
ETTERCAP, 29
event correlation, IDSs (intrusion detection systems), 336

examples

Active Port Scan Results (1-2), 18-19
Analogy as a Standard Access List (5.1), 134
Firewall with Self-Hosted Internal Web Server (7-2), 209-214
Query Via nbstat (1-5), 25
RADIUS Configuration (5-3), 159
Sample Cisco ASA Firewall Rules (7-1), 199
Secure IOS Template (8-1), 235-250
Standard Access List Filtering Packets (5-2), 135
TACACS Configuration (5-4), 160
Telnet to Mail Server, Doing Some Reconnaissance (1-3), 20-21
Using DNS for Passive Reconnaissance via dig Command (1-1), 13-14
Using nbtstat -c to Display Net BIOS Names (1-6), 25
Using Windows Net View (1.4), 24
excessive user rights, 34
Executive Summary reports (CORE IMPACT Pro), 385
Extensible Authentication Protocol (EAP), 321-323
external vulnerability analysis, 371-373
Extranet Connection Policy, 74-77
 Business Case, 75
 Conclusion section, 76-77
 Establishing Connectivity section, 75
 Modifying or Changing Connectivity and Access section, 76
 Point of Contact (POC), 75
 Purpose section, 74
 Scope section, 74-75
 Security Review, 75
 Terminating Access section, 76
 Third-Party Connection Agreement, 75

extranet VPNs, 262
extranets, security policies, 47

F

false negatives, IDSs (intrusion detection systems), 352
false positives, IDSs (intrusion detection systems), 336, 352
fame, hacking, 3
FFS IDS, 230-234
filtering
 malware, 201
 packets, 134-136
 ACLs, 131-136
 reactive, 154-155
 traffic, 149
filtering network traffic, firewalls, 201
filters
 Cisco Web Reputation, 155
 content, 147-150
firewalking, 369-370
Firewall/ASAs, 115-118
Firewall with Self-Hosted Internal Web Server example, 209-214
firewalls, 193-194, 215, 219
 all-in-one, 204
 benefits, 195
 Cisco IOS Firewall IDS, intrusion detection, 229-234
 DMZ (Demilitarized Zone), 206-214
 downstream liability, 195-196
 enterprise, 204
 filtering network traffic, 201
 functions, 196-197
 implementing, 203-205
 inbound access policies, 205-206
 limitations, 214-215
 lost data, 196
 operations, 200-206
 outbound access policies, 206

 personal, 203-204
 proxies, 145
 security policies, 200
 security policy, 197-200
 SPI (Stateful Packet Inspection), 139
 VPNs (virtual private networks), 264
 zone-based, routers, 224-229
FISMA Vulnerability Validation reports (CORE IMPACT Pro), 385
Flash (Adobe), attacks, 34
footprinting, 11-17
 goals, 12-13
fraggle attacks, vulnerability analysis, 364
fragmentation, IDSs (intrusion detection systems), 353
freeware security scanners, 376-382
functionality, PPTP (Point-to-Point Tunneling Protocol), 177-178

G

General Network Access
 Requirements subsection, Wireless Communication Policy, 72-73
General Password Construction
 Guidelines, Password Policy, 66-67
General Policy section, Password Policy, 65-66
General Use and Ownership section, Acceptable Use Policy, 58-59
GFI LANGuard, 29
grocery list analogy, packet filtering via *ACLs*, 132-136

H

hackers, 1-2
 hactivism, 4
 script kiddies, 3, 6
 stereotypes, 7
 tasks, 26

hacking

- attacks, process, 9-32
- motivations, 3-4
- targets, 2-3
 - choice*, 7-8
 - opportunity*, 4-7
- hacking tools, wireless, 325-329
- hactivism, 4
- Hammersley, Ben, 308
- hard drive space, as hacking target, 5
- Hawking, Stephen, 331
- HaXor, 359
- header condition signatures, NIDS (network-based intrusion systems), 342
- Health Insurance Portability and Accounting Act (HIPAA) of 1996, 50, 81
- Heartland Payment Systems, attack on, 50
- helpdesk, forewarning, 100
- HIDS (host-based intrusion detection systems), 337-341
 - detection software, 341
 - versus NIDS, 350-351
- HIPAA (Health Insurance Portability and Accounting Act) of 1996, 50, 81
- History section (security policy), 56
- Home Wireless Device Requirements subsection, Wireless Communication Policy, 73
- honeypots, 331-333, 354-357
 - deception systems, 355
 - design strategies, 356-357
 - limitations, 357
 - multideception systems, 356
 - port monitoring, 355
 - production, 355
 - research, 355
- Host reports (CORE IMPACT Pro), 385

