

NETWORKING

Guide to Networking Essentials

Sixth Edition

Gregory Tomsho



This is an electronic version of the print textbook. Due to electronic rights restrictions, some third party content may be suppressed. Editorial review has deemed that any suppressed content does not materially affect the overall learning experience. The publisher reserves the right to remove content from this title at any time if subsequent rights restrictions require it. For valuable information on pricing, previous editions, changes to current editions, and alternate formats, please visit www.cengage.com/highered to search by ISBN#, author, title, or keyword for materials in your areas of interest.



Guide to Networking Essentials

Sixth Edition

Greg Tomsho

 **COURSE TECHNOLOGY**
CENGAGE Learning™

Australia • Brazil • Japan • Korea • Mexico • Singapore • Spain • United Kingdom • United States

**Guide to Networking Essentials,
Sixth Edition**
Greg Tomsho

Vice President, Editorial: Dave Garza
Director of Learning Solutions: Matt Kane
Executive Editor: Steve Helba
Acquisitions Editor: Nick Lombardi
Managing Editor: Marah Bellegarde
Senior Product Manager: Michelle Ruelos Cannistraci
Developmental Editor: Lisa M. Lord
Editorial Assistant: Sarah Pickering
Vice President, Marketing: Jennifer Ann Baker
Marketing Director: Deborah S. Yarnell
Senior Marketing Manager: Erin Coffin
Associate Marketing Manager: Shanna Gibbs
Production Director: Carolyn Miller
Production Manager: Andrew Crouth
Senior Content Project Manager: Andrea Majot
Senior Art Director: Jack Pendleton
Simulations Designer: Angela Poland

© 2011, 2007, 2004, 2003, 2001, 1998 Course Technology, Cengage Learning

ALL RIGHTS RESERVED. No part of this work covered by the copyright herein may be reproduced, transmitted, stored or used in any form or by any means graphic, electronic, or mechanical, including but not limited to photocopying, recording, scanning, digitizing, taping, Web distribution, information networks, or information storage and retrieval systems, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the publisher.

For product information and technology assistance, contact us at
Cengage Learning Customer & Sales Support, 1-800-354-9706

For permission to use material from this text or product, submit all requests online at **www.cengage.com/permissions**
Further permissions questions can be emailed to
permissionrequest@cengage.com

Library of Congress Control Number: 2010943258

ISBN-13: 978-1-111-31252-7

ISBN-10: 1-111-31252-4

Course Technology

20 Channel Center Street
Boston, MA 02210
USA

Cengage Learning is a leading provider of customized learning solutions with office locations around the globe, including Singapore, the United Kingdom, Australia, Mexico, Brazil, and Japan. Locate your local office at **international.cengage.com/region**.

Cengage Learning products are represented in Canada by Nelson Education, Ltd.

For your lifelong learning solutions, visit **www.cengage.com/coursestechnology**.

Purchase any of our products at your local college store or at our preferred online store **www.cengagebrain.com**.

Visit our corporate website at **cengage.com**.

Some of the product names and company names used in this book have been used for identification purposes only and may be trademarks or registered trademarks of their respective manufacturers and sellers.

Microsoft and the Office logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Course Technology, a part of Cengage Learning, is an independent entity from the Microsoft Corporation, and not affiliated with Microsoft in any manner.

Any fictional data related to persons or companies or URLs used throughout this book is intended for instructional purposes only. At the time this book was printed, any such data was fictional and not belonging to any real persons or companies.

Course Technology and the Course Technology logo are registered trademarks used under license.

Course Technology, a part of Cengage Learning, reserves the right to revise this publication and make changes from time to time in its content without notice.

The programs in this book are for instructional purposes only. They have been tested with care, but are not guaranteed for any particular intent beyond educational purposes. The author and the publisher do not offer any warranties or representations, nor do they accept any liabilities with respect to the programs.



Brief Table of Contents

INTRODUCTION	xvii
READ THIS BEFORE YOU BEGIN	xxiv
CHAPTER 1	
Introduction to Computer Networks	1
CHAPTER 2	
Network Hardware Essentials	61
CHAPTER 3	
Network Topologies and Technologies	109
CHAPTER 4	
Network Media	163
CHAPTER 5	
Network Protocols	207
CHAPTER 6	
Network Reference Models and Standards	261
CHAPTER 7	
Network Hardware in Depth	289
CHAPTER 8	
Network Operating System Fundamentals	331
CHAPTER 9	
Server Management and Administration	395
CHAPTER 10	
Introduction to Network Security	459
CHAPTER 11	
Supporting a Small-Business Network	503
CHAPTER 12	
Wide Area Network Essentials	533
CHAPTER 13	
Troubleshooting and Support	565
APPENDIX A	
Common Networking Standards	619
APPENDIX B	
Older and Obsolete Technologies	629
APPENDIX C	
Network Troubleshooting Guide	645
APPENDIX D	
Networking Resources, Online and Offline	649
APPENDIX E	
A Step-by-Step Guide to Using Server Virtualization Software	651
GLOSSARY	689
INDEX	703

This page intentionally left blank



Table of Contents

INTRODUCTION	xvii
READ THIS BEFORE YOU BEGIN	xxiv
CHAPTER 1	
Introduction to Computer Networks	1
An Overview of Computer Concepts	2
Basic Functions of a Computer	2
Storage Components	4
Personal Computer Hardware	5
Computer Boot Procedure	9
Hands-On Project 1-1	9
How the Operating System and Hardware Work Together	12
The Fundamentals of Network Communication	14
Network Components	14
Hands-On Project 1-2	15
Steps of Network Communication	17
Layers of the Network Communication Process	18
How Two Computers Communicate on a LAN: Some Details	19
Hands-On Project 1-3	21
Hands-On Project 1-4	22
Network Terms Explained	26
LANs, Internetworks, WANs, and MANs	26
Packets and Frames	29
Clients and Servers	32
Network Models	34
Peer-to-Peer/Workgroup Model	34
Server/Domain-Based Model	36
Hands-On Project 1-5	38
Hands-On Project 1-6	41
Hands-On Project 1-7	43
Network Servers	44
Domain Controller/Directory Servers	44
File and Print Servers	45
Application Servers	45
Communication Servers	45
E-Mail/Fax Servers	45
Web Servers	46
Additional Network Services	46
Server Hardware Requirements	46
Specialized Networks	47
Storage Area Networks	47
Wireless Personal Area Networks	47
Hands-On Project 1-8	49
Chapter Summary	50
Key Terms	51
Review Questions	53
Case Projects	59

CHAPTER 2

Network Hardware Essentials 61

- Network Repeaters and Hubs 62
 - Multiport Repeaters and Hubs 63
- Network Switches 66
 - Basic Switch Operation 66
 - Hands-On Project 2-1 69
 - Hands-On Project 2-2 72
 - Hands-On Project 2-3 74
 - Hands-On Project 2-4 75
- Wireless Access Points 77
 - Basic AP Operation 77
- Network Interface Cards 79
 - NIC Basics 79
 - Selecting a NIC 82
 - NIC Drivers 82
 - Wireless NICs 84
 - Hands-On Project 2-5 85
 - Hands-On Project 2-6 86
- Routers 88
 - Routers Connect LANs 90
 - Routers Create Broadcast Domains 92
 - Routers Work with IP Addresses and Routing Tables 93
 - Hands-On Project 2-7 96
 - Hands-On Project 2-8 99
- Chapter Summary 100
- Key Terms 101
- Review Questions 102
- Challenge Labs 107
- Case Projects 108

CHAPTER 3

Network Topologies and Technologies 109

- Physical Topologies 110
 - Physical Bus Topology 110
 - Physical Star Topology 112
 - Physical Ring Topology 115
 - Point-to-Point Topology 116
- Logical Topologies 117
 - Hands-On Project 3-1 120
- Network Technologies 121
 - Network Technologies and Media 121
 - Ethernet Networks 122
 - Ethernet Standards 127
 - Additional Ethernet Standards 130
 - Hands-On Project 3-2 133
 - Hands-On Project 3-3 136
 - 802.11 Wi-Fi 137
 - Hands-On Project 3-4 141
 - Token Ring Networks 144
 - Fiber Distributed Data Interface Technology 145
 - Internet Access Technologies 145

Chapter Summary	150
Key Terms	151
Review Questions	153
Challenge Labs	157
Case Projects	159

CHAPTER 4

Network Media	163
Wired Networking	164
Criteria for Choosing Network Media	164
Coaxial Cable	167
Twisted-Pair Cable	167
Structured Cabling: Managing and Installing a UTP Cable Plant	172
Hands-On Project 4-1	179
Hands-On Project 4-2	181
Hands-On Project 4-3	183
Fiber-Optic Cable	184
Fiber-Optic Connectors	186
Fiber-Optic Installation	187
Fiber-Optic Cable Types	187
Wireless Networking	188
Wireless Benefits	188
Types of Wireless Networks	189
Wireless LAN Components	190
Wireless LAN Transmission	190
LAN Media Selection Criteria	196
Chapter Summary	197
Key Terms	198
Review Questions	200
Case Projects	204

CHAPTER 5

Network Protocols	207
TCP/IP's Layered Architecture	208
Hands-On Project 5-1	210
Hands-On Project 5-2	211
Role of the Network Access Layer	213
Role of the Internetwork Layer	213
Protocols at the Internetwork Layer	215
Hands-On Project 5-3	219
Hands-On Project 5-4	219
Hands-On Project 5-5	221
Hands-On Project 5-6	221
Role of the Transport Layer	222
TCP: The Reliable Transport Layer	224
Role of the Application Layer	225
Hands-On Project 5-7	229
Hands-On Project 5-8	231
IP Addressing	232
IP Address Classes	232
Private IP Addresses	233

Network Address Translation	234
Classless Interdomain Routing	236
Subnet Masks	236
Binary Arithmetic	239
Calculating a Subnet Mask	243
Supernetting	247
Introduction to Internet Protocol Version 6	248
Chapter Summary	251
Key Terms	252
Review Questions	254
Challenge Labs	258
Case Projects	260
CHAPTER 6	
Network Reference Models and Standards	261
Introducing the OSI and IEEE 802 Networking Models	262
Role of a Reference Model	263
Structure of the OSI Model	264
Application Layer	268
Presentation Layer	269
Session Layer	270
Transport Layer	270
Network Layer	272
Data Link Layer	273
Physical Layer	274
Summary of the OSI Model	274
IEEE 802 Networking Standards	275
IEEE 802 Specifications	276
IEEE 802 Extensions to the OSI Reference Model	277
Hands-On Project 6-1	278
Hands-On Project 6-2	279
Hands-On Project 6-3	279
Hands-On Project 6-4	280
Chapter Summary	280
Key Terms	281
Review Questions	282
Challenge Labs	286
Case Projects	286
CHAPTER 7	
Network Hardware in Depth	289
Network Switches in Depth	290
Switch Port Modes of Operation	291
Creating the Switching Table	292
Frame Forwarding Methods	293
Advanced Switch Features	294
Hands-On Project 7-1	299
Routers in Depth	300
Router Interfaces	302
Routing Tables	303
Routing Protocols	305

Access Control Lists	308
Hands-On Project 7-2.	308
Wireless Access Points in Depth.	310
Basic Wireless Settings	310
Wireless Security Options	312
Advanced Wireless Settings	313
Network Interface Cards in Depth.	314
PC Bus Options	314
Advanced Features of NICs.	317
Chapter Summary	319
Key Terms.	319
Review Questions.	322
Challenge Labs	326
Case Projects	330

CHAPTER 8

Network Operating System Fundamentals.	331
Operating System Fundamentals	332
The File System	332
Hands-On Project 8-1.	334
Hands-On Project 8-2.	337
Processes and Services.	340
The Kernel.	341
Hands-On Project 8-3.	342
Hands-On Project 8-4.	344
Network Operating System Overview	345
The Role of a Client Operating System	346
Hands-On Project 8-5.	350
Hands-On Project 8-6.	351
Hands-On Project 8-7.	354
Hands-On Project 8-8.	355
The Role of a Server Operating System	357
Centralized User Account and Computer Management	358
Centralized Storage	360
Infrastructure Services.	361
Server and Network Fault Tolerance	364
Additional Server Features	365
Operating System Virtualization	366
Hosted Virtualization	367
Bare-Metal Virtualization	376
Installing an OS	380
Planning for and Installing Windows Server 2008.	380
Planning for and Installing Linux.	384
Chapter Summary	386
Key Terms.	387
Review Questions.	388
Challenge Labs	392
Case Projects	393

CHAPTER 9

Server Management and Administration 395

- Managing User and Group Accounts 396
 - Account and Password Conventions 396
 - Working with Accounts in Windows 397
 - Hands-On Project 9-1. 401
 - Hands-On Project 9-2. 406
 - Working with Accounts in Linux. 409
 - Hands-On Project 9-3. 411
 - Hands-On Project 9-4. 412
- Storage and File System Management 414
 - Volumes and Partitions. 414
 - The FAT File System 415
 - The NTFS File System 416
 - Hands-On Project 9-5. 421
 - Hands-On Project 9-6. 423
 - Hands-On Project 9-7. 425
 - The Linux File System 427
- Working with Shared Files and Printers 428
 - Sharing Files and Printers in Windows. 429
 - Hands-On Project 9-8. 431
 - Sharing Files and Printers in Linux 434
- Monitoring System Reliability and Performance 435
 - Event Viewer 435
 - Performance Monitor 437
 - Hands-On Project 9-9. 439
 - Hands-On Project 9-10. 440
 - Windows System Resource Manager 442
- Backup and Fault Tolerance 442
 - Windows Backup. 443
 - Hands-On Project 9-11. 444
 - Protecting Data with Fault Tolerance. 446
- Chapter Summary 450
- Key Terms. 451
- Review Questions. 452
- Challenge Labs 456
- Case Projects 457

CHAPTER 10

Introduction to Network Security. 459

- Network Security Overview and Policies 460
 - Developing a Network Security Policy 460
 - Determining Elements of a Network Security Policy 461
 - Understanding Levels of Security 462
- Securing Physical Access to the Network 463
 - Physical Security Best Practices 463
- Securing Access to Data 465
 - Implementing Secure Authentication and Authorization 466
 - Hands-On Project 10-1. 468
 - Hands-On Project 10-2. 470
 - Securing Data with Encryption 472

Securing Communication with Virtual Private Networks	474
Protecting Networks with Firewalls	477
Hands-On Project 10-3	479
Hands-On Project 10-4	481
Protecting a Network from Worms, Viruses, and Rootkits	483
Protecting a Network from Spyware and Spam	484
Hands-On Project 10-5	485
Implementing Wireless Security	486
Using an Attacker's Tools to Stop Network Attacks	487
Hands-On Project 10-6	492
Chapter Summary	494
Key Terms	495
Review Questions	497
Challenge Labs	500
Case Projects	501
CHAPTER 11	
Supporting a Small-Business Network	503
Addressing the Needs of Small-Business Networks	504
Data and Application Sharing in a Small Business	504
Hands-On Project 11-1	508
Equipment Sharing in a Small Business	510
Equipping Small-Business Networks	511
Servers and Desktops	511
Networking Equipment	512
Communicating with the Outside World	515
Identifying Requirements for Small-Business Applications	517
Accounting Software	518
Sales and Contact Management Software	518
Windows Small Business Server	519
Hosted Applications	519
Is Linux a Viable Desktop Alternative to Windows?	520
Supporting a Small Business	522
Entrepreneurs Wanted	522
Securing a Small-Business Network	523
Managing a Small-Business Network	524
Chapter Summary	525
Key Terms	526
Review Questions	527
Challenge Labs	530
Case Projects	531
CHAPTER 12	
Wide Area Network Essentials	533
Wide Area Network Fundamentals	534
WAN Devices	534
WAN Connection Methods	536
Circuit-Switched WANs	537
Leased Lines	539
Packet-Switched WANs	541
WANs over the Internet	546
WAN Equipment	547

- Remote Access Networking** 549
 - Making a VPN Connection in Windows 550
 - Making a Dial-Up Connection 551
 - Remote Access Networking via the Web 552
 - Hands-On Project 12-1. 553
 - Hands-On Project 12-2. 555
- Cloud Computing** 555
 - Hosted Applications 556
 - Hosted Platforms 556
 - Hosted Infrastructure 557
- Chapter Summary** 558
- Key Terms**. 558
- Review Questions**. 561
- Challenge Labs** 564
- Case Projects** 564

CHAPTER 13

- Troubleshooting and Support** 565
 - Documenting Your Network**. 566
 - Documentation and Network Changes. 567
 - Documentation and Troubleshooting 567
 - Documentation and IT Staffing 568
 - Documentation and Standards Compliance. 568
 - Documentation and Technical Support. 568
 - Documentation and Network Security 568
 - What Should Be Documented? 569
 - Approaches to Network Troubleshooting**. 571
 - Trial and Error 571
 - Solve by Example. 573
 - The Replacement Method. 574
 - Step by Step with the OSI Model. 574
 - The Problem-Solving Process**. 576
 - Step 1: Determine the Problem Definition and Scope. 578
 - Step 2: Gather Information 579
 - Step 3: Consider Possible Causes 580
 - Step 4: Devise a Solution 581
 - Step 5: Implement the Solution 582
 - Step 6: Test the Solution 583
 - Step 7: Document the Solution 584
 - Step 8: Devise Preventive Measures 584
 - Making Use of Problem-Solving Resources**. 584
 - Experience. 584
 - The World Wide Web 586
 - Network Documentation 587
 - Network Troubleshooting Tools** 589
 - Ping and Trace Route. 590
 - Hands-On Project 13-1. 593
 - Network Monitors 594
 - Hands-On Project 13-2. 594
 - Hands-On Project 13-3. 595
 - Protocol Analyzers 597
 - Hands-On Project 13-4. 598
 - Time-Domain Reflectometer 600