- host unreachable (ICMP) attacks, 38
- host-based IDSs (intrusion detection systems), 337-341

hotfixes, 97

- best practices, 101
- uninstalling, 99

ICMP flood attacks, 38**identity theft, 4**

- IDSs (intrusion detection systems), 331-333, 335-346. *See also* HIDS (host-based intrusion detection systems); NIDS (network-based intrusion systems)
 - anomaly detection, 337, 346-347
 - attack patterns, 351
 - attack signatures, 351
 - centralized sensor management, 336
 - combining methods, 347
 - DoS (Denial of Service) attacks, 353
 - elimination of false positives, 336
 - event correlation, 336
 - fragmentation, 353
 - host-based, 337-341
 - intrusion prevention, 347-348
 - limitations, 350-353
 - NBA (network behavior analysis), 338-339
 - network-based, 338-339, 341-343
 - origins, 335
 - pattern detection, 346
 - products, 348-350
 - signature detection, 346
 - signatures, 336
 - matching*, 337
 - standards-based implementation, 336
 - stateful protocol analysis, 347
 - thresholds, 336
 - wireless, 338-339, 343-344

- IEEE 802.1x, 162-164
- IKE (Internet Key Exchange), VPNs (virtual private networks), 274-275
- IM (instant messaging), attacks, 34
- implementation, firewalls, 203-205
- implementing, VPNs (virtual private networks), 264-265
- inbound access policies, firewalls, 205-206
- inbound telnet, limiting access, 112-113
- incident response teams, 130-131
- incidents, defining, 92
- industry best practices, 98-102
 - change control processes, 98
- industry standards, security policies, 79-82
- Information Asset Sensitivity, 48
- Information System Audit Logging, 48
- Information Technology Law (IT Law), 51
- infrastructure, wireless networking, 305
- Inge, William Ralph, 359
- inline wiretap, NIDS (network-based intrusion systems), 341
- instant messaging, attacks, 34
- intercepting data, wireless networks, 313-314
- Internal Lab security, 48
- internal vulnerability assessment, 370-371
- Internet lawyers, 51
- Internet Security Association Key Management Protocol (ISAKMP), 272-273
 - configuring, 281-283
- Internet Storm Center, 41
- Internet usage policies, 48
- intrusion detection, 331-345. *See also*
 - IDSs (intrusion detection systems)
 - Cisco IOS Firewall IDS, 229-234
 - IDSs (intrusion detection systems), 331-333, 336-339, 345-346
 - methods*, 345-353
 - signature detection*, 346
 - NBA (network behavior analysis), 344-345
 - wireless, 343-344
 - intrusion detection systems (IDSs). *See* IDSs (intrusion detection systems), 347-348
- IOS best practices, Cisco, 110-111
- IP spoofing, 37, 362-363
- IPsec, 276-280
 - Configuring, 282-286
 - Diffie-Hellman algorithm, 279-280
 - encryption modes, 271-272
 - IKE Phase 1, 277-278
 - IKE Phase 2, 278
 - PFS (perfect forward secrecy), 278-279
 - protocols, 272-273
 - SAs (security associations), 275-276
 - transforms, 284-285
 - tunneling data, 269-270
 - VPNs (virtual private networks), 257-259, 265-267, 271-273
 - authentication*, 268-269
 - configuring routers*, 281-286
 - data integrity*, 268-269
 - IKE (Internet Key Exchange)*, 274-275
 - security considerations*, 293-295
 - versus SSL VPNs*, 290-293
 - zone-based policy firewalls, 224-225
- ISAKMP (Internet Security Association Key Management Protocol), 272-273
 - configuring, 281-283
 - presared keys, 282

ISE (Identity Services Engine), 166
ISO certification, 77-79
ISO/IEC 27002 information security
standard, 78-79
IT professionals, arrogance, 2-3

J-K

Java, attacks, 37
Jones, Matt, 308

keystroke loggers, 31

L

L2TP (Layer 2 Tunneling Protocol),
179-182
Lab and Isolated Wireless Device
Requirements subsection, Wireless
Communication Policy, 72
Land (C) attacks, 37
LAND (Local Area Network Denial)
attacks, 369
LANs (local area networks), 301-304
LAND (Local Area Network Denial),
369
WLANs (wireless LANs), 301-304
benefits, 303
radio frequency, 303-304
standard characteristics, 301-302
Wi-Fi (Wireless Fidelity),
302-303
large ICMP packet attacks, 38
lawyers, Internet, 51
Layer 2 Tunneling Protocol (L2TP),
179-182
layered security, 128
VPNs (virtual private networks),
270-271
LEAP (Lightweight Extensible
Authentication Protocol), 321-322
legal precedences, 50-51