Basic Cable Testers	601
Advanced Cable Testers	601
Advanced Monitoring Tools	601
Common Troubleshooting Situations	602
Cabling and Related Components	602
Power Fluctuations	603
Upgrades	603
Poor Network Performance	603
Disaster Recovery	604
Backing Up Network Data	604
Backup Types	605
System Repair and Recovery in Windows	606
Hands-On Project 13-5	608
Chapter Summary	611
Key Terms	612
Review Questions	613
Challenge Labs	616
Case Projects	618

APPENDIX A

Common Networking Standards	619
Standards-Making Process	619
Important Standards Bodies	620
American National Standards Institute	621
Comité Consultatif International Téléphonique et Télégraphique	622
Electronic Industries Alliance	623
Internet Architecture Board	623
Institute of Electrical and Electronics Engineers, Inc.	624
International Organization for Standardization	625
Object Management Group	625
The Open Group	626
The World Wide Web Consortium	627
Internet Corporation for Assigned Names and Numbers	627

APPENDIX B

Older and Obsolete Technologies	629
Thinwire Ethernet (Thinnet)	629
Thickwire Ethernet (Thicknet)	630
10Base5 Ethernet	631
10Base2 Ethernet	633
100VG-AnyLAN	634
Ethernet Frame Types	635
Ethernet 802.3 Frame Type	636
Ethernet 802.2	636
Ethernet SNAP	636
The Token Ring Architecture	636
Token Ring Function	637
Beaconing	637
Hardware Components	638

The AppleTalk Environment 639
 LocalTalk 640
 EtherTalk and TokenTalk 640
The Fiber Distributed Data Interface (FDDI) Architecture 640
Routable Versus Nonroutable Protocols 642
 The IPX/SPX Protocol Suite 642
 IPX/SPX Implementations 642
 NetBIOS and NetBEUI 643
 AppleTalk 644

APPENDIX C

Network Troubleshooting Guide **645**
 General Questions for Troubleshooting 645
 Cabling Problems 646
 Problems with NICs 646
 Driver Problems 646
 Problems with Network Operations 646
 Problems with Network Printing and Fax Services 647
 Problems with Client/Server Computing 647
 Problems with Network Accounts 647
 Problems with Data Security 647
 Problems with WAN Communication 648

APPENDIX D

Networking Resources, Online and Offline **649**
 Printed Materials 649
 Online/Electronic Materials 649

APPENDIX E

A Step-by-Step Guide to Using Server Virtualization Software **651**
 Microsoft Virtual PC 652
 Requirements for Microsoft Virtual PC 652
 Guest OSs Supported 652
 Downloading Microsoft Virtual PC 653
 Installing Microsoft Virtual PC 653
 Creating a Virtual Machine and Installing a Guest OS 654
 Installing an OS from an ISO Image 657
 Configuring Networking and Hardware Options 657
 Host Key Options 659
 Microsoft Virtual Server 660
 Guest OSs Supported 660
 Host OSs Supported 661
 Requirements for Microsoft Virtual Server 661
 Downloading Microsoft Virtual Server 661
 Installing Microsoft Virtual Server 662
 Creating a Virtual Machine and Installing a Guest OS 664
 Installing an OS from an ISO Image 667
 Configuring Networking and Hardware Options 668
 Host Key Options 671
 VMware Server 672
 Guest OSs Supported 672
 Host OSs Supported 673
 Requirements for VMware Server 673

Downloading VMware Server	674
Installing VMware Server	674
Creating a Virtual Machine and Installing a Guest OS	676
Installing an OS from an ISO Image	680
Configuring Networking Options	680
Configuring Hardware Options	681
Installing VMware Tools	682
Other Virtual Systems	684
VMware Workstation	685
Microsoft Hyper-V	686
Glossary	689
Index	703

CD RESOURCE CONTENTS

Flash Movies

1. Using VMware Workstation
2. Installing Windows Server 2008
3. Installing CentOS 5.4

Simulations

1. Layers of the Network Communication Process
2. Communication Between Two Computers
3. Basic Operation of a Hub
4. Basic Operation of a Switch
5. How a NIC Works
6. Router Operation in a Simple Internetwork
7. Ethernet Operation Using CSMA/CD
8. Wireless LAN Operation
9. The Changing Frame Header
10. Demonstrating NAT/PAT
11. Peer Communication with the OSI Model
12. OSI Model: Layer Names Activity
13. OSI Model: Layer Descriptions Activity
14. Build a Data Frame Activity
15. STP Prevents Switching Loops
16. How Switches Use Trunk Ports with VLANs
17. Routers Use Multiple Paths in an Internetwork
18. How E-Mail Works

This page intentionally left blank



Introduction

***Guide to Networking Essentials, Sixth Edition*, is intended to serve the needs of students, instructors, aspiring information technology professionals, and others who are interested in learning more about networking technologies but who might have little or no background in this subject matter. This book's extensive and broad coverage of computer networking technologies gives you a solid networking background to pursue a number of certifications, including Network+, CCNA, MCTS, and Security+. With the extensive use of tables that compare important properties of networking technologies, this book also makes an excellent reference.**

What's New to This Edition

The sixth edition is completely updated and rewritten, giving a fresh face to a long-standing title. Recognizing that many students are learning computer concepts at the same time they're learning about networking, the first chapter includes a refresher on computer components and terminology. This new edition includes coverage of the latest networking technologies and operating systems, including 60/100 Gigabit Ethernet, cloud computing, Windows 7, Windows Server 2008, and Ubuntu Linux 10.4. A new chapter, "Network Hardware in Depth," delves deeper into switches, routers, APs, and NICs than in previous editions. In keeping with the latest trends in networking, this edition has updated coverage on the 802.11 and 802.16 wireless standards and virtual private networks. A new section on virtualization covers both server and desktop virtualization technologies.

All new hands-on projects that are interwoven throughout the chapter text allow students to apply the concepts they learn in the chapter. Challenge labs, a new feature of this edition, are

placed at the end of many chapters. They give students an opportunity to apply what they have learned from the chapter material and hands-on projects in a format that requires additional research and problem-solving skills.

A CD packed with new simulations is included. These completely redesigned simulations with audio narrations give visually oriented students an innovative tool to help them grasp difficult networking concepts. The simulations cover topics ranging from basic LAN communication to Network Address Translation (NAT) and Internet e-mail operation. New drag-and-drop exercises reinforce concepts on the OSI model and network frame formats.

Intended Audience

Guide to Networking Essentials, Sixth Edition, is intended for people who are getting started in computer networking and want to gain a solid understanding of a broad range of networking technologies. This book is ideal for would-be information technology professionals who want to pursue certifications in a variety of computer networking fields as well as those in a managerial role who want a firm grasp of networking technology concepts. To understand the material in this book, you should have a background in basic computer concepts and have worked with the Windows and/or Linux operating system. This book is ideal for use in a classroom or an instructor-led training environment and is also an effective learning tool for self-paced training.

Coping with Change on the Web

Sooner or later, all the specifics on Web-based resources mentioned in this book will become outdated or be replaced by newer information. In some cases, the URLs listed in this book might lead to their replacements; in other cases, they'll lead nowhere, resulting in the dreaded 404 error message "File not found."

When that happens, please don't give up! There's always a way to find what you want on the Web, if you're willing to invest some time and energy. Most large or complex Web sites offer a search engine. As long as you can get to the site itself, you can use this tool to help you find what you need. In addition, try using general search tools, such as *www.google.com*, *www.yahoo.com*, or *www.bing.com*, to find related information. The bottom line is if you can't find something where the book says it should be, start looking around. It's likely to be somewhere!

Chapter Descriptions

Here's a summary of the topics covered in each chapter of this book:

- **Chapter 1**, "Introduction to Computer Networks," introduces many of the computer and networking terms and technologies discussed in detail in later chapters.
- In **Chapter 2**, "Network Hardware Essentials," you learn about the basic operation of hubs, switches, access points, network interface cards, and routers.
- **Chapter 3**, "Network Topologies and Technologies," discusses logical and physical topologies and the LAN technologies that use them.

- **Chapter 4**, “Network Media,” covers the cables and connectors required to connect network devices and discusses wireless networking.
- In **Chapter 5**, “Network Protocols,” you learn about the purpose and operation of network protocols, focusing on the TCP/IP protocol suite. Special emphasis is given to IP addressing and subnetting, and a section on IPv6 is included.
- **Chapter 6**, “Network Reference Models and Standards,” discusses the OSI model’s seven-layer architecture and gives you an overview of the IEEE 802 networking standards.
- **Chapter 7**, “Network Hardware in Depth,” revisits the hardware components of networks discussed in Chapter 2, providing more in-depth coverage of each type of device.
- In **Chapter 8**, “Network Operating System Fundamentals,” you learn about network operating system features and the most common types of services provided by current server operating systems. A new section on server and desktop virtualization discusses using virtual machines in data centers and on the desktop.
- **Chapter 9**, “Server Management and Administration,” discusses everyday tasks that network and server administrators perform, including working with user and group accounts, creating and managing file shares, and monitoring system performance and reliability.
- In **Chapter 10**, “Introduction to Network Security,” you learn about Trojan programs, worms, spam, denial-of-service and other attacks, spyware, backdoors, rootkits, and more security concerns. In addition, you learn how to develop a security policy.
- **Chapter 11**, “Supporting a Small-Business Network,” discusses the unique technology requirements of small businesses to give you more insight into addressing a small business’s computer and networking needs.
- In **Chapter 12**, “Wide Area Network Essentials,” you learn how WAN technologies, such as frame relay and ISDN, work and how to implement these technologies to create networks that can extend across your town or across the country. In addition, you’re introduced to remote access protocols and cloud computing, the latest networking model.
- **Chapter 13**, “Troubleshooting and Support,” discusses what you can do to prevent network downtime, data loss, and system failures. In addition, you learn about the problem-solving process, several different approaches to solving network problems, and the tools for troubleshooting networks.
- **Appendix A**, “Common Networking Standards,” explains the standards-making process as it applies to networking. It also covers the most important and influential standards-making bodies in the United States and worldwide.
- **Appendix B**, “Older and Obsolete Technologies,” gives you an overview of older or obsolete technologies included in previous editions of this book. Details of these technologies have been removed in chapters to make room for current topics but might still be important for some readers.
- **Appendix C**, “Network Troubleshooting Guide,” summarizes advice on how to recognize, isolate, and diagnose trouble on a network, whether it’s related to media, hardware, or software.

- **Appendix D**, “Networking Resources, Online and Offline,” is a compilation of printed and online resources you can use for additional research into networking essentials.
- **Appendix E**, “A Step-by-Step Guide to Using Server Virtualization Software,” discusses how to use some of the most common virtualization programs.

Features

To help you understand networking concepts thoroughly, this book incorporates many features designed to enhance your learning experience:

- *Chapter objectives*—Each chapter begins with a detailed list of the concepts to be mastered. This list is a quick reference to the chapter’s contents and a useful study aid.
- *Hands-on projects*—Although understanding the theory behind networking technology is important, nothing can improve on real-world experience. Each chapter has projects interwoven throughout the text aimed at giving you hands-on practice.
- *Challenge labs*—After performing the chapter’s step-by-step hands-on projects, you get an opportunity to challenge your knowledge and skills with these labs included at the end of most chapters.
- *Screenshots, illustrations, and tables*—Numerous illustrations of networking components aid you in visualizing common network setups, theories, and architectures. In addition, tables summarize details in an at-a-glance format and provide comparisons of both practical and theoretical information. Examples of concepts and system features encompass Windows desktop operating systems, Windows Server, Linux, and Novell NetWare. When client-side functionality is important, desktop OSs are used for examples; when the focus is on server-side functionality, server OSs are used. Because most campus labs use Windows OSs, these products have been used for most screenshots and hands-on projects.
- *Simulations*—In many chapters, you’ll find references to simulations included on the CD accompanying this book. These simulations provide an audio and visual learning experience, demonstrating concepts such as basic LAN communication, Ethernet switches, routing, Network Address Translation, Internet e-mail operation, and more.
- *Chapter summary*—Each chapter ends with a summary of the concepts introduced in the chapter. These summaries are a helpful way to recap and revisit the material covered in the chapter.
- *Key terms*—All terms in the chapter introduced with bold text are gathered together in the Key Terms list at the end of the chapter. This list gives you an easy way to check your understanding of important terms and is a useful reference.
- *Review questions*—The end-of-chapter assessment begins with review questions that reinforce the concepts and techniques covered in each chapter. Answering these questions helps ensure that you have mastered important topics.
- *Case projects*—Each chapter closes with one or more case projects that describe networking situations. You’re asked to evaluate these situations and decide on a course of action to remedy the problems. This valuable tool helps you sharpen decision-making and troubleshooting skills, which are important aspects of network administration.

Text and Graphic Conventions

Additional information and exercises have been added to this book to help you better understand what's being discussed in the chapter. Icons throughout the text alert you to these additional materials:



Notes draw your attention to additional helpful material related to the topic being discussed.



Tips offer extra information based on the author's experience about useful resources, how to approach a problem, and what to do in real-world situations.



The Caution icon identifies important information about potential mistakes or hazards.



Each hands-on project in this book is preceded by this icon.



Simulation icons refer you to simulations on the CD accompanying this book. These simulations reinforce the concepts being discussed.



Challenge labs identify labs that apply your knowledge and skills to networking tasks.



This icon marks the end-of-chapter case projects, which are scenario-based assignments that ask you to apply what you have learned in the chapter.

Instructor Support

The following supplemental materials are available when this book is used in a classroom setting. These supplements are included on the Instructor Resources CD (ISBN 1111312532), or instructors can download them from <http://login.cengage.com>. In addition, the author

maintains a Web site at <http://books.tomsho.com> with lab notes, errata, and the latest lab setup guide as well as hints and tips for teaching with this book.

- *Electronic Instructor's Manual*—The Instructor's Manual that accompanies this book includes additional instructional material to assist in class preparation, including suggestions for lecture topics and lab activities and tips on setting up a lab for hands-on projects.
- *Solutions*—Includes answers to end-of-chapter material, including review questions and, when applicable, challenge labs and case projects.
- *ExamView®*—This book is accompanied by ExamView, a powerful testing software package that instructors can use to create and administer printed, computer (LAN-based), and Internet exams. ExamView includes hundreds of questions that correspond to the topics covered in this book, enabling students to generate detailed study guides that include page references for further review. The computer-based and Internet testing components allow students to take exams at their computers, and they save instructors time by grading each exam automatically.
- *PowerPoint presentations*—This book comes with Microsoft PowerPoint slides for each chapter. They're included as a teaching aid for classroom presentation, to make available to students on the network for chapter review, or to be printed for classroom distribution. Instructors, please feel free to add your own slides for additional topics you introduce to the class.
- *Figure files*—All figures and tables in the book are reproduced on the Instructor Resources CD. Similar to PowerPoint presentations, they're included as a teaching aid for classroom presentation. In this edition, selected figure files are available in color on the Instructor Resources CD.

Contact the Author

I would like to hear from you. Please e-mail me with any problems, questions, suggestions, or corrections. I even accept compliments! This book has staying power, so I wouldn't be surprised to see a seventh edition in the future. Your comments and suggestions are invaluable for shaping the content of that next edition. You can contact me at NetEss@tomsho.com. In addition, please visit my Web site at <http://books.tomsho.com>, where you can access lab notes, errata, and submit comments and suggestions.

Online Resources

Additional materials designed especially for you might be available online for your course. Go to www.cengage.com/coursetechnology and search for this book title periodically for more details.

You can also visit www.cengagebrain.com to find additional course materials, including CourseMate. At the CengageBrain.com home page, search for your book's ISBN (listed on the back cover) by using the search box at the top of the page. This search takes you to the product page where you can find these resources. Cengage Learning's Networking

CourseMate brings course concepts to life with interactive learning and study and exam preparation tools that support the printed book. The Networking CourseMate for this book offers many helpful resources, including an interactive eBook, Simulations, Flashcards, Quizzes, Crossword Puzzles, and Engagement Tracker.

Acknowledgements

I would like to thank the team at Course Technology for this opportunity to improve and expand on the fifth edition of this book. This team includes but is not limited to Michelle Ruelos Cannistraci, Senior Product Manager; Andrea Majot, Senior Content Project Manager; Nick Lombardi, Acquisitions Editor; and Green Pen Quality Assurance for testing projects and labs for accuracy. Thanks especially to my development editor, Lisa Lord, for her excellent guidance in creating a polished product. The simulations and Flash movies could not have been done without the excellent artwork and Flash animation expertise of Angela Poland, so to her a hearty thanks and congratulations for a job well done. Additional praise and special thanks goes to my beautiful wife, Julie, our daughters, Camille and Sophia, and our son, Michael, who all deserve medals for their patience and support while going husbandless and fatherless for almost 6 months.

I would also like to thank the following reviewers, who guided me with excellent and helpful feedback on each chapter: Dave Braunschweig, Harper College; R. Scott Domowicz, Erie Institute of Technology; Mark Highum, Bay College; Mitchell Kofi, Northern Virginia Community College; Nestor Reyes, Southwest Florida College; and Richard Tuttle, Portland Community College.

Read This Before You Begin

The hands-on projects in this book are intended to reinforce the networking concepts students learn while studying the chapter material. Many projects require students to have administrative access to their computer. Because many classrooms are shared with other classes, using virtualization is highly recommended so that students can have administrative control over their virtual machines without affecting the host operating system. Chapter 8 and Appendix E discuss some widely used virtualization software packages. VMware Player is recommended because it's a free, robust solution for creating and running virtual machines. Microsoft Virtual PC and the open-source VirtualBox are also useful products.

The author not only wrote the book, but also uses it in his own classrooms. You can find lab notes, teaching tips, errata, updates, and news on his Web site at <http://books.tomsho.com>. You can also submit comments and suggestions.

Classroom Computers

There should be one classroom computer per student, and these computers can be virtual machines.

- Used for most hands-on projects and challenge labs.
- Windows 7 Enterprise, Professional, or Ultimate Edition preferred.
- A second partition or hard drive formatted as NTFS of at least 1.5 GB and assigned the drive letter D.
- A second (or third, if D is a separate drive from C) unallocated drive.
- Computer name: NET-XX (replacing XX with the number assigned to the student).
- Workgroup name: NetEss.
- IP address via DHCP initially (static addresses also work).
- In the hands-on project in Chapter 5 for setting the IP address configuration to 192.168.100.XX/24, instructors need to give students a valid default gateway address and DNS server address that work with this configuration or provide alternative addresses.
- Valid router and DNS server configured with Internet access.
- Administrator account named NetAdmin with the password Password01 (set to never expire).

Instructor Computer

This computer can be a virtual machine and is the same as the student computer except for the following:

- Computer name: Net-instr

- Shared folder named NetDocs, allowing NetAdmin users Read and Change sharing permissions and Modify NTFS permissions
- User account name NetAdmin with the password Password01

Separate Lab Computers

These computers must be physical computers, not virtual machines.

- Used for Hands-On Projects 1-2, 2-2, 2-3, 2-4, 2-5, and 2-7.
- Minimum of three computers (Students can work in groups, with each group having three computers, or instructors can do projects requiring physical computers as a demonstration.)
- Windows 7 preferred (Projects are written for Windows 7, but Windows Vista or XP can be substituted by modifying instructions slightly.)
- Wireshark installed
- 10/100 Ethernet NICs to install (USB is acceptable.)

Networking Equipment

- Two 10/100 hubs (minimum of four ports)
- Two 10/100 switches (minimum of four ports)
- One wireless access point/router 802.11b/g/n with at least WPA support (Linksys WRT54GL or similar preferred); SSID set to NetEss initially, and no security protocols enabled
- 802.11b/g/n NICs (USB is acceptable) with at least WPA support
- Four patch cables, one crossover cable

Additional Tools and Software

- Network diagram software, such as Visio (optional).
- A Linux Live CD or a computer or virtual machine with Linux installed. (Ubuntu Linux 10.4 is used in projects.)
- Windows Server 2008 (for Challenge Lab 12-1).
- A shared printer for students to connect to (optional).
- Wireshark (downloaded and installed by students on classroom computers or installed ahead of time by instructors).
- NetInfo (downloaded and installed by students on classroom computers or installed ahead of time by instructors).
- Simple Server Monitor (downloaded and installed by students on classroom computers or installed ahead of time by instructors).

- RJ-45 crimping tool.
- Punchdown tool.
- Cable stripper.
- Cat 5e or better patch panel for punching down.
- Cat 5e jacks for punching down (optional).
- Cat 5e or better cable.
- RJ-45 plugs (at least four per student if making patch and crossover cables).

Back of the Book CD

The simulations/Flash movies on the CD show common networking processes with animations. Each animation is narrated, but you can mute the sound, if you want, by clicking the speaker icon. The simulations require enabling Flash in your Web browser; if you don't have Flash installed or need to upgrade, go to <http://get.adobe.com/flashplayer/> to download and install the latest player.

Open the simulation menu by double-clicking index.html. After you choose a simulation from the menu, click the Play button. You can pause a simulation by clicking the Pause button and skip to a particular point in the simulation by dragging the slider.

System Requirements for CD Simulations

Minimum system requirements for PCs:

- Operating system: Windows XP with SP2, Vista with SP1, 7, Linux, MAC OS X
- Memory: 512 MB
- Hard drive space: 200 MB
- Screen resolution: 1024 × 768 or higher
- CD-ROM drive
- Sound card and listening device required for audio features
- An Internet connection and a Flash-enabled Web browser (go to <http://get.adobe.com/flashplayer/> to download the most recent Flash Player plug-in, if necessary)

PC setup instructions:

1. Insert the disc in the CD-ROM drive. The program should start automatically. If it doesn't, go to Step 2.
2. In the Computer window, double-click the CD drive icon.
3. Double-click the start.exe file to start the program or double-click index.html.

Technical support:

1-800-648-7450, 8:30 a.m. to 6:30 p.m. EST

E-mail: delmar.help@cengage.com

Introduction to Computer Networks

After reading this chapter and completing the exercises, you will be able to:

- Describe basic computer components and operations
- Explain the fundamentals of network communication
- Define common networking terms
- Compare different network models
- Identify the functions of various network server types
- Describe specialized networks

In only a couple of decades, computer networks have evolved from being a complex technology accessible to only the most tech-savvy of users to being part of most people's everyday lives. Computer networks can be found in almost every business, school, and home. The use of networks is available to anyone with a computer and a network connection, but installation and upkeep of all but the smallest of networks still require a considerable degree of know-how. This chapter starts you on the path toward acquiring the skills to manage a large corporate network or simply configure a home network with a wireless router.

This chapter begins by discussing the computer and its role in a network to give you a foundation for the topics in this book. Next, you examine the components of a network and the fundamentals of communication between computers. Many new terms are introduced and defined, and the varied types of networks and network servers you might encounter are described. Finally, some specialized network types are introduced.

An Overview of Computer Concepts

At the heart of a computer network is the computer. Networks were created to facilitate communication between computing devices, which ultimately facilitates communication between people. So to better understand computer networks, how they work, and how to support them, you must have a solid understanding of computer operations. In fact, most of the devices you encounter when working with a network involve a computer. The most obvious are network servers and workstations that run operating systems, such as Windows, Linux, UNIX, and Mac OS X. Not as obvious are devices such as routers and switches, which move network data from computer to computer and network to network. These complex devices are also computers, although they're specialized computers for performing specific tasks. The next sections discuss the basic functions of a computer and its associated components, along with computer hardware, the boot procedure, and the basic functions of an operating system.

Basic Functions of a Computer

A computer's functions and features can be broken down into the three basic tasks all computers perform: input, processing, and output. Information is input to a computer from a device such as a keyboard or from a storage device such as a hard drive; the central processing unit (CPU) processes the information, and then output is usually created. The following example illustrates the process:

- *Input*—A user running a word-processing program types the letter A on the keyboard, which results in sending a code representing the letter A to the computer.
- *Processing*—The computer's CPU determines what letter was typed by looking up the keyboard code in a table.
- *Output*—The CPU sends instructions to the graphics cards to display the letter A, which is then sent to the computer monitor.

Some components of today's computers are designed to perform only one of these three functions; others are designed to perform two or all three functions. For example, a standard keyboard and mouse perform input functions, and storage devices, such as hard drives, perform

both input (when files are read from the drive) and output (when files are written to the drive). Network cards can perform all three functions. A network card is an output device when data is sent from the computer to the network and an input device when data comes from the network to the computer. In addition, many network cards have rudimentary processors that perform actions on incoming and outgoing data to help supplement the computer's main CPU.

Input Components Before a computer can do any processing, it requires input, commonly from user-controlled devices, such as keyboards and mice, but includes devices such as microphones, Web cameras, and scanners. External interfaces, such as serial, FireWire, and USB ports, can also be used to get input from peripheral devices.

Input is also generated by storage devices, such as hard disks and CDs/DVDs that store computer programs and data files containing computer instructions and data. For example, a spreadsheet program, such as Microsoft Excel, might contain instructions for the CPU to calculate formulas for adding the values of two columns of data and a spreadsheet file called *MyBudget.xls* containing the numbers and formulas the spreadsheet program should use. Both the program (Microsoft Excel) and the data file (*MyBudget.xls*) are used as input to the CPU, which then processes the program instructions and data.

Of course, a spreadsheet program is normally started only when a user double-clicks the spreadsheet program icon or the icon representing the spreadsheet data file. These actions are instigated by user input. Sometimes, however, your computer seems to start performing actions without user input. For example, you might have noticed that your hard drive sometimes shows activity without any obvious action from you to initiate it. However, inputs to a computer can include timers that cause programs to run periodically and data arriving from network cards, for example, that cause a program or process to run. So although it sometimes seems as though your computer has a mind of its own, computers don't actually do anything without first getting input to jolt them into action.

Processing Components A computer's main processing component is the CPU, which executes instructions from computer programs, such as word-processing programs and Web browsers. It also runs the instructions composing the operating system (OS), which provides a user interface and the environment in which applications run. Aside from the CPU, modern computers usually include ancillary processors associated with input/output (I/O) devices, such as graphics cards. These processors are often referred to as onboard processors. The processor on a graphics card, called a graphics processing unit (GPU), takes a high-level graphics instruction, such as "draw a circle," and performs the calculations needed to draw the circle on the display device. With an onboard GPU, the main CPU doesn't have to handle many of the complex calculations current graphical applications require, thereby improving overall system performance. Other devices, such as network interface cards and disk controller cards, might also include onboard processors.

CPUs now are often composed of two or more processors, called **cores**, in one package. A **multicore CPU** is like a person with two brains. With only one brain, you could add four numbers together, but you would probably do it in three sequential summing operations: Add the first number to the second number, take the first sum and add it to the third number, and add that sum to the fourth number to arrive at the final sum. If you had two brains, you'd still need three summing operations, but two could be done simultaneously:



The first brain adds the first two numbers while the second brain is adding the third and fourth numbers; then the second brain gives its results to the first brain, and the first brain sums the results of the first two summing operations. So multicore CPUs enable computers to carry out multiple instructions simultaneously, which results in better overall performance when running demanding applications.

Output Components Output components include monitors and printers, but they also include storage devices, network cards, and speakers, to name a few. The external interfaces mentioned previously as input components can be used as output components, too. For example, a disk drive connected to a USB port allows reading files from the disk (input) and writing files to the disk (output).

Storage Components

Storage components are a major part of a computer's configuration. Generally speaking, the more storage a computer has, the better the performance is. As you saw in the previous section, most storage components are both input and output devices, allowing data to be saved (output) and then accessed again later (input). When most people think of storage, they think of disk drives, CD/DVD drives, and USB flash drives. However, there are two main categories of storage: short-term storage and long-term storage.

RAM: Short-Term Storage Short-term storage is the random access memory (RAM) on a computer. RAM is short-term storage because when power to the computer is turned off, RAM's contents are gone, just as though you erased a whiteboard. When power is restored, RAM has no data stored until the CPU begins to write data to it.

The amount of RAM, or memory, in a computer is crucial to the computer's capability to operate efficiently. RAM is also referred to as "working storage." Everything the CPU is currently processing must be available in RAM, including program instructions and the data the current application requires. So to run a spreadsheet program, there must be enough RAM to load both the spreadsheet program and the data in the spreadsheet. If there's not enough available memory, the spreadsheet program won't run, or the computer will use the disk drive to supplement RAM temporarily.

Neither option is desirable. The reason temporary use of the disk drive isn't optimal is because RAM is thousands of times faster than the fastest disk drives. The time required to access data in RAM is measured in nanoseconds (billionths of a second), but access to data on a disk drive is measured in milliseconds (thousandths of a second). So if the disk drive must be used to supplement RAM while running an application, that application, and indeed the entire computer, slows down precipitously.

On current computers, the amount of RAM installed is usually 1 GB or more. More is generally better, but the amount of RAM that a system can use effectively depends on the OS installed. The 32-bit version of an OS can usually access a maximum of 4 GB of RAM, whereas the 64-bit version can access many thousands of gigabytes. The amount of RAM you actually need depends on how you use your computer. If you usually have only one or two typical business applications open at once, 1 GB or even less is probably enough.

However, if you run complex graphics applications or games or have several applications open simultaneously, you'll likely benefit from having more RAM.

Long-Term Storage Long-term storage maintains its data even when there's no power. Examples include hard disks, CDs/DVDs, and USB flash drives as well as other types of removable media. Long-term storage is used to store document and multimedia files as well as the files that make up applications and the OS. The amount of storage a computer needs depends on the type and quantity of files to be stored. In general, office documents, such as word-processing files, spreadsheets, and presentations, require comparatively little space. Multimedia files—pictures, music files, and videos—require much more space. Long-term storage is plentiful and extremely inexpensive. Hard drive specifications are in units of tens or hundreds of gigabytes, with terabyte (1000 GB) drives quite commonplace now. More details about hard disks are discussed later in “Personal Computer Hardware.”

Data Is Stored in Bits Whether storage is long term or short term, data on a computer is stored and processed as binary digits (“bits,” for short). A bit holds a 1 or 0 value, which make representing bits with electrical pulses easy. For example, a pulse of 5 volts of electricity can represent a 1 bit, and a pulse of 0 volts (or absence of a pulse) can represent a 0 bit. Bits can also be stored as pulses of light, as with fiber-optic cable: A 1 bit is represented by the presence of light and a 0 bit as the absence of light.

Data in a computer, such as the letters in a word-processing document or the music you hear when you play an MP3 music file, is represented by collections of 8 bits, called a byte. You can look at each byte as a printable character in a document. A single byte from an MP3 file plays about 1/17 thousandth of a second of music. To put it another way, one second of MP3 music takes more than 17,000 bytes.

Personal Computer Hardware

Most people are familiar with personal computer (PC) hardware. Other types of computers, such as minicomputers and mainframes, are usually locked away in a heavily air-conditioned room and privy only to the eyes of IT staff. Besides, the basic hardware used to build a PC or a mainframe differs only in the details. This section describes four major PC components housed in a computer case:

- Motherboard
- Hard drive
- RAM
- BIOS/CMOS

The Motherboard and Its Components The motherboard is the nerve center of a computer, much like the spinal cord is the nerve center of the human body. It's a network of wires and controlling circuits that connects all computer components, including the CPU, RAM, disk drives, and I/O devices, such as network interface cards. Some key components of a motherboard are labeled in Figure 1-1 and explained in Table 1-1.






Figure 1-1 A PC motherboard

Courtesy of Course Technology/Cengage Learning

Table 1-1 Key components of a motherboard

Component	Description
CPU socket	The CPU is installed in this socket.
PCI bus expansion slots	Used to add functionality to a PC by adding expansion cards that have a Peripheral Component Interconnect (PCI) connector.
PCI-Express expansion slots	PCI-Express supersedes PCI and supports faster data transfer speeds. The larger slots are suitable for high-performance expansion cards, such as graphics cards and disk controllers. The smaller slots are best suited to sound cards and network interface cards.
RAM slots	Slots for installing RAM on the motherboard.
Chipset with heat sinks	The chipset consists of two chips referred to as the Northbridge and the Southbridge. These chips control data transfers between memory, expansion slots, I/O devices, and the CPU. The heat sink sits on top of the chipset to prevent it from overheating.
SATA connectors	Used for connecting hard drives and CD/DVD drives that use the Serial AT Attachment (SATA) specification.

(continues)

Table 1-1 Key components of a motherboard (continued)

Component	Description
IDE connector	Used for connecting Integrated Drive Electronics (IDE) hard drives and CD/DVD-ROM drives. Most systems now use SATA for hard drives and IDE for CD/DVD drives.
Main power connector	This connector is where the motherboard receives power from the system power supply.

All data that goes into or comes out of a computer goes through the motherboard because all storage and I/O devices are connected to the motherboard, as is the CPU, which processes data going in and coming out of a computer.

Computer Bus Fundamentals Table 1-1 mentions PCI bus expansion slots as a component of a motherboard. So what is a bus? A **bus** is a collection of wires carrying data from one place to another on the computer. There are many bus designs and formats, each designed for a particular purpose. Although bus types come and go, it's safe to say that replacements for an older bus design will almost certainly be faster than their predecessor.

In a computer, there are buses between the CPU and RAM, between the CPU and disk drives, and between the CPU and expansion slots, among others. For the purposes of this book, you're most interested in the bus connecting expansion slots to the motherboard because you usually connect a network interface card (NIC) into one of these slots. NIC installation and expansion slot bus types are discussed in Chapters 2 and 7. What you need to know now is that not all motherboards come with all types of expansion slots, and the faster and busier your computer is, the faster its bus type needs to be.