Lightweight Extensible Authentication
Protocol (LEAP), 322
limitations, honeypots, 357
line access controls, limiting, 111
long-term states, IDSs (intrusion
detection systems), 352
loose route attacks, 38
lost data, firewalls, 196

M

MAC address filtering, 320-321
main mode (IKE), 274
malicious web pages, 147
malware, filtering, 201
man-in-the-middle attacks, 367
Massachusetts 201: Standards for the
Protection of Personal Information
of Residents of the Commonwealth,
81-82
matching signatures, IDSs
(intrusion detection systems), 337
MD5 (Message Digest 5)
algorithm, 173-175
MD5 route authentication, OSPF
(Open Shortest Path First), 253-254
media players, attacks, 33
Message Digest 5 algorithm, 173-175
Metasploit Framework, 376
METASPLOIT PRO, 30
Microsoft, security, 121-125
Microsoft KB Articles, 98
Microsoft security bulletins, 89-90
Microsoft Security Compliance
Manager, 124-125
Microsoft Windows, NSA (National
Security Agency) Security
Configuration Guides, 119-121
misconfiguration attacks, 28
MitM (man-in-the-middle) attacks, 367
modes of operation, wireless
networking, 305-306

Modifying or Changing Connectivity and Access section, Extranet Connection policy, 76
 motivations, hacking, 3-4
 multideception systems, honeypots, 356
 multi-factor authentication, 161-167

N

NAC (Network Access Control), 162-164
NAT (Network Address Translation), 140-144, 205-206
 dynamic, 142
 limitations, 143-144
 overloading, 142
 static, 142
National Institute of Standards and Technology (NIST), 258
National Vulnerability Database (NVD), 41
NBA (network behavior analysis), 338-339, 344-345
nbstat command, 25
neighbor authentication, OSPF (Open Shortest Path First), 252
Nessus, 29, 377
Ncat, 31
NetStumbler, 325-326
Network Access Control (NAC), 162-164
Network Address Translation (NAT). See NAT (Network Address Translation)
network security organizations, 39-42
network security standards, 105
 Cisco SAFE 2.0, 106
 CVD (Cisco Validated Design) program, 107-110
network traffic, filtering, firewalls, 201
network-based IDSs (intrusion detection systems), 338-339, 341-343

networks, downtime, 196
NIDS (network-based intrusion systems), 338-339, 341-343
 header condition signatures, 342
 versus HIDS, 350-351
 inline wiretap, 341
 port mirroring, 342
 port signatures, 342
 string signatures, 342
NIST (National Institute and Technology), 258
NIST security documents, 90
NMAP (Network Mapper), 29, 376-377
NSA (National Security Agency) Security Configuration Guides, 118-121
 Apple, 121
 Cisco Systems, 119
 Microsoft Windows, 119-121
NVD (National Vulnerability Database), 41

O

office firewalls, 204
office software, attacks, 34
OmniPeek, 327-329
operating systems, attacks, 27
operations, SSH (Secure Shell), 186-187
organizations, responsibilities and expectations, 50-53
origins, attacks, 32-33
OSPF (Open Shortest Path First)
 authentication, 251-254
 MD5 route authentication, 253-254
 plaintext route authentication, 253
outbound access policies, firewalls, 206
outbound telnet, limiting access, 112-113

overloading NAT, 142

Overview section, 56

Acceptable Use Policy, 57-58

Password Policy, 64

P

P2P (peer-to-peer), attacks, 35

packet analyzers, 363

packet filtering

ACLs, 131-136

grocery list analogy, 132-136

Layer 3, 131

limitations, 136

packet filters, placement, 135

packet flow, proxies, 144

packet inspector, edge routers, 220

packet sniffers, wireless, 326-327

packet sniffing, wireless, 313-314

packets

sniffing, 39

SPI (Stateful Packet Inspection),
136-140

parameter problem on datagram

(ICMP) attacks, 38

Password Policy, 64-69

Conclusion section, 68-69

Enforcement section, 68

General Password Construction
Guidelines, 66-67

General Policy section, 65-66

Overview section, 64

Password Protection Standards, 67-68

Purpose section, 64

Scope section, 64-65

Password Protection Standards,

Password Policy, 67-68

passwords

brute force guess, 29

capturing, 29

policies, 48

securing, 110-111

try and sniff, 29

PAT (Port Address Translation), 142

patches, uninstalling, 99

pattern detection, IDSs (intrusion
detection systems), 346

pattern evasion, IDSs (intrusion
detection systems), 353

Payment Card Industry Data Security
Standard (PCI DSS), 80

PCI DSS (Payment Card Industry
Data Security Standard), 80

PCI Vulnerability Validation reports
(CORE IMPACT Pro), 385

peer-to-peer (P2P), attacks, 35

penetration assessment, 370-373

penetration testing, 370-375

perfect forward secrecy (PFS),
IPsec, 278-279

perimeter routers, configuring, 220

personal communication devices,
policies, 48

personal employee information, as
hacking target, 6

personal firewalls, 203-204

PFS (perfect forward secrecy),
IPsec, 278-279

phishing, 35

spear phishing, 35

vishing, 35

whaling, 35

physical security assessment, 373-374

ping of death, 367

ping pong attacks, 369

ping scans, 37

PKI (Public Key Infrastructure)
encryption, 150-152

plaintext route authentication, OSPF
(Open Shortest Path First), 253

Point of Contact (POC), Extranet
Connection Policy, 75

Point-to-Point Tunneling Protocol (PPTP), 177-179**policies (security), 45-49**

- Acceptable Use Policy, 46, 57-64
 - Conclusion section, 63-64*
 - Enforcement section, 63*
 - General Use and Ownership section, 58-59*
 - Overview section, 57-58*
 - Purpose section, 58*
 - Scope section, 58*
 - Security and Proprietary Ownership Information section, 59-60*
 - Unacceptable Use section, 60-63*
- common, 48-49
- corporate, 53-57
- Definitions section, 56
- delivering, 77-78
- Enforcement section, 56
- Extranet Connection Policy, 74-77
 - Business Case, 75*
 - Conclusion section, 76-77*
 - Establishing Connectivity section, 75*
 - Modifying or Changing Connectivity and Access section, 76*
 - Point of Contact (POC), 75*
 - Purpose section, 74*
 - Scope section, 74-75*
 - Security Review, 75*
 - Terminating Access section, 76*
 - Third-Party Connection Agreement, 75*
- firewalls, 197-200
- History section, 56
- industry standards, 79-82
- ISO certification, 77-79
- Microsoft, 121-124

Overview section, 56

Password Policy, 64-69

- Conclusion section, 68-69*
- Enforcement section, 68*
- General Password Construction Guidelines, 66-67*
- General Policy section, 65-66*
- Overview section, 64*
- Password Protection Standards, 67-68*
- Purpose section, 64*
- Scope section, 64-65*
- Policy section, 56*
- Purpose section, 56*

RBAC (role based access control), 55
relevant, 54

Revision section, 56

samples, 79

Scope section, 56

SLAs (service-level agreements), 45

Virtual Private Network (VPN)

- Security Policy, 31, 69-71
- Conclusion section, 71*
- Policy section, 70-71*
- Purpose section, 69*
- Scope section, 69*

Wireless Communication Policy, 71-74

- Definitions section, 73*
- Enforcement section, 73*
- Policy Statement, 72-73*
- Revision History, 73*
- Scope section, 72*

Policy section

- Virtual Private Network (VPN)
- Security Policy, 70-71

Policy Statement, Wireless**Communication Policy, 72-73****port forwarding, SSH (Secure Shell), 187-188**

- port mirroring, NIDS (network-based intrusion systems), 342
- port monitoring, honeypots, 355
- port scan attacks, 36
- port signatures, NIDS (network-based intrusion systems), 342
- PPTP (Point-to-Point Tunneling Protocol), 177-179
 - functionality, 177-178
 - limitations, 178-179
- presheared keys, ISAKMP (Internet Security Association Key Management Protocol), 282
- privileges, escalating, 30
- procedural risk assessment, 374
- procedures, 85-86
 - establishing, 94
- processes, 85-86, 102-103
 - attacks, 9-32
 - covering tracks*, 31-32
 - enumeration*, 23-26
 - escalating privilege*, 30
 - footprinting*, 11-17
 - gaining access*, 26-30
 - reconnaissance*, 9-11
 - scanning*, 18-23
 - change control, 98
- production honeypots, 355
- protocols
 - authentication, EAP (Extensible Authentication Protocol), 321-323
 - IPsec, 272-273
 - Message Digest 5 algorithm, 173-175
 - routing, security, 251-254
 - security, 169-171, 192
 - AES (Advanced Encryption Standard)*, 172-173
 - L2TP (Layer 2 Tunneling Protocol)*, 179-182
 - PPTP (Point-to-Point Tunneling Protocol)*, 177-179