Hard Drive Fundamentals The hard drive is the primary long-term storage component on your computer. Hard drives consist of magnetic disks, called platters, that store data in the form of magnetic pulses. These magnetic pulses are maintained even when power is turned off. Each pulse represents a single bit of data.

The platters spin at extremely fast speeds, with some of the fastest disks having rotational speeds of 15,000 revolutions per minute (rpm). A read/write head is attached to an actuator arm that moves across the spinning platters in response to commands from the computer to read or write a file (see Figure 1-2). Generally, the faster the rotational speed, the better the hard drive performance is. When a file is requested to be written or read, its location is determined, and then the read/write heads are moved over the corresponding spot on the platter. After the platter spins to the file's starting location, the read/write heads are activated to read or write the data. The average amount of time platters take to spin into position is called the rotational delay or latency. The amount of time required to move read/write heads to the correct place is referred to as the seek time, and the time it takes to read or write data is called the transfer time. The average amount of time between the request to read or write data and the time the action is completed is referred to as the access time.



The terms used to measure hard drive performance aren't universal among manufacturers, but the terms used in the preceding paragraph represent most specifications.




Figure 1-2 Inside a hard drive

Courtesy of © 2010 Western Digital Technologies, Inc.

Hard disks store the documents you use with your computer as well as the applications that open these documents. In addition, the hard disk stores the OS your computer loads when it boots. As mentioned, the hard disk acts as an input device when files are read. When the computer boots, the OS files are read from the disk, and instructions in these files are processed by the CPU. However, the files don't go directly from the hard disk to the CPU; first, they're transferred to short-term storage (RAM).

RAM Fundamentals RAM, the main short-term storage component on your computer, consists of capacitors to store data and transistors to control access to data. Capacitors require power to maintain the bits they store. Because RAM requires continuous power to store data, it's referred to as “volatile memory.”

RAM has no moving parts, so as mentioned, accessing data in RAM is much faster than accessing data on a hard drive—there's no seek time or rotational delay. Because RAM is so much faster than a hard drive, any information the CPU processes should be in RAM. If data the CPU requires is located on the hard drive, it's loaded into RAM first, which takes considerable time. Therefore, the more RAM your system has, the more likely it is that all the data running programs need can be stored in RAM, making the system perform much faster.

BIOS/CMOS Fundamentals A key component of every computer is its basic input/output system (BIOS), which is a set of instructions located in a chip on the motherboard. A main function of the BIOS is to tell the CPU to perform certain tasks when power is first applied to the computer, including initializing motherboard hardware, performing a power-on self test (POST), and beginning the boot procedure.

Because of the complexity of motherboards, configuring some of their hardware components and tuning performance parameters are often necessary. When a computer begins to boot, the BIOS program offers the user an opportunity to run the Setup utility to perform this configuration. The configuration data the user enters is stored in complementary metal oxide semiconductor (CMOS) memory. It holds information such as on which devices the CPU should look for an OS to boot, the status of hardware devices, and even a system password, if needed. CMOS is a type of low-power memory that requires only a small battery to maintain its data. It's also referred to as nonvolatile memory because it doesn't require power from the computer's main power supply.

Computer Boot Procedure

The following six steps are necessary to take a computer from a powered-off state to running a current OS, such as Windows or Linux:

1. Power is applied to the motherboard.
2. The CPU starts.
3. The CPU carries out the BIOS startup routines, including the POST.
4. Boot devices, as specified in the BIOS configuration, are searched for an OS.
5. The OS is loaded into RAM.
6. OS services are started.

These steps apply to almost every type of computer, including very small computing devices, such as cell phones and iPods. Probably the biggest difference between computers is what occurs in the last step. OS services are programs that are part of the OS rather than applications a user starts. The particular services an OS starts can vary greatly, depending on which OS is loaded and how it's configured. The number and type of services started on a system are what, at least in part, account for the time it takes a system to boot completely. Examples of common OS services include the user interface, the file system, and, of course, networking services.



The projects in this book involving a Windows client OS use Windows 7 Enterprise Edition. Other editions of Windows 7 can be used, except Windows 7 Home Edition. Windows Vista can also be used, with some small changes to step-by-step instructions. Windows XP can be used in most cases but might require additional changes.



Hands-On Project 1-1: Examining a Computer's Boot Procedure

Time Required: 10 minutes

Objective: Examine the computer boot procedure and BIOS setup utility.

Required Tools/Equipment: Your classroom computer and access to the BIOS Setup utility

Description: In this project, you examine the computer boot procedure from beginning to end, using a Windows computer. You also examine the BIOS Setup utility and view the configuration that specifies which devices the BIOS should search for an OS. Because the BIOS is different for different computers, your instructor might have to assist with the specific keystrokes you enter to run the BIOS Setup utility and view the boot order menu. This project uses a

virtual machine and the BIOS Setup utility in VMware Workstation 6.x. If you aren't using virtual machines for the projects in this book, the BIOS on most computers is similar.



Your computer must be turned off before you begin this project. Read the first step carefully before turning on the computer, as you need to act quickly to enter the BIOS Setup utility.

1. Turn on your computer. Watch the screen carefully for a message telling you what key to press to activate the BIOS Setup utility. On many systems, this key is F1, F2, or Delete. If you don't press the key in time, the OS boots normally. If this happens, shut down the computer and try again.
2. When you have entered the BIOS Setup utility, your screen should look similar to Figure 1-3. Before continuing, write down the steps of the boot procedure that have taken place to this point:

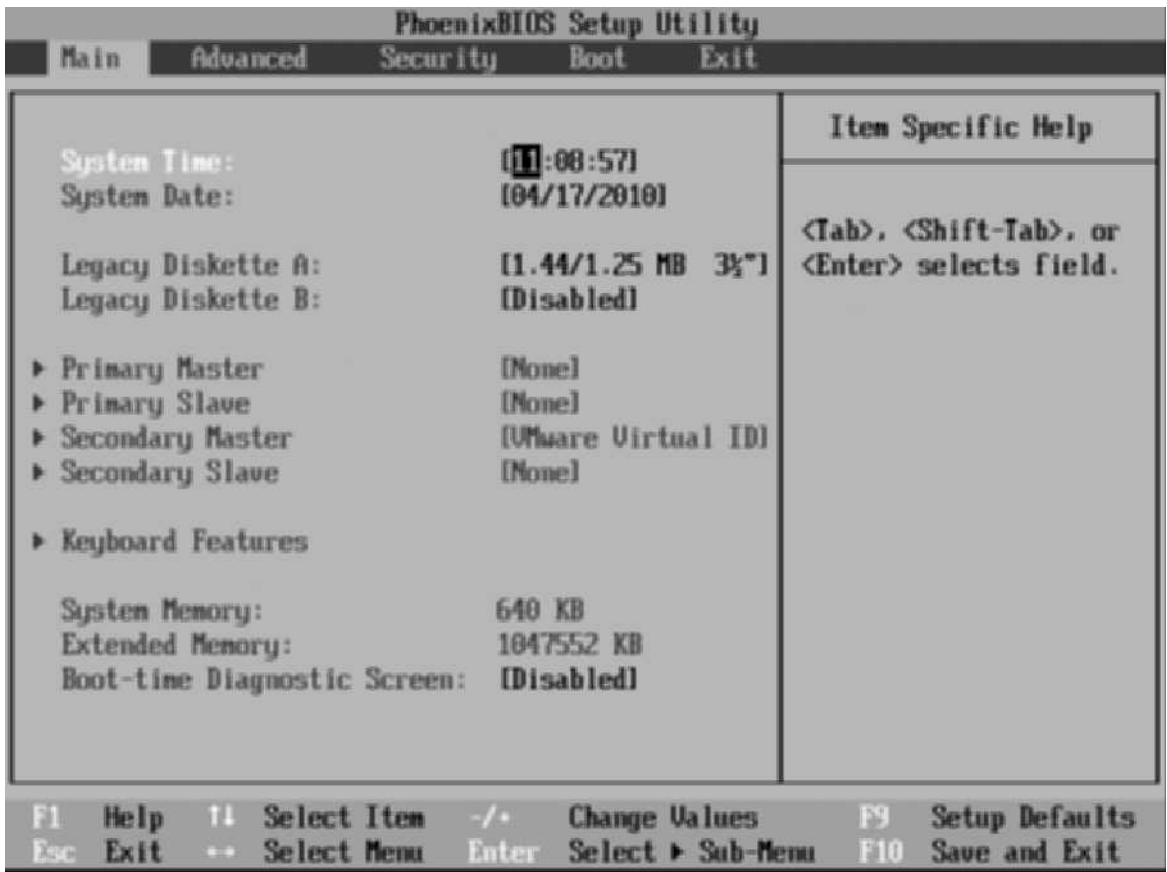


Figure 1-3 The BIOS Setup utility

Courtesy of Course Technology/Cengage Learning

- Navigate the BIOS Setup utility until you find the boot order menu (see Figure 1-4). From this menu, you can change the order in which the BIOS looks for boot devices, or you can exclude a device from the boot order. The BIOS boots from the first device in which it finds an OS. You might need to change the boot order if, for example, you have an OS installed on the hard drive but want to boot from an installation CD/DVD to install a new OS. In this case, you move the CD/DVD device to the first entry in the boot order.

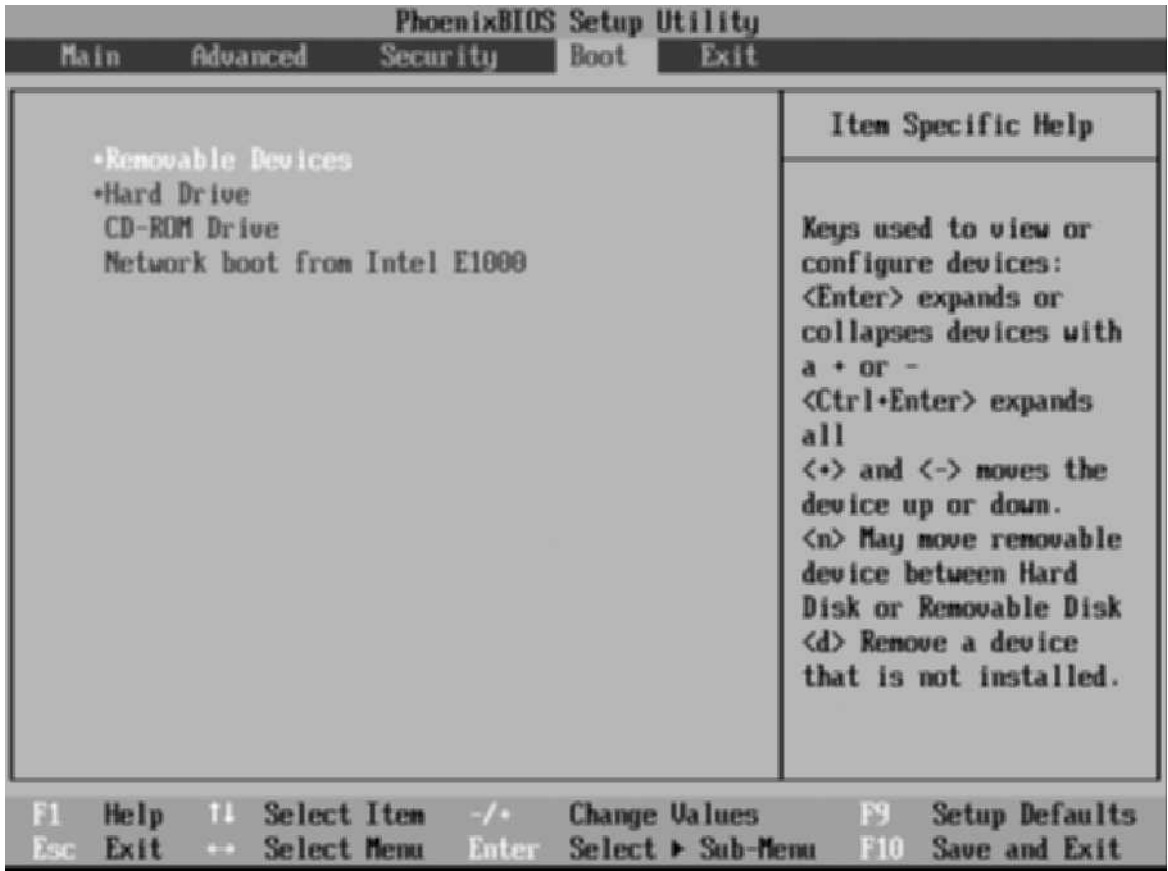


Figure 1-4 The BIOS boot order menu

Courtesy of Course Technology/Cengage Learning

- For now, you can leave the boot order as it is. To quit the Setup utility, press the correct key (usually specified at the bottom of the screen). In Figure 1-4, you press Esc to exit without saving changes or F10 to save the changes before exiting. In either case, when you exit, the computer restarts. Press the key for exiting without saving changes.
- Write the final steps of the boot procedure that occurred as Windows started:

- Shut down the computer for the next project.

How the Operating System and Hardware Work Together

A computer's OS provides a number of critical services, including a user interface, memory management, a file system, multitasking, and the interface to hardware devices. Without an OS, each application would have to provide these services, and if a user wanted to run multiple applications at once (multitasking), the applications would have to run cooperatively. In short, without an OS, computing would still be in the proverbial Stone Age. The following sections describe these services briefly, and Chapter 8 discusses OS components in more detail.

User Interface The user interface enables people to interact with computers. With graphical user interfaces (GUIs), users can point and click their way around the computer to run applications, access network services, manage hard drives and files, and configure the working environment to their liking. In short, users provide input, and the OS, along with the CPU, processes that input, whether it's mouse clicks or keystrokes, and generates output. Without a user interface, computers could process only information that has been programmed into memory or storage. If something went wrong, there would be no way to indicate the problem to a person, making a computer without a user interface of little value except when it has a narrowly defined task, such as running a piece of machinery.

Memory Management Computers are now equipped with memory measured in hundreds of megabytes or gigabytes, whereas in the early 1990s, the typical amount of memory was about 1 megabyte. Each application requires a certain amount of memory in which to run. When the OS loads an application, memory must be allocated for the application to run in, and when the application exits, the memory it was using must be marked as available. The OS handles these memory management tasks. Without a central memory manager, an application could use any memory in the system, and it might be memory already being used by a running application or the OS itself. If this happens, the system can crash or perform erratically. Today's OSs usually detect an application's attempt to access another process's memory and force the offending application to terminate.

File System The file system is used to organize space on storage devices, such as disk drives and flash drives, for the purpose of storing and locating files. Contemporary file systems typically have the following objectives:

- Provide a convenient interface for users and applications to open and save files.
- Provide an efficient method to organize space on a drive.
- Provide a hierarchical filing method to store files.
- Provide an indexing system for fast retrieval of files.
- Provide secure access to files by authorized users.

When a user double-clicks a file to open it, the user interface calls the file system with a request to open the file. The file type determines exactly how the file is opened. If the file is an application, the application is loaded into memory and run by the CPU. If the file is a document, the application associated with the document type is loaded into memory and opened by the application. For example, if you double-click the Budget.xls file, the Excel

application is loaded into memory and then opens the Budget.xls document file. If a user creates a new file or changes an existing file and wants to save it, the application calls the file system to store the new or changed file on the disk. Most users of an OS interact with the file system by using Windows Explorer or a similar file manager program on another OS, but as a future computer or network professional, you need to have a deeper understanding of how a file system works so that you can make informed choices when you need to install a file system or troubleshoot file system-related problems. You can find more discussion on this topic in Chapter 8.

Multitasking Quite simply, **multitasking** is an operating system's capability to run more than one application or process at a time. Multitasking is what allows you to listen to a music file while browsing the Web, for example. Computer hardware can't do that by itself. The OS is designed to look for applications that have some kind of work to do (such as load a new Web page or continue playing the current music file) and then schedule CPU time so that the work gets done. For example, if you're browsing the Web and reading the current page loaded in your Web browser, the computer isn't really doing any work.

However, if you click a link on the Web page, you're telling the Web browser you want to load a new page. The OS responds by telling the CPU to start executing the part of the Web browser application responsible for loading a new Web page. You might wonder how can you play a music file at the same time the CPU is loading a new Web page. There are two possible answers: The computer contains more than one CPU or a multicore CPU and can literally do two things at once (in this case, load a Web page and play a music file), or the OS instructs the CPU to switch between the two tasks rapidly, giving the illusion that they're happening simultaneously. Because CPUs can execute hundreds of millions of instructions per second, this illusion isn't difficult to carry off.

Interface to Hardware Devices When an application needs to communicate with computer hardware, as when writing information to the display device or sending data to the network, it calls on the OS, which then calls on a device driver. A **device driver** is software that provides the interface between the OS and computer hardware. The reason the application can't simply read or write data directly to hardware is that other applications might also need to communicate with the same device at the same time. If this were allowed to happen, it would be akin to two or more people on different extensions of the same land line trying to dial a different number. Nobody's phone call would go through, or one person might call an unintended destination. The OS queues up each request and sends it to the device driver when it's not busy. This procedure ensures that every application's request is taken care of in a nice orderly fashion.

Every device performing an input or output function requires a device driver. When an input device has data ready for processing, or when an output device is ready to accept data, the device must signal the OS. Most devices use a signal called an interrupt to let the OS know it has data ready to be read or is ready for more data to be written. Computers spend a considerable amount of time servicing interrupts on a busy computer. For example, when the mouse is moved or a key on the keyboard is pressed, an interrupt is generated so that the OS knows the mouse pointer must be redrawn onscreen or a character must be written to



the screen. On a networked computer, an interrupt is generated by the NIC when a packet arrives.

Every time an interrupt occurs, the OS must stop what it's doing to service the interrupt. It takes many instructions for an OS to stop what it's doing, service the interrupt, and then resume what it was doing before the interrupt occurred. Because computers can execute millions of instructions per second, users don't usually notice the interruption. If enough interrupts occur simultaneously and for a prolonged period, however, a system can become noticeably sluggish or even seem to freeze. Malfunctioning hardware and network errors that generate excessive packets are two of the many possible causes of this problem. Remember this idea about excessive interrupts caused by the NIC; it's an important point later when you learn about network protocols in Chapter 5.

Networking is, of course, the focus of this book, but your grasp of the fundamentals of computer components and operations will facilitate your understanding of networking components and operations.

The Fundamentals of Network Communication

A computer **network** consists of two or more computers connected by some kind of transmission medium, such as a cable or air waves. After they're connected, correctly configured computers can communicate with one another. The primary motivation for networking was the need for people to share resources, such as printers and hard drives, and information such as word-processing files and to communicate by using applications such as e-mail. These motivations remain, especially for businesses, but another motivating factor for networking for both businesses and homes is to get “online”—to access the Internet. The Internet, with its wealth of information, disinformation, fun, and games, has had a tremendous impact on how and why networks are used today. Indeed, many of the networking technologies used now that you learn about in this book were developed as a result of the Internet explosion.

You might know how to use a network already; in particular, you probably know how to use programs that access the Internet, such as Web browsers and e-mail programs. To understand *how* networks work, however, you need to learn about the underlying technologies and processes that are put into action when you open a Web browser or an e-mail program. A good place to start is with the components that make a stand-alone computer a networked computer.

Network Components

Imagine a computer with no networking components—no networking hardware, no networking software. It's hard to imagine in this age of seemingly everything and everybody being connected. However, not too long ago, when you bought a computer, its main purpose was to run applications such as word-processing and spreadsheet programs, not Web browsers and e-mail. In fact, the computer had neither the necessary hardware nor software to run these programs. These computers were called **stand-alone computers**. If you wanted to network such a computer, you had to add the necessary components:

- *Network interface card*—A NIC is an add-on card that’s plugged into a motherboard expansion slot and provides a connection between the computer and the network. Most computers now have a NIC built into the motherboard, so no additional card is necessary. NICs are discussed in more detail in Chapter 2.
- *Network medium*—A cable that plugs into the NIC and makes the connection between a computer and the rest of the network. In the simplest of networks, with just two computers, the other end of the cable can plug into the second computer’s NIC. More likely, the other end of the cable plugs into an interconnecting device that accommodates several computer connections. Network media can also be the air waves, as in wireless networks. In this case, the connection is between the antenna on the NIC and the antenna on another NIC or interconnecting device. Network media are discussed in more detail in Chapter 4.
- *Interconnecting device*—Although this component isn’t always necessary because two computers can be connected directly with a cable and small wireless networks can be configured without an interconnecting device, most networks include one or several of these components. Interconnecting devices allow two or more computers to communicate on the network without having to be connected directly to one another. They include switches, hubs, routers, and wireless access points, all discussed in Chapter 2. A small network connected to a switch is shown in Figure 1-5.




Figure 1-5 A network of computers connected to a switch

Courtesy of Course Technology/Cengage Learning



Hands-On Project 1-2: Upgrading a Stand-alone Computer to a Networked Computer

Time Required: 30 minutes

Objective: Upgrade a stand-alone computer to a networked computer.

Required Tools/Equipment: Lab computers (as specified in the book’s lab setup instructions), NICs, patch cables, and hub or switch.

Description: In this project, you install a NIC and connect it to an interconnecting device with a cable. This project can be done in groups or as an instructor demonstration. It's intended only to familiarize you with the hardware components needed to make a stand-alone computer a networked computer.

1. Install the NIC, following the steps your instructor provides. This process might involve opening the computer case or simply plugging a USB NIC into a USB slot.
2. Turn on the computer. If necessary, insert a disk containing the NIC driver and follow the instructions for installing it.
3. Using the supplied cable, plug one end into the NIC and the other end into the interconnecting device, which should be a hub or a switch.
4. Examine the indicator lights on the NIC and the hub or switch. There might be one or two lights on each port of the device, depending on its features. There's at least one indicator on the NIC and on each port of the hub or switch that's usually referred to as a "link light." The link light glows when a data connection has been made between the NIC and the hub or switch. Your instructor can supply more details about the indicator lights available on your hub or switch. List the status of indicators on the NIC and the hub or switch port into which the NIC is plugged:

-
5. Shut down the computer and unplug and put away the cables.

The previous list of components satisfies the hardware components needed to make a stand-alone computer a networked computer. The computer must also have the necessary software to interact with network hardware and communicate with other computers on the network. Network software transforms a stand-alone OS into a network OS. It's the software that allows a word-processing program to print to a networked printer or open a document on a server or knows how to request a Web page or send an e-mail. It's also the software that communicates between the OS and network hardware. Network software can be divided into the following categories:

- *Network clients and servers*—**Network client software** requests information that's stored on another network computer or device. **Network server software** allows a computer to share its resources by fielding resource requests generated by network clients. Network client software can be an integral part of well-known applications, such as Web browsers and e-mail programs. A Web browser, for example, sends a request for a Web page to a Web server. Network client software can also run in the background, usually installed as a networking service. In this case, it enables programs without built-in client software to access shared network resources on other computers. For example, Client for Microsoft Networks, which is installed automatically in Windows, allows a word processor to open a file that's shared on another Windows computer or print to a printer attached to another Windows computer. In this setup, the server software called File and Printer Sharing for Microsoft Networks receives the request the client generates and provides access to the shared file or printer. When network clients and servers need to send information on the network, they must pass it to network protocols.
- *Protocols*—**Network protocols** define the rules and formats a computer must use when sending information across the network. A network protocol can be likened

to a human language. Just as two people who want to communicate must speak the same language, two computers that want to communicate must use the same protocol. Examples of network protocols include TCP/IP and IPX/SPX. Network protocols do all the behind-the-scenes tasks required to make networking work. Most of the complexity in networking is handled by these protocols, and they're discussed in depth in Chapter 5. After a network protocol has formatted the message correctly, it hands the data off to the NIC device driver for transmission onto the network.



The term “NIC device driver” is often shortened to simply “NIC driver,” which is the term used throughout this book.

- *NIC driver*—NIC drivers receive data from protocols and then forward this data to the physical NIC, which transmits data onto the medium. The reverse is also true. When data arrives at the NIC from the medium, the NIC hands it off to the NIC driver, which then hands it off to network protocols. Every NIC card installed in a computer must have an associated device driver installed in the OS. The device driver software manages the details of communicating with the NIC hardware to send and receive data to and from network media.

Each of these software components plays a role in the steps of network communication, described in the next section.

Steps of Network Communication

Most network communications start by a user needing to access a resource on another computer, such as a Web server or file server. A user's attempt to access network resources is summarized in these basic steps:

1. An application tries to access a network resource by attempting to send a message to it.
2. Network client software detects the attempt to access the network. Client software formats the message generated by the application and passes the message on to the network protocol.
3. The protocol packages the message in a format suitable for the network and sends it to the NIC driver.
4. The NIC driver sends the data in the request to the NIC card, which converts it into the necessary signals to be transmitted across the network medium.

Remember that there are two sides to a communication session, and most of them involve a client trying to access network resources and a server providing those resources. The steps taken on the server side are essentially the reverse of those on the client side:

1. The NIC card on the server receives signals from the network medium and converts them into message data, which is read by the NIC driver.
2. The NIC driver passes the message to the network protocol.

3. The network protocol determines which server software the message is targeting and passes the message to this designated software. Remember that a computer can have many clients and many servers running at the same time. For example, a computer running Windows Server 2008 might be acting as a mail server and a file server. Each server function requires different server software.
4. The server software receives the message and responds by sending the requested data to the client computer, using the four steps outlined previously.

Layers of the Network Communication Process

Each step of a client accessing network resources is often referred to as a “layer” in the network communication process. Each layer has a specific function to accomplish, and all the layers work together. Figure 1-6 depicts this process, and Simulation 1 on the book’s CD shows an animation of this process. Keep in mind that the steps outlined previously simplified the communication process, which is one reason the layered approach is so effective: Complex concepts can be described in simple steps. Chapter 6 discusses the layered approach to networking in more detail, and Chapter 5 explains the role of protocols in network communication.



Simulation 1: Layers of the network communication process

As you’ll see in Chapter 6 when the OSI model of networking is discussed, the layers are given different names and divided into additional pieces. What’s important now is grasping the idea of a layered approach, in which a complex process is broken into manageable steps, each with a specific role to play. Table 1-2 maps the resource access steps listed previously to the four layers in Figure 1-6.

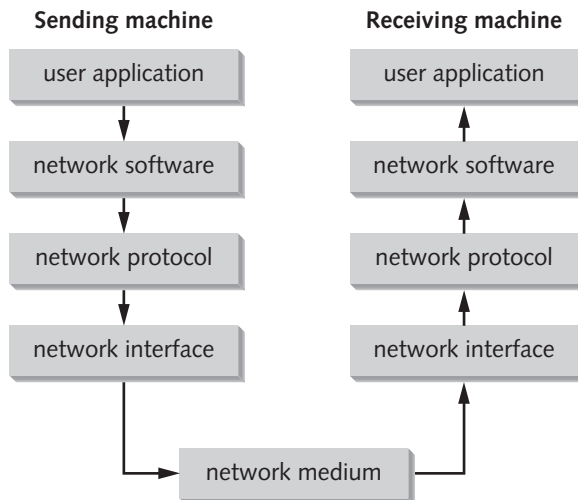


Figure 1-6 Layers of the network communication process

Courtesy of Course Technology/Cengage Learning

Table 1-2 Layers of the network communication process

Step	Description	Layer
1	An application tries to access a network resource.	User application
2	Client software detects the attempt to access the network and passes the message on to the network protocol.	Network software
3	The protocol packages the message in a format suitable for the network and sends it to the NIC driver.	Network protocol
4	The NIC driver sends the data in the request to the NIC, which converts it into the necessary signals to be transmitted across the network medium.	Network interface

How Two Computers Communicate on a LAN: Some Details

The layers of the network communication process give an overview of how network communication works. However, there are few details on what each layer accomplishes. This discussion focuses on computer addresses and how they're used during network communication.

In a network using a protocol such as TCP/IP (the most common protocol used on networks), computers have two addresses: a logical address and a physical address. With TCP/IP, the logical address is the IP address, and the physical address is called the Media Access Control (MAC) address. You can look at these two addresses much like the addresses used to send mail through the postal system. When a letter is mailed in the United States, it requires a street address and a zip code. The zip code gets the letter to the correct region of the country, and the street address gets the letter to the correct home or business.



The MAC address is stored as part of the NIC, and no two MAC addresses in the world should be the same.

You can liken the zip code to the logical or IP address and the street address to the physical or MAC address. When a message is sent on a network, the IP address is used to get the message to the correct network, and the MAC address is used to get the message to the correct computer on this network. If the sender and receiver are on the same network, the IP address in the message is used primarily as a means to ascertain the destination computer's MAC address.

For example, Figure 1-7 shows two computers connected to a switch. Computer A wants to communicate with Computer B. One of the simplest forms of communication is a ping. The ping command sends a message from one computer to another, essentially asking the other computer whether it's listening on the network. If a computer receives a ping, it replies so that the sending computer knows the message was received. It's like the commercial with a cell phone user asking "Can you hear me now?" Following are the steps of this communication process:

1. A user at Computer A types `ping 10.1.1.2` at a command prompt.
2. The network software creates a ping message.
3. The network protocol packages the message by adding IP addresses of the sending and destination computers and acquires the destination computer's MAC address.

4. The network interface software adds MAC addresses of the sending and destination computers and sends the message to the network medium as bits.
5. Computer B receives the message, verifies that the addresses are correct, and then sends a reply to Computer A, using Steps 2 through 4.




Figure 1-7 Communication between two computers

Courtesy of Course Technology/Cengage Learning

Simulation 2 on the book's CD shows this communication process in action.



Simulation 2: Communication between two computers

In most cases, users don't usually initiate network communication by using a computer's IP address; instead, they use the computer name. However, just as you can't mail a letter with only the recipient's name, you can't communicate over a network with only the computer's name. You certainly know the name of the person you're writing to, but you might have to look up his or her address in your address book before you can address the envelope. Similarly, computers use an address book of sorts, called a **name server**, to get the IP address of a computer, given its name. TCP/IP provides name server functions through its Domain Name System (DNS, discussed in more detail in Chapter 5). With this information in mind, the preceding steps can be expanded as follows:

1. A user at Computer A types `ping Computer B` at a command prompt.
2. A name lookup is performed to retrieve Computer B's IP address.

3. The network software creates a ping message.
4. The network protocol packages the message by adding IP addresses of the sending and destination computer and acquires the destination computer's MAC address.
5. The network interface software adds MAC addresses of the sending and destination computers and sends the message to the network medium as bits.
6. Computer B receives the message, verifies that the addresses are correct, and then sends a reply to Computer A, using Steps 3 through 5.

Quite a few details in some of these steps have been left out for now, but they're expanded on in the TCP/IP discussion in Chapter 6. Now that you have a solid idea of how network communication takes place and how networks are depicted in figures, you can learn some common terms for describing networks and network components in the next section. Along the way, you'll see more figures of different types of networks.



Student computers should be named Net-XX, with XX representing a two-digit number assigned to the student. Wherever you see XX in an activity step, substitute your student number.



Hands-On Project 1-3: Viewing Network Software Layers

Time Required: 10 minutes

Objective: View the properties of your computer's network connection and identify the layers of the network communication process.

Required Tools/Equipment: Your classroom computer and a user account with administrative access named "NetAdmin" with the password "Password01"

Description: In this project, you view the properties of your computer's local area connection and identify the layers of the network communication process. Each network connection in Windows contains the software responsible for the steps of the network communication process.

1. Start your computer and log on as **NetAdmin**, if necessary.
2. Open the Network and Sharing Center by clicking **Start, Control Panel**. Under Network and Internet, click **View network status and tasks**.
3. In the left pane of the Network and Sharing Center, click **Change adapter settings**. Right-click **Local Area Connection** and click **Properties** to open the Local Area Connection Properties dialog box (see Figure 1-8).
4. The Connect using text box displays the NIC. In the list box under it, you see several items. Client for Microsoft Networks, File and Printer Sharing for Microsoft Networks, and Internet Protocol Version 4 are the items you're most interested in right now, as they're the most necessary software components to make network communication work.



Figure 1-8 The Local Area Connection Properties dialog box

Courtesy of Course Technology/Cengage Learning

5. Assume a user is running a word-processing program and saves a file to a Windows server. Use the information you learned in this chapter to write which of the four layers of the network communication process each component corresponds to. Note that some layers can be used more than once.
 - Word-processing program: _____
 - NIC displayed in the Connect using text box: _____
 - Client for Microsoft Networks: _____
 - File and Printer Sharing for Microsoft Networks: _____
 - Internet Protocol Version 4: _____
6. Close all open windows, but leave your computer running for the next project.



Hands-On Project 1-4: Using Ipconfig, Ping, and ARP

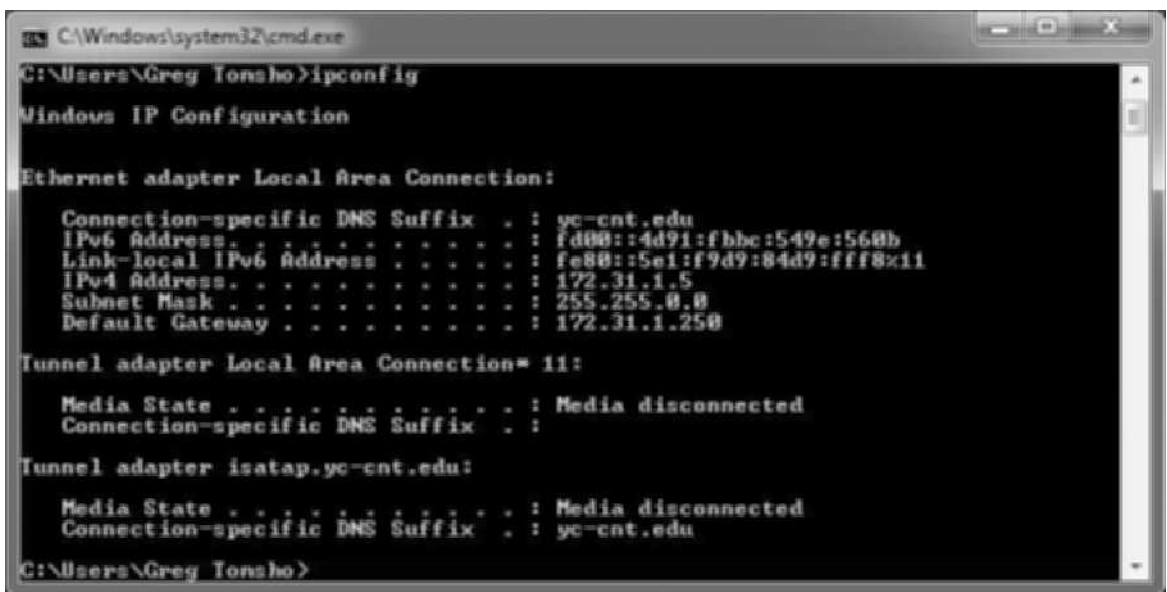
Time Required: 15 minutes

Objective: Use Ipconfig, Ping, and ARP to view and test network addresses and connectivity.