- SHA (Secure Hash Algorithm)*, 175-177
- SNMP v3 (Simple Network Management Protocol Version 3)*, 188-191
- SSH (Secure Shell)*, 182-188
- Triple DES*, 171-172
- proxies, 144-147
 - firewalls, 145
 - limitations, 146-147
 - packet flow, 144
- Public Key Infrastructure (PKI) encryption, 150-152
- public libraries, content filtering, 147
- Purpose section, 56
 - Acceptable Use Policy, 58
 - Extranet Connection Policy, 74
 - Password Policy, 64
 - Virtual Private Network (VPN) Security Policy, 69

Q-R

- Qoncert, 375
- Query Via nbstat example, 25
- radio frequency, 303-304
- RADIUS (Remote Authentication Dial-In User Service), 158-159
 - versus TACACS, 160
- RADIUS Configuration example, 159
- RBAC (role based access control), 28, 55, 128
- reactive filtering, 154-155
- Reader (Adobe), attacks, 34
- Reagan, Ronald, 127
- reconnaissance, 9-11
 - goals, 12-13
- record route attacks, 37
- redirect (ICMP) attacks, 38
- relevant security policies, 54

remote access policies, 48
 remote access VPNs, 259-261
 removable media
 attacks, 35
 policies, 48
 reporting, security scanners, 379
 reports, CORE IMPACT Pro, 384-385
 reputation-based security, 152-156
 research honeypots, 355
 resource limitations, IDSs (intrusion detection systems), 352
 responding to security advisories, 87-98
 awareness, 88
 responsibilities and expectations, organizations, 50-53
 Retina version 5.1.1.10, 380
 Revision History, Wireless Communication Policy, 73
 Revision section (security policy), 56
 risk assessment, 48
 risks, common, 33-36
 rogue access points, wireless networks, 316-317
 role based access control (RBAC), 28, 55, 128
 roles, establishing, 91
 routers, 217-220, 254-255
 configuring, IPsec VPNs (virtual private networks), 281-286
 edge
 as a choke point, 220-224
 as a packet inspector, 220
 perimeter, configuring, 220
 policies, 49
 VPNs (virtual private networks), 264
 zone-based firewalls, 224-229
 routing protocols, security, OSPF (Open Shortest Path First), 251-252

S

SAFE (Cisco) 2.0, 106
 SAINT scanner, 29, 377
 Sample Cisco ASA Firewall Rules, 199
 SANS (SysAdmin, Audit, Network, Security) Institute, 40
 Sarbanes-Oxley Act of 2002, 80-81
 SAs (security associations)
 IPsec, 275-276
 VPNs (virtual private networks), 273
 SAS (Statement on Auditing Standards) series, 82
 SATAN, 29
 scanning, 18-23
 scanners (security), 375-382
 documentation, 379
 reporting, 379
 scan and detection accuracy, 378
 vulnerability updates, 379-380
 scheduled downtime, backups, 100
 Scope section, 56
 Acceptable Use Policy, 58
 Extranet Connection Policy, 74-75
 Password Policy, 64-65
 Virtual Private Network (VPN) Security Policy, 69
 Wireless Communication Policy, 72
 SCORE, 41
 script kiddies, 3, 6
 script source route attacks, 38
 scripted attacks, 29-30
 Secure Consulting Services (Cisco), 375
 Secure Hash Algorithm (SHA), 175-177
 Secure IOS template, 234-250
 Secure Shell (SSH). *See* SSH (Secure Shell).