Required Tools/Equipment: Your classroom computer

Description: In this project, you use command-line tools to view your network configuration and test your computer's capability to communicate with other computers. Ipconfig displays the IP address configuration of your network interfaces. Ping sends a message to a computer to verify the capability to communicate with it. ARP displays the MAC (physical) addresses your computer has discovered.

1. Start your computer and log on as **NetAdmin**, if necessary.
2. Click **Start**, type **cmd**, and press **Enter** to open a command prompt window. At the command prompt, type **ipconfig** and press **Enter**. You should see a screen similar to Figure 1-9, although the specific numbers you see will vary. Ipconfig lists the IP address configuration for your network interfaces as well as other network settings.



```
C:\Windows\system32\cmd.exe
C:\Users\Greg Tonsho>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : yc-ent.edu
    IPv6 Address . . . . .           : fd00::4d91:fbbc:549e:568b
    Link-local IPv6 Address . . . . . : fe80::5e1:f9d9:84d9:fff8::11
    IPv4 Address. . . . .            : 172.31.1.5
    Subnet Mask . . . . .            : 255.255.0.0
    Default Gateway . . . . .         : 172.31.1.250

Tunnel adapter Local Area Connection* 11:

    Media State . . . . .            : Media disconnected
    Connection-specific DNS Suffix  . :

Tunnel adapter isatap.yc-ent.edu:

    Media State . . . . .            : Media disconnected
    Connection-specific DNS Suffix  . : yc-ent.edu

C:\Users\Greg Tonsho>
```

Figure 1-9 The ipconfig command output

Courtesy of Course Technology/Cengage Learning

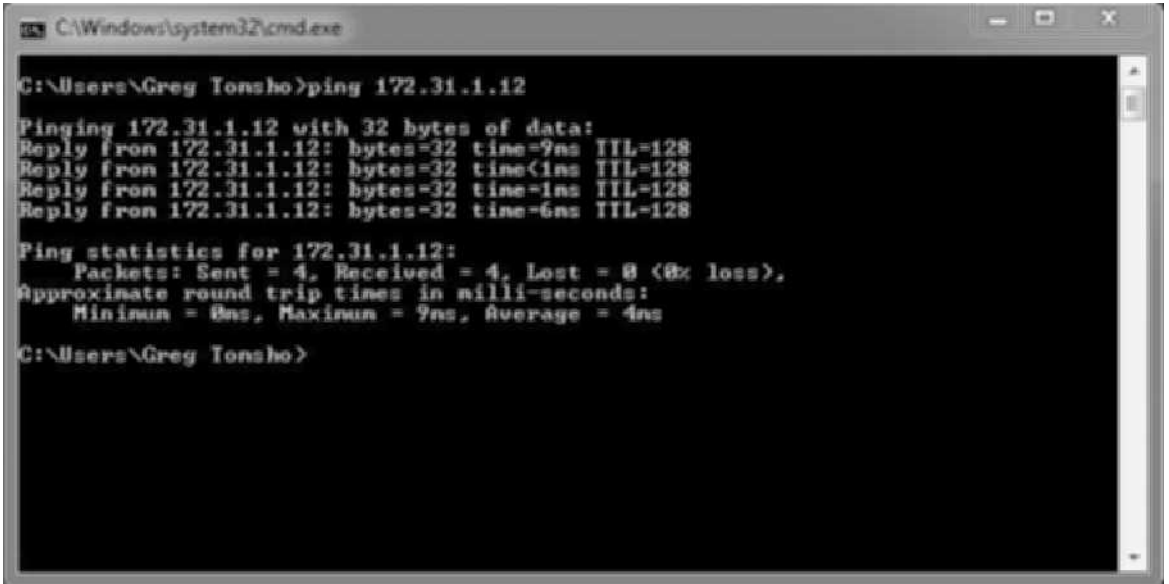
3. To see more details about your network configuration, type **ipconfig /all** and press **Enter**. You can scroll up the command prompt window to see all the output. Under the heading “Ethernet adapter Local Area Connection,” find the row labeled Physical Address (see Figure 1-10). The number you see in this row is the MAC address, a 12-digit hexadecimal value. Also, find the IP address in the IPv4 Address row. Write down these two addresses:



Figure 1-10 Using ipconfig /all to list physical (MAC) and IP addresses

Courtesy of Course Technology/Cengage Learning

4. Tell your partner what your IP address is and make note of your partner’s IP address. At the command prompt, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with your partner’s IP address). You should see output similar to Figure 1-11.
5. Remember that your computer needs both the destination IP address and MAC address to communicate with another computer. You supplied the IP address by typing it at the command prompt. Your computer discovered the MAC address of your partner’s computer by using Address Resolution Protocol (ARP). To see this address, type **arp -a** and press **Enter**. The output should be similar to Figure 1-12. You might see more lines of output, depending on what other devices your computer has been communicating with. ARP is discussed in more detail in Chapter 5, but for now, just know that it works automatically without user intervention.
6. Use the **ping** command to communicate with other computers and devices on your network, and use **ipconfig /all** to find the addresses of your default gateway (a router in your network) and your DNS servers. Write the MAC addresses of your default gateway and your DNS servers:
 - Default gateway: _____
 - DNS servers: _____
7. Close all open windows, but leave your computer running for the next project.



```
C:\Windows\system32\cmd.exe

C:\Users\Greg Tonsho>ping 172.31.1.12

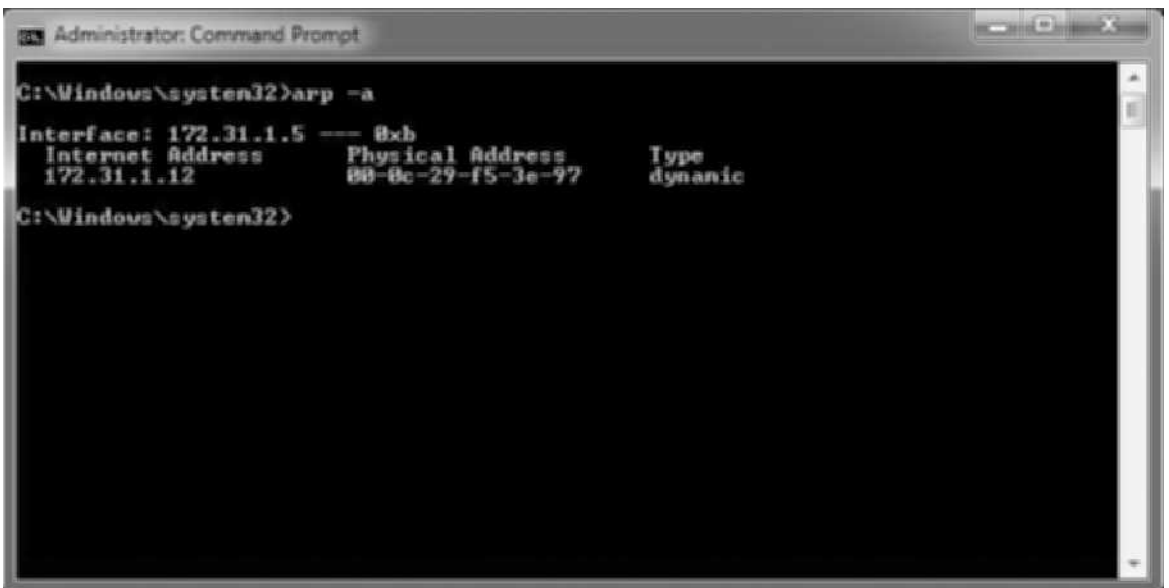
Pinging 172.31.1.12 with 32 bytes of data:
Reply from 172.31.1.12: bytes=32 time=9ms TTL=128
Reply from 172.31.1.12: bytes=32 time<1ms TTL=128
Reply from 172.31.1.12: bytes=32 time=1ms TTL=128
Reply from 172.31.1.12: bytes=32 time=6ms TTL=128

Ping statistics for 172.31.1.12:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 6ms, Maximum = 9ms, Average = 4ms

C:\Users\Greg Tonsho>
```

Figure 1-11 Results of the ping command

Courtesy of Course Technology/Cengage Learning



```
Administrator: Command Prompt

C:\Windows\system32>arp -a

Interface: 172.31.1.5 --- 8xb
Internet Address      Physical Address      Type
172.31.1.12           00-0c-29-f5-3e-97    dynamic

C:\Windows\system32>
```

Figure 1-12 The arp -a command displays MAC addresses

Courtesy of Course Technology/Cengage Learning

Network Terms Explained

Every profession has its own language with its own unique terms and acronyms. Learning this language is half the battle of becoming proficient in a profession, and it's no different in computer and networking technology. The following sections explain some common terms used in discussing computer networks. Because some of these terms are often associated with network diagrams, a number of figures are included in the following sections to show different ways of depicting networks.

LANs, Internetworks, WANs, and MANs

A small network, limited to a single collection of machines and connected by one or more interconnecting devices in a small geographic area, is called a **local area network (LAN)**. LANs also form the building blocks for constructing larger networks called “internetworks.” In Figure 1-13, the computers in a LAN are interconnected by a hub. LANs are represented in other ways, as in Figure 1-14; note the different symbols for a hub and a switch. Figure 1-15 shows a logical depiction of the same network; a logical depiction of a network leaves out details such as interconnecting devices, showing only the computers making up the network.




Figure 1-13 A LAN with computers interconnected by a hub

Courtesy of Course Technology/Cengage Learning

An **internetwork** is a networked collection of LANs tied together by devices such as routers, discussed in Chapters 2 and 7. Figure 1-16 shows two LANs interconnected by a router (represented by the standard symbol). Internetworks are usually created for these reasons:

- Two or more groups of users and their computers should be logically separated on the network yet still allow the groups to communicate. For example, in a school, you might want to logically separate the LAN containing student computers from the LAN containing faculty computers. Routers provide this logical separation but still allow communication between groups, as you see in Chapter 2.




Figure 1-14 A LAN with a symbolic hub (left) and a symbolic switch (right)

Courtesy of Course Technology/Cengage Learning




Figure 1-15 A logical depiction of a LAN

Courtesy of Course Technology/Cengage Learning

- The number of computers in a single LAN has grown to the point that network communication is no longer efficient. The nature of certain network protocols and devices make network communication increasingly less efficient as the number of computers on a LAN grows. Routers can be used to separate the computers into two or more smaller LANs, thereby increasing communication efficiency.
- The distance between two groups of computers exceeds the capabilities of most LAN devices, such as hubs and switches. This problem can happen, for example, when a company has multiple buildings or multiple floors in a building. Routers are often used to communicate between groups of computers that are separated geographically.




Figure 1-16 An internetwork with two LANs connected by a router

Courtesy of Course Technology/Cengage Learning

You might not realize it, but the computer you have at home is probably part of an internetwork. Every time you go online to browse the Web or check your e-mail, your computer (or LAN, if you have a home network) becomes part of the largest internetwork in the world: the Internet.

As a network's scope expands to encompass LANs in geographically dispersed locations, internetworks become classified as **wide area networks (WANs)**. A WAN spans distances measured in miles and links two or more separate LANs. WANs use the services of third-party communication providers, such as phone companies, to carry network traffic from one location to another. So although both internetworks and WANs connect LANs, the difference lies mainly in the LANs' proximity to each other and the technologies used to communicate between LANs. Therefore, the Internet is both an internetwork and, because it spans the globe, a very large WAN.

Occasionally, you might encounter a network type called a **metropolitan area network (MAN)**. Essentially, MANs use WAN technologies to interconnect LANs in a specific geographic region, such as a county or city. It's not uncommon to find large, complex networks involving all four network types: LANs and internetworks for purely local access, MANs for regional or citywide access, and WANs for access to remote sites elsewhere in the country or around the world. Take, for example, a nationwide bank. The main branch in a large

city has a building with multiple floors and hundreds of computers. Each floor constitutes a LAN, and these LANs are connected to form an internetwork. The internetwork at the main branch is then connected to other branches throughout the city to form a MAN. In addition, the main branch is connected to other branches in other cities and states to form a WAN.

In network drawings, WANs are often depicted with a jagged or thunderbolt-shaped line representing the connection between two devices, usually routers, and the Internet is usually represented as a cloud. A cloud is used to obscure the details of a large network, as if to say, “There’s some collection of networks and network devices, but the details aren’t important.” Figure 1-17 shows a WAN connection between two routers with a connection to the Internet. A grouping of three computers is often used to represent multiple computers on a LAN when the exact number doesn’t matter.




Figure 1-17 A WAN with a connection to the Internet

Courtesy of Course Technology/Cengage Learning

MANs have recently received a boost in popularity because of the growing trend in some major cities to implement a citywide wireless network. In these cases, wireless networking is possible in almost any part of the city, allowing users to stay connected anywhere. Some wireless technologies that make this type of networking possible are discussed in Chapter 4. Figure 1-18 shows a typical drawing of a wireless network.

Packets and Frames

When computers transfer information across a network, they do so in short bursts of about 1500 bytes of data. Each burst, or chunk, of data has the same basic structure; specifically, each chunk of data contains the MAC addresses and IP addresses of both the sending (source) and receiving (destination) computers. So to transfer a small word-processing file, only one burst of data transfer might be needed, but large photo or music files are first




Figure 1-18 A wireless LAN

Courtesy of Course Technology/Cengage Learning

broken into several hundred or even thousands of chunks before they're transferred. After each chunk of data is sent, the computer pauses momentarily. Data is transferred in this way for a number of reasons:

- The pause between bursts might be necessary to allow other computers to transfer data during pauses.
- The pause allows the receiving computer to process received data, such as writing it to disk.
- The pause allows the receiving computer to receive data from other computers at the same time.
- The pause gives the sending computer an opportunity to receive data from other computers and perform other processing tasks.
- If an error occurs during transmission of a large file, only the chunks of data involved in the error have to be sent again, not the entire file.

To use another analogy, you can look at chunks of data as sentences people use when speaking. Pauses in conversation give listeners an opportunity to register what has been said and possibly get a word in themselves.



TIP

To get an idea of how many chunks of data are involved in transferring a typical file, a 3-minute music file is about 3 million bytes (3 MB) of data, which takes about 2000 chunks of data.

Packets The chunks of data sent across the network are usually called packets or frames. **Packet** is the more well-known term and is often used generically to mean a chunk of data sent over the network. However, the term “packet” does have a particular meaning: It’s a

chunk of data with source and destination IP addresses (as well as other IP protocol information) added to it. Figure 1-19 shows a representation of the original data to be transferred, and Figure 1-20 shows the packets created after the data has been broken into chunks and IP addresses added.



Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies, purus lectus malesuada libero, sit amet commodo magna eros quis urna.

Nunc viverra imperdiet enim. Fusce est. Vivamus a tellus.

Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci.

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies, purus lectus malesuada libero, sit amet commodo magna eros quis urna.

Nunc viverra imperdiet enim. Fusce est. Vivamus a tellus.

Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci.

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies, purus lectus malesuada libero, sit amet commodo magna eros quis urna.

Nunc viverra imperdiet enim. Fusce est. Vivamus a tellus.

Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci.

Figure 1-19 Original data

Courtesy of Course Technology/Cengage Learning

Dest: IP: 172.16.1.2, Source IP: 172.16.1.1	Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies.
Dest: IP: 172.16.1.2, Source IP: 172.16.1.1	purus lectus malesuada libero, sit amet commodo magna eros quis urna. Nunc viverra imperdiet enim. Fusce est. Vivamus a tellus.
Dest: IP: 172.16.1.2, Source IP: 172.16.1.1	Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci
Dest: IP: 172.16.1.2, Source IP: 172.16.1.1	Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies
Dest: IP: 172.16.1.2, Source IP: 172.16.1.1	Pellentesque habitant tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci.

Figure 1-20 Data broken into several packets

Courtesy of Course Technology/Cengage Learning

Using the U.S. mail analogy, you can look at a packet as an envelope that has had the zip code added to the address but not the street address. In relation to the layers of the network communication process, packets are generated by and processed by the network protocol. You learn more details about this process in Chapters 5 and 6.

Frames A **frame** is a packet with the source and destination MAC addresses added to it. In addition, frames have an error-checking code added to the back end of the packet, which is why they're called frames. The packet is "framed" by MAC addresses (and other network interface information) on one end and an error-checking code on the other. A frame is like a letter that's been addressed and stamped and is ready to deliver.

Frames are essentially the final state of data before it gets placed on the network medium as bits. The network interface is the layer of the network communication process that works with frames. Figure 1-21 shows what the packets from Figure 1-20 would look like after the frame information is added.