security

- advisories, 86-98
 - Apple*, 89
 - awareness*, 88
 - Cisco*, 89
 - incidents, responses*, 90-91
 - Microsoft*, 89-90
 - NIST security documents*, 90
 - responding to*, 87-98
 - roles*, 91
- alerts, 86-98
- Microsoft, 121-125
- wireless networking, 329-330
- Security and Proprietary Ownership Information section, Acceptable Use Policy**, 59-60
- security assessments**, 370-375
- security associations (SAs)**
 - IPsec, 275-276
 - VPNs (virtual private networks), 273
- Security Compliance Manager (Microsoft)**, 124-125
- Security Configuration Guides (NSA)**, 118-121
 - Apple, 121
 - Cisco Systems, 119
 - Microsoft Windows, 119-121
- security design zone guides**, 109-110
- security patches, uninstalling**, 99
 - Acceptable Use Policy, 46, 57-64
 - Conclusion section*, 63-64
 - Enforcement section*, 63
 - General Use and Ownership section*, 58-59
 - Overview section*, 57-58
 - Purpose section*, 58
 - Scope section*, 58
 - Security and Proprietary Ownership Information section*, 59-60
 - Unacceptable Use section*, 60-63
- common, 48-49
- corporate, 53-57
- Definitions section, 56
- delivering, 77-78
- Enforcement section, 56
- Extranet Connection Policy, 74-77
 - Business Case*, 75
 - Conclusion section*, 76-77
 - Establishing Connectivity section*, 75
 - Modifying or Changing Connectivity and Access section*, 76
 - Point of Contact (POC)*, 75
 - Purpose section*, 74
 - Scope section*, 74-75
 - Security Review*, 75
 - Terminating Access section*, 76
 - Third-Party Connection Agreement*, 75
- firewalls, 197-200
- History section, 56
- industry standards, 79-82
- ISO certification, 77-79
- Microsoft, 121-124
- Overview section, 56
- Password Policy, 64-69
 - Conclusion section*, 68-69
 - Enforcement section*, 68
 - General Password Construction Guidelines*, 66-67
 - General Policy section*, 65-66
 - Overview section*, 64
 - Password Protection Standards*, 67-68
 - Purpose section*, 64
 - Scope section*, 64-65
 - Policy section*, 56
 - Purpose section*, 56
- RBAC (role based access control), 55

- relevant, 54
- Revision section, 56
- samples, 79
- Scope section, 56
- SLAs (service-level agreements), 45
- Virtual Private Network (VPN)
 - Security Policy, 31, 69-71
 - Conclusion section*, 71
 - Policy section*, 70-71
 - Purpose section*, 69
 - Scope section*, 69
- Wireless Communication Policy, 71-74
 - Definitions section*, 73
 - Enforcement section*, 73
 - Policy Statement*, 72-73
 - Revision History*, 73
 - Scope section*, 72
- security protocols, 169-171, 192**
 - AES (Advanced Encryption Standard), 172-173
 - L2TP (Layer 2 Tunneling Protocol), 179-182
 - Message Digest 5 algorithm, 173-175
 - PPTP (Point-to-Point Tunneling Protocol), 177-179
 - SHA (Secure Hash Algorithm), 175-177
 - SNMP v3 (Simple Network Management Protocol Version 3), 188-191
 - SSH (Secure Shell), 182-188
 - Triple DES, 171-172
- Security Review, Extranet Connection Policy, 75**
- security scanners, 375-382**
 - documentation, 379
 - reporting, 379
 - scan and detection accuracy, 378
 - vulnerability updates, 379-380
- security updates, 97**
 - applying, 99
 - best practices, 101-102
- SecurityFocus, 42**
- sensor blindness, IDSs (intrusion detection systems), 352**
- server-based filtering, 149**
- servers**
 - as hacking targets, 5
 - policies, 49
- service packs, 97**
 - best practices, 101
 - keeping up with, 100
 - uninstalling, 99
- Service Set Identifier (SSID). *See* SSID (Service Set Identifier).**
- service-level agreements (SLAs), 45**
- session hijacking, 362-363**
- session timeouts, establishing, 113**
- SHA (Secure Hash Algorithm), 175-177**
- signature detection, IDSs (intrusion detection systems), 346**
- signatures**
 - IDSs (intrusion detection systems), 336
 - matching*, 337
- site-to-site VPNs, 258, 261-262**
- SLAs (service-level agreements), 45**
- small-to-medium office firewalls, 204**
- Smurf attacks, 37**
 - vulnerability analysis, 364
- sniffing packets, 39**
 - wireless networks, 313-314
- SNMP v3 (Simple Network Management Protocol Version 3), 188-191**
- Snort IDS/IPS, 348-350**
- social messaging, 34**
- source quench (ICMP) attacks, 38**
- source routing attacks, 37-38**

spam, content filtering, 147
 spamming, warspamming, 311-312
 spear phishing, 35
 SPI (Stateful Packet Inspection),
 136-140
 detailed packet flow, 138-139
 firewalls, 139
 limitations, 139-140
 split tunneling, VPNs (virtual
 private networks), 265
 SSH (Secure Shell), 182-188
 Limitations, 188
 operation, 186-187
 port forwarding, 187-188
 versus Telnet, 184-185
 tunneling, 187-188
 SSID (Service Set Identifier), 310
 wireless networks, 318
 SSL (Secure Sockets Layer)
 attacks, 35
 VPNs (virtual private networks),
 290-293
 security considerations, 293-295
 Standard Access List Filtering Packets
 example, 135
 standard proxy firewalls, 145
 Stateful Packet Inspection (ISP). *See*
 SPI (Stateful Packet Inspection)
 stateful protocol analysis, IDSs
 (intrusion detection systems), 347
 Statement on Auditing Standards
 (SAS) 70 series, 82
 static ACLs, creating, 224
 static NAT (Network Address
 Translation), 142
 stereotypes, hackers, 7
 storage limitations, IDSs (intrusion
 detection systems), 352
 string signatures, NIDS (network-based
 intrusion systems), 342
 switches, policies, 49

SYN flood attacks, 36
 vulnerability analysis, 364-365
 System and Network Activities
 subsection, Acceptable Use Policy,
 61-62
 System Message Logging (syslog), 229

T

TACACS (Terminal Access Control
 Access Control System), 159-160
 versus RADIUS, 160
 TACACS Configuration example, 160
 TACACS+, 112
 targets, hacking, 2-3, 27
 choice, 7-8
 opportunity, 4-7
 TCP wrappers, 341
 teardrop attacks, 37, 365
 Telnet, 20-21, 184-185
 versus SSH (Secure Shell), 184-185
 Telnet to Mail Server, Doing Some
 Reconnaissance example, 20-21
 telnetting, 22
 templates, Secure IOS, 234-250
 Terminating Access section, Extranet
 Connection Policy, 76
 testing, 99
 Third-Party Connection Agreement,
 Extranet Connection Policy, 75
 threat agents, 86
 threats. *See also* attacks
 common, 33-36
 DoS (Denial of Service), wireless
 networks, 315
 packet sniffing, wireless networking,
 313-314
 rogue access points, wireless
 networks, 316-317
 unauthorized access points, wireless
 networks, 316-317

vulnerability analysis, 361-370
 wireless networking, 312-321
 thresholds, IDSs (intrusion detection systems), 336
 time exceed for a datagram (ICMP) attacks, 38
 timeouts, establishing, 113
 TJX Companies, attack on, 52
 traffic filtering, 149
 transforms, IPsec, 284-285
 Transport mode (IPsec), 272
 trapdoors, 368-369
 Trend reports (CORE IMPACT Pro), 385
 Triple DES encryption, 171-172
 Trojan horses, 147
 TrustSec (Cisco), 164-167
 try and sniff passwords, 29
 Tunnel mode (IPsec), 271
 tunneling, SSH (Secure Shell), 187-188
 tunneling data, IPsec VPNs (virtual private networks), 269-270
 Twain, Mark, 193
 two-factor authentication, 161-167

U

UDP flood attacks, 36
 Unacceptable Use section, Acceptable Use Policy, 60-63
 uninstalling service packs, 99
 University of East Anglia, hacking scandal, 4
 updates (security), 97
 applying, 99
 URL-filtering, 154
 user awareness education, 54-55
 user rights, excessive, 34
 users, awareness, 128

Using DNS for Passive Reconnaissance via dig Command example, 13-14

Using nbtstat -c to Display NetBIOS Names example, 25

Using Windows Net View example, 24

V

Virtual Private Network (VPN) Security Policy, 31, 69-71

 Conclusion section, 71

 Policy section, 70-71

 Purpose section, 69

 Scope section, 69

viruses, 147

vishing, 35

VNC (Virtual Network Computing), 31

vos Savant, Marlene, 45

VPN (Virtual Private Network) Security Policy. *See* Virtual Private Network (VPN) Security Policy