Dest MAC, Source MAC	Dest IP, Source IP	Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies,	Error check
Dest MAC, Source MAC	Dest IP, Source IP	purus lectus malesuada libero, sit amet commodo magna eros quis urna. Nunc viverra imperdiet enim. Fusce est. Vivamus a tellus.	Error check
Dest MAC, Source MAC	Dest IP, Source IP	Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci	Error check
Dest MAC, Source MAC	Dest IP, Source IP	Lorem ipsum dolor sit amet, consectetur adipiscing elit. Maecenas porttitor congue massa. Fusce posuere, magna sed pulvinar ultricies	Error check
Dest MAC, Source MAC	Dest IP, Source IP	Pellentesque habitant tristique senectus et netus et malesuada fames ac turpis egestas. Proin pharetra nonummy pede. Mauris et orci.	Error check

Figure 1-21 The packets are now frames and ready for delivery

Courtesy of Course Technology/Cengage Learning

The process of adding IP addresses and then MAC addresses to chunks of data is called **encapsulation**. Information added at the front of data is called a **header**, and information added at the end of data is called a **trailer**. Data is encapsulated several times as it works its way down from the sending application until it makes it to the network interface as a frame. When the destination computer receives the frame, the process is reversed as the network interface deencapsulates (has the header and trailer removed) the frame so that it becomes a packet again. This process continues until the packet arrives at the receiving application or service as the original data. This process is all part of the layered approach to networking.

Clients and Servers

You've already learned about the role of client network software and server network software. Unfortunately, the world of networking sometimes uses the same terms to discuss two different things. The following sections attempt to clarify what these terms mean and how their meaning can differ depending on how they're used.

Client A **client**, in networking terms, can be a workstation running a client OS, such as Windows 7, or as you have seen, it can also refer to the network software on a computer that requests network resources from a server. Further, you can refer to a physical computer as a client computer. What the term “client” means, therefore, depends on the context in which it’s used. To clarify, it’s usually used in these three contexts:

- *Client operating system*—The OS installed on a computer is designed mainly to access network resources, even though it might be capable of sharing its own resources. Windows XP, Windows Vista, Windows 7, and Mac OS X all fit this description, as do certain distributions of Linux. A client OS is also frequently referred to as a “desktop OS.”
- *Client computer*—This computer’s primary role in the network is to run user applications and access network resources. Most computers in a network fit this description.
- *Client software*—It’s the software that requests network resources from server software running on another computer. For example, a Web browser, an e-mail client (such as Microsoft Outlook), and Client for Microsoft Networks fit into this category.

Server When most people hear the word “server,” they conjure up visions of a large tower computer with lots of hard drives and memory. This image is merely a computer hardware configuration that may or may not be used as a server, however. In short, a computer becomes a **server** when software is installed on it that provides a network service to client computers. In other words, you can install certain software on an inexpensive laptop computer and make it act as a server. By the same token, a huge tower computer with six hard drives and 16 GB of RAM can be used as a workstation for a single user. So although some computer hardware configurations are packaged to function as a server, and others are packaged as client or desktop computers, what makes a computer a server is the software installed on it. Just as there are three contexts in which the term “client” is used, so it is with the term “server”:

- *Server operating system*—This term is used when the OS installed on a computer is designed mainly to share network resources and provide other network services, some of which are discussed later in “Network Servers.” A server OS is tuned to be able to share files efficiently and perform network operations in response to client requests, even though the OS might also be able to run user applications and client software. Windows Server 2008, Mac OS X Server, UNIX, and many Linux distributions fit this description.
- *Server computer*—This term is used when a computer’s primary role in the network is to give client computers access to network resources and services. The computers that most often fit this description are usually found in the IT computer room or locked away in a closet.
- *Server software*—It’s the software that responds to requests for network resources from client software running on another computer. A Web server, such as Internet Information Services (IIS), an e-mail server, such as Microsoft Exchange, and File and Printer Sharing for Microsoft Networks fit into this category.



Microsoft refers to server software components as “services.” Other operating systems use other terms; for example, in Linux/UNIX, server software components are referred to as “daemons.”



As you can see, the lines between a client computer and a server computer are often blurred because OSs are now designed as network operating systems, and most can take on both the roles of server and client. As you're learning, however, the language of networking is often imprecise, and you must pay attention to the context in which networking terms are used to grasp their meaning. As you get more comfortable with all the terms and better understand how networks work, the nuances of the terminology will fall into place.

Network Models

A **network model** defines how and where resources are shared and how access to these resources is regulated. Networks models fall into two major types: peer-to-peer and server-based (also called client/server). This discussion of network models addresses the role that computers play on the network and how these roles interact. Server-based networks are the most common in business settings, but understanding both types is essential, especially as they compare with one another.



Peer-to-peer networks running Windows operating systems are referred to as "workgroup networks," and server-based networks running Windows Server are referred to as "domain-based networks."

In a **peer-to-peer network**, most computers function as clients or servers, as circumstances dictate. For example, a computer can act as a server by sharing a printer it's connected to and simultaneously act as a client by accessing a file shared by another computer on the network. In this type of network, there's no centralized control over who has access to network resources; each user of a computer maintains control over his or her own shared resources. The computers in peer-to-peer networks usually run desktop or client OSs, such as Windows 7, Mac OS X, and Linux distributions configured to run as a desktop OS.

In a **server-based network**, certain computers take on specialized roles and function mainly as servers, and ordinary users' machines tend to function mainly as clients. Windows Server 2008, Linux, and UNIX are operating systems designed primarily for server use. In these networks, servers have centralized authority over who has access to network resources.

Peer-to-Peer/Workgroup Model

As you have learned, computers on a peer-to-peer network can take both a client and a server role. Because all computers on this type of network are peers, these networks impose no centralized control or security over shared resources. Any user can share resources on his or her computer with any other user's computer, and each user can determine what level of access other users have to his or her shared resources. Physically, a peer-to-peer network looks just like a server-based network; mainly, location and control over resources differentiate the two.

In a peer-to-peer network, every user must act as the administrator of his or her computer's resources. Users can give everyone else unlimited access to their resources or grant restricted (or no) access to other users on the network. To grant this access, users must create user accounts and passwords for each user who will access shared resources on

their computer. The username and password used to access a computer are called **credentials**. If you have five computers in a peer-to-peer network, each user might have to remember as many as five different sets of credentials. Because of the lack of centralized authority over resources, controlled chaos is the norm for all but the smallest peer-to-peer networks, and security can be a major concern because not all users might be educated in creating secure passwords.

On a Windows-based peer-to-peer network, computers are members of a workgroup, but a workgroup is simply an identifier and doesn't constitute a network security boundary. In other words, users on computers in Workgroup A can access resources on computers in Workgroup B as long as they have the correct credentials.

Although this system can work on small networks, as the number of users and computers grows, these networks can become unworkable—not because they don't operate correctly, but because users can't cope with the complexity of having to remember multiple sets of credentials to access resources spread out over several computers. This limitation is in contrast to a server-based network, in which security of all resources is administered centrally.

Most peer-to-peer networks consist of collections of desktop PCs linked by a common network medium and network connectivity device, such as a switch. The machines and the OS installed on them aren't tuned to provide network services as efficiently as dedicated network servers configured with server operating systems. They can bog down easily under increasing loads, as more users try to access resources from a particular machine. The user whose machine is being accessed across the network also has to endure a performance reduction while his or her machine is busy handling network information requests. For example, if a user's machine has a network-accessible printer attached, the machine slows down every time someone sends a job to that printer. In addition, if a user restarts the machine not knowing that someone is accessing a resource on it, the network user's access fails or, even worse, data loss can occur.

Another issue that affects peer-to-peer networks is data organization. If every machine can be a server, how can users keep track of what information is stored on which machine? If five users are responsible for a collection of documents, any of those users might have to search through files on all five machines to find a document. The decentralized nature of peer-to-peer networks makes locating resources more difficult as the number of peers increases. Likewise, decentralization makes backup considerably trickier: Instead of backing up a single server that holds the shared documents, each machine must be backed up to protect shared data.

Given these issues and complexities, peer-to-peer networks might not seem worth using. However, they offer some advantages, particularly for small organizations. Peer-to-peer networks are the easiest and most inexpensive to install. Most require only a desktop OS (such as Windows 7 or Vista) on desktop computers along with cabling and connectivity devices. After computers are connected and configured correctly, users can begin sharing information immediately. Desktop computers and desktop/client operating systems cost considerably less than their server counterparts.

Peer-to-peer networks are well suited to small organizations, which tend to have small networks and small operating budgets. They're also easy to use and don't require extensive staff training or a dedicated network administrator. With no centralized control, the loss of



a single machine means only the loss of access to the resources on it; otherwise, a peer-to-peer network continues to function when one computer fails. However, because managing resources and their security is difficult on a peer-to-peer network, even small networks of a few computers sometimes opt to use the server or domain network model.



TIP

Windows desktop OSs limit the number of simultaneous network connections to 10 (20 in Windows 7), making the use of peer-to-peer/workgroup networking with only the desktop version of Windows OSs impractical when there are more than 10 computers on the network.

Server/Domain-Based Model

Server-based networks provide centralized control over network resources, mainly by providing an environment in which users log on to the network with a single set of credentials maintained by one or more servers running a server OS. Server OSs are designed to handle many simultaneous user logons and requests for shared resources efficiently. In most cases, servers are dedicated to running network services and shouldn't be used to run user applications. You want to reserve servers' CPU power, memory, and network performance for user access to network services.

When you're using Windows Server OSs in a server-based network with centralized logon, you're running a Windows domain. A **domain** is a collection of users and computers whose accounts are managed by Windows servers called **domain controllers**. Users and computers in a domain are subject to network access and security policies defined by a network administrator and enforced by domain controllers. The software managing centralized access and security is referred to as a **directory service**. On Windows servers, the directory service software is Active Directory, and it's what makes a Windows server a domain controller.



NOTE

The Windows domain model came about with Windows NT in the early 1990s. However, the Active Directory implementation of the domain model was first used in Windows 2000.

The Linux OS supports a centralized logon service called **Network Information Service (NIS)**, but more often Linux administrators use a service compatible with Active Directory called **Lightweight Directory Access Protocol (LDAP)** if they want to use a directory service. A directory service is one of several network services usually found only on server OSs running in a server-based network. Others include the following:

- *Naming services*—Translate computer names to their address.
- *E-mail services*—Manage incoming and outgoing e-mail from client e-mail programs.
- *Application services*—Grant client computers access to complex applications that run on the server.
- *Communication services*—Give remote users access to a corporate network.
- *Web services*—Provide comprehensive Web-based application services.

Unlike peer-to-peer networks, server-based networks are easier to expand. Peer-to-peer networks should be limited to 10 or fewer users, but server-based networks can handle

anywhere from a handful to thousands of users. In addition, multiple servers can be configured to work together, which enables administrators to add more servers to share the load when an application's performance wanes or to provide fault tolerance if a server's hardware malfunctions.

Like peer-to-peer networks, server-based networks have some disadvantages. The most obvious is the additional overhead of operating a server-based network. Server-based networks require one or more dedicated computers to run the server OS. Computers sold as servers usually have features that improve reliability and performance and cost more than desktop computers. In addition, these networks usually require at least part-time support from a person skilled in managing server OSs. Acquiring the skills to manage a server-based network or hiring a trained network administrator adds quite a bit to operating costs.

Housing all your network resources and services on a single server makes administration of resources easier in the long run, but it also creates a single point of failure. Fortunately, most server OSs now have redundancy features that allow taking a single server offline while other machines assume that server's duties. Naturally, having redundant hardware is costly. You must carefully weigh the costs of lost productivity in the event of server failure with the additional hardware and software costs of providing this redundancy.

Table 1-3 summarizes the strengths and weaknesses of peer-to-peer/workgroup and server/domain-based networks.

Table 1-3 Peer-to-peer versus server-based networks

Network attribute	Peer-to-peer network	Server-based network
Resource access	Distributed among many desktop/client computers; Makes access to resources more complex	Centralized on one or more servers; streamlines access to resources
Security	Users control their own shared resources and might have several sets of credentials to access resources; not ideal when tight security is essential	Security is managed centrally, and users have a single set of credentials for all shared resources; best when a secure environment is necessary
Performance	Desktop OS not tuned for resource sharing; access to shared resources can be hindered by users running applications	Server OS tuned for resource sharing; servers are usually dedicated to providing network services
Cost	No dedicated hardware or server OS required, making initial costs lower; lost productivity caused by increasing complexity can raise costs in the long run	Higher upfront costs because of dedicated hardware and server OSs; additional ongoing costs for administrative support

Peer-to-peer networks and server-based networks both have advantages. For this reason, using some combination of the two models isn't uncommon. For example, a user might want to share a printer with a group of users in close proximity or a document folder with a department colleague. With this arrangement, a user is in control of a shared resource yet can still assign permissions to this resource by using accounts from the central user database on the server. Although sharing the resource is decentralized, the logon credentials needed to access the resource are still centralized.



Hands-On Project 1-5: Exploring Peer-to-Peer Networking

Time Required: 15 minutes

Objective: View other computers and shared resources on a peer-to-peer network.

Required Tools/Equipment: Your classroom computer

Description: In this project, you view other computers and shared resources in a peer-to-peer network. You also view and, if necessary, change the type of network (public or private) you're connected to.



All students should use the same username and password.

1. Start your computer and log on as **NetAdmin**, if necessary.
2. Click **Start**, right-click **Computer**, and click **Properties** to open the System control panel. In the “Computer name, domain, and workgroup settings” section, examine the current settings. Does your computer belong to a workgroup or a domain? Is this computer operating in a peer-to-peer or server-based environment? Write your answers on the following lines:

3. Your computer name should be NET-XX (with XX representing your student number) and your workgroup should be NETESS. Verify with your instructor whether they're the right settings for your environment. If the settings are incorrect, click **Change settings**; otherwise, close the System control panel and skip to Step 4. In the System Properties dialog box, click **Change**. Type NET-XX in the Computer name text box (replacing XX with your student number), type NETESS in the Workgroup text box, and click **OK**. Click **OK** in the message box welcoming you to the NETESS workgroup, and then click **OK** in the message box stating that you must restart your computer. Click **Close**, and then click **Restart Now** to restart your computer. When the computer restarts, log on.
4. To see other computers on the network and share files, you need to verify that certain network settings are correct. Open the Network and Sharing Center by clicking **Start**, **Control Panel**, and under Network and Internet, click **View network status and tasks**. In the “View your active networks” section, verify that the network is listed as Work or Home. If the network is listed as Public, click **Public network** to open the Set Network Location dialog box. Click **Work network**, and then click **Close**.
5. Click **Change advanced sharing settings** to open the Advanced sharing settings dialog box (see Figure 1-22).
6. Under Network discovery, click the **Turn on network discovery** option button, if necessary. Under File and printer sharing, click the **Turn on file and printer sharing** option button, if necessary. Click the **Save changes** button, and then close the Network and Sharing Center.



Figure 1-22 The Advanced sharing settings dialog box

Courtesy of Course Technology/Cengage Learning

7. Windows 7 and Windows Vista no longer group computers by workgroup name in the GUI, as in Windows XP and earlier. However, you can see the list of computers in your workgroup by using the command line. To open a command prompt window, click **Start**, type **cmd** in the Search programs and files text box, and press **Enter**.
8. At the command prompt, type **net view** and press **Enter**. You should see a list of computers in your workgroup, similar to Figure 1-23.
9. To view shared resources on a computer, you use the `net view computername` command. For example, to see whether there are any shared folders or printers on the instructor's computer, type **net view net-instr** and press **Enter**. You should see a screen similar to Figure 1-24, in which the share name is listed as NetDocs and the type is listed as Disk.
10. Close the command prompt window, but leave your computer running for the next project.




Figure 1-23 Using the net view command to list computers in a workgroup

Courtesy of Course Technology/Cengage Learning

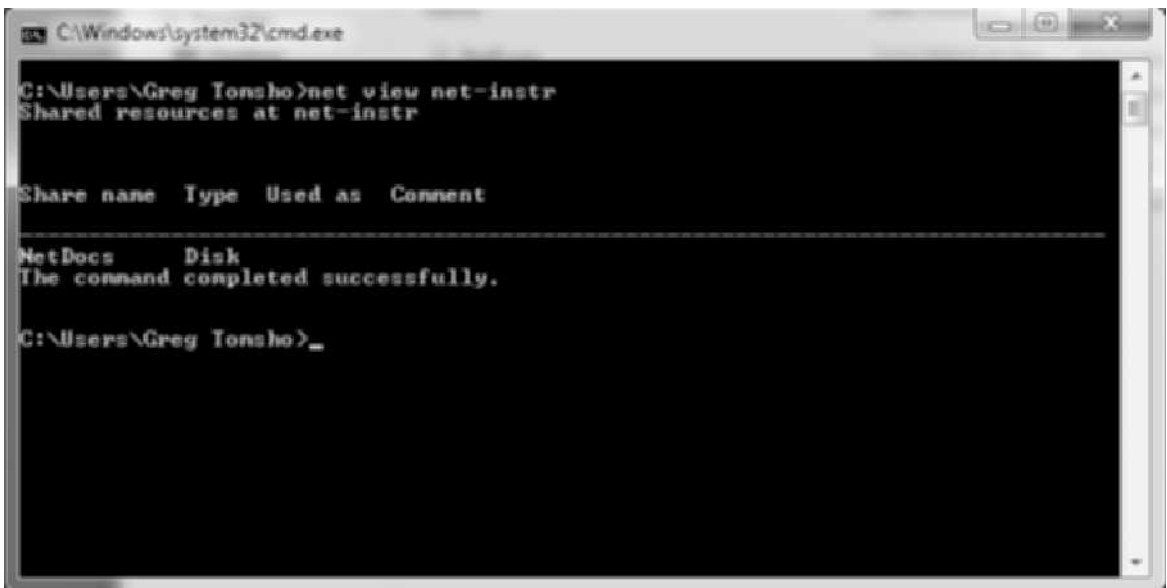


Figure 1-24 Viewing shared resources with the net view command

Courtesy of Course Technology/Cengage Learning



Hands-On Project 1-6: Creating a Shared Folder

Time Required: 15 minutes

Objective: Create a new folder on your computer and share it with the rest of the network.