VPN peers, routers, configuring, 281-286

VPNs (virtual private networks), 261-265, 296

 ASAs (Adaptive Security Appliances), 264

 benefits, 263-264

 client software, 264

 configuring, firewalls, 286-289

 deployment, 270-271

 extranet, 262

 firewalls, 264

 goals, 263-264

 implementation strategies, 264-265

 IPsec VPNs (virtual private networks), 257-259, 265-267

authentication, 268-269

data integrity, 268-269

- encryption modes, 271-272*
 - IKE (Internet Key Exchange), 274-275*
 - tunneling data, 269-270*
- layered security, 270-271
- remote access, 259-261
- routers, 264
- SAs (security associations), 273
- security policies, 49
- site-to-site, 258-262
- split tunneling, 265
- SSL (Secure Sockets Layer), 289-290
- vulnerability analysis, 361-370**
 - ARP spoofing, 367-368
 - back doors, 368-369
 - common, 33-36
 - CORE IMPACT Pro, 382-386
 - DDoS (Distributed Denial of Service) attacks, 365
 - disaster recovery, 374
 - DoS (Denial of Service) attacks, 363
 - external vulnerability, 371-373
 - firewalking, 369-370
 - fraggle attacks, 364
 - information handling security assessment, 375
 - internal vulnerability, 370-371
 - IP spoofing, 362-363
 - LAND (Local Area Network Denial), 369
 - MitM (man-in-the-middle), 367
 - packet analyzers, 363
 - penetration assessment, 370-371
 - penetration testing, 370-375
 - physical security assessment, 373-374
 - ping of death, 367
 - ping pong attacks, 369
 - procedural risk assessment, 374
 - security assessments, 370-375

- security scanners, 375-382
- session hijacking, 362-363
- Smurf attacks, 364
- SYN flood attacks, 364-365
- teardrop attacks, 365
- Xmas tree attacks, 369
- Vulnerability reports (CORE IMPACT Pro), 385**
- vulnerability scanners, 375-382**
 - scan and detection accuracy, 378
- vulnerability updates, CORE IMPACT Pro, 386**

W

- WAN design zone guides, 107**
- WAP (wireless access points), 304**
- warchalking, wireless networking, 307-309**
- wardriving, wireless networking, 309-311**
- WarGames, 307-308**
- warspamming, 311-312**
- warspying, 312**
- Web Application Vulnerability reports (CORE IMPACT Pro), 385**
- web browsers, attacks, 36**
- WEP (Wired Equivalent Privacy), 319-320**
- whaling (phishing), 35**
- whois tool, 16**
- Wi-Fi (Wireless Fidelity), 302-303**
- Wiki-Leak, 4**
- Windows, enumerating, 24-26**
- Windows 7, security policies, 122-123**
- Windows Server 2003, security policies, 122**
- Windows Server 2008, security policies, 123**
- Windows XP Professional, security policies, 122**

Wired Equivalent Privacy (WEP),
 wireless networks, 319-320
wireless access point (WAP), 304
Wireless Communication Policy, 71-74
 Definitions section, 73
 Enforcement section, 73
 Policy Statement, 72-73
 Revision History, 73
 Scope section, 72
wireless hacking tools, 325-329
wireless IDSs (intrusion detection systems), 338-339, 343-344
wireless networking
 access points, association, 319
 ad-hoc, 306
 AP deployment guidelines, 317-318
 bandwidth availability, 307
 coverage, 306-307
 device associations, 319
 EAP (Extensible Authentication Protocol), 321-323
 infrastructure, 305
 MAC address filtering, 320-321
 modes of operation, 305-306
 security, 304-307, 323-324, 329-330
 SSID (Service Set Identifier), 318
 threats, 312-321
 DoS (Denial of Service) attacks,
 315
 packet sniffing, 313-314
 rogue/unauthorized access
 points, 316-317
 warchalking, 307-309
 wardriving, 309-311
 warspamming, 311-312
 warspying, 312
 WEP (Wired Equivalent Privacy),
 319-320
 wireless hacking tools, 325-329

 wireless packet sniffers, 326-327
Wireless Penetration Test reports
 (CORE IMPACT Pro), 385
wireless security, 299-300
 wireless networking, 304-307
 threats, 312-321
 WLANs (wireless LANs), 301-304
Wireshark, 329
WLANs (wireless LANs), 301-304
 benefits, 303
 radio frequency, 303-304
 standard characteristics, 301-302
 Wi-Fi (Wireless Fidelity), 302-303
World Trade Organization (WTO),
 denial of service attack, 4

X

Xmas tree attacks, 369

Z

zero day attacks, 36
 Zeus botnet, 31
 ZFW (zone-based firewalls), 224-229



FREE Online Edition

Your purchase of **Network Security First-Step** includes access to a free online edition for 45 days through the Safari Books Online subscription service. Nearly every Cisco Press book is available online through Safari Books Online, along with more than 5,000 other technical books and videos from publishers such as Addison-Wesley Professional, Exam Cram, IBM Press, O'Reilly, Prentice Hall, Que, and Sams.

SAFARI BOOKS ONLINE allows you to search for a specific answer, cut and paste code, download chapters, and stay current with emerging technologies.

Activate your FREE Online Edition at www.informit.com/safarifree

- **STEP 1:** Enter the coupon code: OLHVFWH.
- **STEP 2:** New Safari users, complete the brief registration form. Safari subscribers, just log in.

If you have difficulty registering on Safari or accessing the online edition, please e-mail customer-service@safaribooksonline.com