Required Tools/Equipment: Your classroom computer

Description: In this project, you create a new folder and then share it so that other users can add files to the folder via the network. Your instructor might assign you a partner.

1. Start your computer and log on as **NetAdmin**, if necessary.
2. Click **Start, Computer**. Double-click the **D** drive (or another drive specified by your instructor). Click **New folder**, type **MyData**, and press **Enter** to name the folder.
3. Right-click **MyData** and click **Properties**. In the Properties dialog box, click the **Sharing** tab.
4. Click **Share**. In the File Sharing dialog box, click the down arrow and click **Everyone**. Click **Add**. Notice that the default permission level is **Read**. Click the **Read** down arrow and then click the **Read/Write** permission level. Notice that the account you used to create the share has the permission level **Owner**, which grants the user full access to the share, including the ability to change its permissions.
5. Click **Share** to finish sharing the folder. In the confirmation dialog box shown in Figure 1-25, notice the notation under the share name: **\\NET-XX\MyData**. It's the network path to the share that users on other computers can use to access the shared folder. This notation is



Figure 1-25 The confirmation dialog box displayed after creating a share

Courtesy of Course Technology/Cengage Learning

- referred to as the Universal Naming Convention (UNC) path, which you learn more about in Chapter 8. Click **Done**, and then click **Close**.
6. Try opening another student's shared folder by clicking **Start**, typing `\\NET-XX\MyData` (substituting your partner's student number for `XX`), and pressing **Enter**. An Explorer window should open. To create a new file, right-click the Explorer window, point to **New**, and click **Text Document**. Type your initials and press **Enter** to name the file.
 7. To verify that your partner created a new folder in your MyData share, click **Start**, **Computer**, double-click the **D** drive, and then double-click the **MyData** folder. If your partner finished Step 6, a new file should be there.
 8. Close all open windows. You just performed some basic tasks associated with maintaining a network: creating shared folders and assigning permissions. You assigned Read/Write permissions to the Everyone group, which is a special group in Windows. All user accounts created on your computer belong to the Everyone group automatically, and you can't change this setting. You were able to access your partner's shared folder because you were both logged on to your computers with the same username and password, so you had the correct credentials.
 9. To view the current users on your computer, click **Start**, **Control Panel**. Click **User Accounts and Family Safety**, and then click **Add or remove user accounts** to display the Manage Accounts dialog box (see Figure 1-26), where you can create new accounts and change the properties of existing user accounts. In Chapter 8, you work more with user accounts.

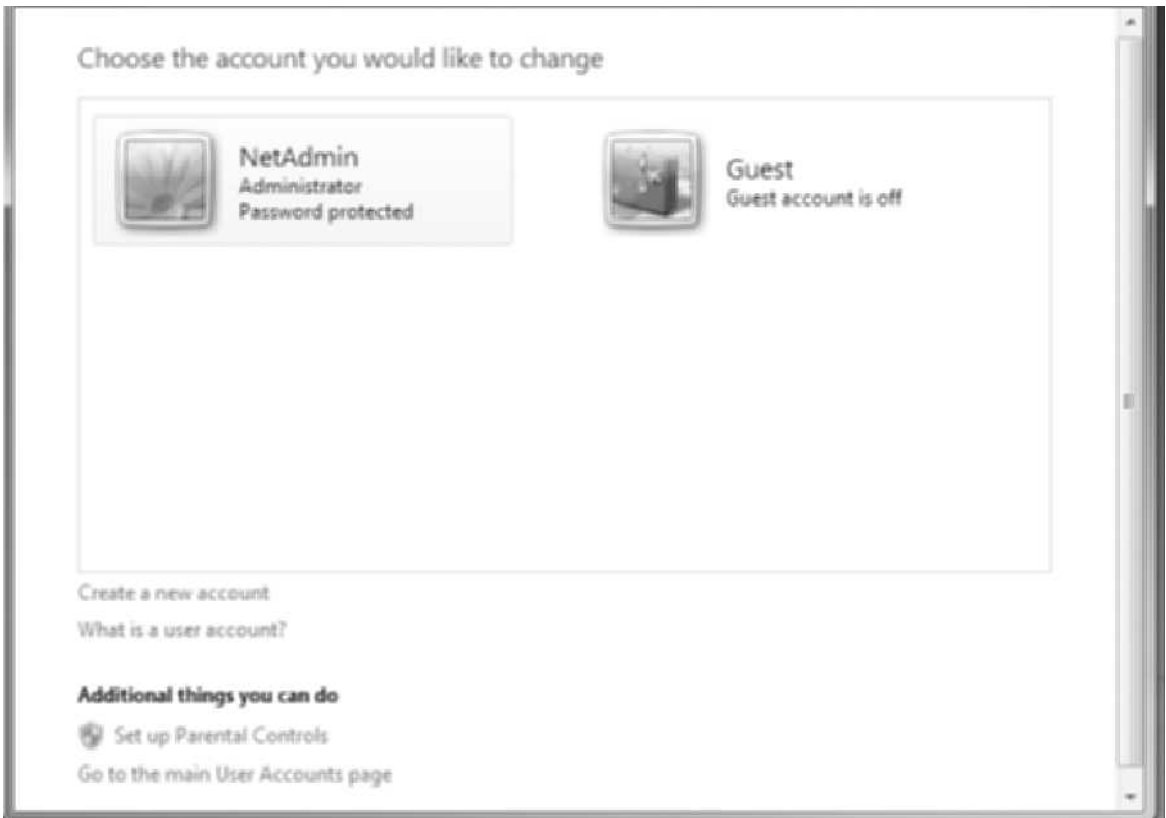


Figure 1-26 Managing user accounts

Courtesy of Course Technology/Cengage Learning

10. Write down answers to the following questions:

- What type of networking service was used in this activity?
-

- Which network model was used in this activity?
-

11. Close all open windows, but leave your computer running for the next project.



Hands-On Project 1-7: Transferring a Document to Another Computer

Time Required: 15 minutes

Objective: Create a document and copy it to your instructor's computer.

Required Tools/Equipment: Your classroom computer

Description: This project requires some setup by your instructor, so verify that the setup has been finished before continuing. In this project, you write a memo to your instructor containing the information specified in the following steps. Then you copy the file you created to a file share on your instructor's computer (or some other computer your instructor designates).

1. Start your computer and log on as **NetAdmin**, if necessary.
2. Start Microsoft Word or another word-processing program; even a simple text editor, such as Notepad, will do. Write a letter to your instructor that includes the following:
 - The reason you're taking this class
 - What you hope to get out of this class
 - How much time you expect to put into this class each week outside classroom hours
 - Whether you expect to take more computer and networking classes
3. Save the document in your **Documents** folder (or a folder your instructor designates), naming it *yourname*. For example, if your name is Bill Smith, name the document **billsmith**.
4. Start Windows Explorer and navigate to the folder where you saved the letter. Right-click the document you created and click **Copy**.
5. To paste the document to the instructor's shared folder, use the UNC path of your instructor's computer, which should be \\Net-Instr\NetDocs, unless your instructor specifies otherwise. Click **Start**, type \\Net-Instr\NetDocs, and press **Enter**.
6. Click **OK**. You should see a Windows Explorer window open. (The folder might already contain documents if some of your classmates have already completed the activity.) Right-click a blank space on the right and click **Paste**. Your document should now be available on your instructor's computer.
7. Close all open windows, but leave your computer running for the next project.



Network Servers

A server is at the heart of any network that's too large for a peer-to-peer configuration. In fact, most large networks with more than a few dozen workstations probably rely on several network servers. A network server can fulfill many roles on your network. Most roles entail the server providing one or more network services. A single server can be configured to satisfy a single role or several roles at once. Following are the most common server roles found on networks, described in more detail in the subsequent sections:

- Domain controller/directory servers
- File and print servers
- Application servers
- Communications server
- E-mail/fax servers
- Web servers

Because you see different servers depicted in network drawings, Figure 1-27 shows some common representations of various types of servers.




Figure 1-27 Different types of network servers

Courtesy of Course Technology/Cengage Learning

Domain Controller/Directory Servers

Directory services make it possible for users to locate, store, and secure information about a network and its resources. Windows servers permit combining computers, users, groups, and resources into domains. Any user belonging to a specific domain can access all resources and information he or she has permission to use simply by logging on to the domain. In Windows, the server handling this logon service and managing the collection of computers, users, and so on in a domain is a domain controller. As mentioned, the software needed to make a Windows server a domain controller is Active Directory, and the Linux directory service add-on that's compatible with Active Directory is LDAP. LDAP is included in most Linux distributions and can be used for centralized logon and resource management on a Linux network or to integrate Linux computers into a Windows network.

File and Print Servers

File and print servers are the mainstay of the server world because they provide secure centralized file storage and sharing and access to networked printers. With these servers, users can run applications locally but keep data files on the server. Any Windows or Linux computer can act as a file and print server. However, using the Server version of Windows provides advanced sharing features, such as fault tolerance, load balancing, and disk quotas. Chapter 8 covers these advanced features in more detail.

Application Servers

Application servers supply the server side of client/server applications, and often the data that goes along with them, to network clients. A database server, for instance, maintains a database of information and provides network clients with a method to send a query to retrieve data.

Application servers differ from basic file and print servers by providing processing services as well as handling requests for file or print services. In file and print services, the client does its own file handling and print processing. Generally, clients must run specialized client-side applications to communicate with an application server. For these applications, typically the client side formulates requests and sends them to the application server, which handles all the request's background processing and then delivers the results back to the client side. The client side then formats and displays these results to the user. Application servers can also be specialized Web servers. For example, when you connect to a shopping site such as Amazon.com, the processing required to find items, process the shopping cart, and handle payment is handled by the application servers at Amazon.com. Your Web browser is simply a client to the application with the main job of displaying information onscreen.

Communication Servers

Communication servers provide a mechanism for users to access a network's resources remotely. They enable users who are traveling or working at home to dial in to the network via a modem or, more commonly, through their existing Internet connection. Windows Server includes a powerful communication server, called Routing and Remote Access Services (RRAS), for handling dial-up network connections and virtual private network (VPN) connections. A VPN provides a secure connection to a private network through the Internet. Similar add-on products are available for Linux/UNIX.

E-Mail/Fax Servers

Mail servers handle sending and receiving e-mail messages for network users. Microsoft Exchange Server is sophisticated mail server software that runs on Windows servers, and Lotus Notes is a mainstay in many organizations. Mail servers generally handle at least two widely used e-mail protocols: Post Office Protocol version 3 (POP3) and Simple Mail Transfer Protocol (SMTP). POP3 is used by client e-mail programs to contact the mail server to download new messages. SMTP is used by client e-mail programs to send e-mail messages and by the mail server to transfer messages from one server to another.

Fax servers manage fax traffic for a network. They receive incoming faxes via telephone, distribute them to recipients over the network, and collect outgoing faxes across the network before sending them via telephone. Some fax servers integrate with the e-mail system so that users can receive and sometimes send faxes with their e-mail client. Windows Server 2008 includes a fax server service; however, earlier Windows Server versions do not.



Web Servers

The World Wide Web is the most well-known aspect of the Internet, made up of millions of documents that can be interlinked by using hyperlinks. Being able to view and retrieve documents with the click of a mouse makes the Internet's resources widely available. Windows Server includes a complete **Web server** called Internet Information Services (IIS) as well as File Transfer Protocol (FTP) services. The excellent Apache Web Server is available as part of most Linux distributions. In fact, Apache remains the most widely used Web server in the world.

Additional Network Services

In addition to the common server roles discussed, most networks require additional support services to function efficiently. The most common are Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP). As discussed, DNS provides name resolution services that allow users to access both local and Internet servers by name rather than address. DHCP handles automatic addressing for network clients, relieving network administrators from having to assign computer addresses manually. Both services are discussed in Chapter 5.

As networks grow larger and more complex, specialization of server roles is increasing. Microsoft has included an installation option called Server Core in Windows Server 2008 to address this trend. It's designed with a limited user interface and is targeted to organizations that need to install a server to run a specialized role, such as Active Directory, DNS, or DHCP.

The myriad functions a network server can perform give rise to wide variations in server hardware requirements. The next section discusses the basic requirements for Windows server and desktop OSs. Keep in mind, however, that the actual requirements depend on the tasks required of the server.

Server Hardware Requirements

A server's primary function is to handle client requests for network resources and other network services. Handling service requests across a network adds to a machine's processing load. The higher this load, the more important it is to purchase computers with additional power to handle demands for network resources. To get an idea of what's involved, review Table 1-4, which compares the system hardware requirements for Windows 7, Windows Vista, and Windows Server 2008.

Table 1-4 Hardware requirements for Windows operating systems

Item	Windows 7	Windows Vista	Windows Server 2008
RAM	1 GB	1 GB	512 MB
Disk type	SATA	SATA	SCSI or SATA
Disk space	40 GB, with 16 GB available	40 GB, with 16 GB available	32 GB
CPU speed	1 GHz	1 GHz	1.4 GHz
Graphics	128 MB graphics memory; DirectX 9 support	128 MB graphics memory; DirectX 9 support	Super VGA



Serial Advanced Technology Attachment (SATA) and Small Computer System Interface (SCSI) are disk controller/disk drive types. SATA is used in desktops and low- to mid-range servers, and SCSI is used primarily in mid-range to high-end servers.

When you look at Table 1-4, you might come to the conclusion that Windows desktop OSs require more power than a server OS. At first glance, this is true. However, the higher RAM and disk space requirements of Windows 7 and Vista are mainly caused by a more graphics-intensive user interface. The Windows Server 2008 requirements are for a Windows server installation in which few network services are running. If you install the directory service, file and printer sharing, DNS, and perhaps an application server and a few dozen users start using these services, the requirements increase substantially. That being said, Microsoft has done a fine job of keeping Windows Server 2008's hardware requirements reasonable. A server with 512 MB of RAM can handily support Active Directory, DNS, and file sharing as long as it's supporting fewer than about 20 users.

Interestingly, the requirements for Windows 7, Microsoft's most recent OS, and Windows Vista are identical. This is a departure from the past, when a new version of a Windows OS meant you had to have a much faster CPU and much more RAM to run it.

Specialized Networks

LANs, internetworks, and WANs are the focus of this book, but you might come across other network types that are intended to connect peripheral devices rather than desktop and server computers. The two specialized networks discussed in the next sections serve very different purposes; storage area networks are used to connect storage devices via very high-speed links, and wireless personal area networks are used to connect one person's personal devices.

Storage Area Networks

A **storage area network (SAN)** uses high-speed networking technologies to provide servers with fast access to large amounts of disk storage. The storage managed by a SAN appears to the server OS as though it's physically attached to the server. However, the storage is connected to a high-speed network technology and can be shared by multiple servers. The most common network technologies used in SANs are Fibre Channel and iSCSI. These technologies are designed to connect large arrays of hard drive storage that can be accessed and shared by servers. Client computers access the shared data by contacting the servers via the usual method, and the servers retrieve the requested data from the SAN devices and pass it along to the client computer. Figure 1-28 shows a LAN with three servers connected to a SAN.

Wireless Personal Area Networks

With all the wireless devices people carry and their need to be connected at all times, it's no wonder that a networking technology designed to connect these devices was developed. A **wireless personal area network (WPAN)** is a short-range networking technology designed to connect personal devices to exchange information. These devices include cell phones, personal digital assistants (PDAs), global positioning system (GPS) devices, MP3 players, and even watches.




Figure 1-28 A typical SAN arrangement

Courtesy of Course Technology/Cengage Learning



A PDA is a handheld computer used for personal organization tasks, such as appointment and address book management. Today's smart-phones have largely replaced PDAs.

A WPAN can connect devices you wear or come in close contact with and transmit to outside devices over a short range, typically within 10 meters or less, by using a secure access method. For example, your WPAN-enabled digital phone can communicate wirelessly with your car's communication/entertainment system. When your phone is in your car, your WPAN-enabled car allows you to make and receive calls by using the controls on the steering wheel or dashboard and the car's audio system.

IEEE 802.15 is the standard for WPANs. This standard covers several areas, including both high and low data transfer rates. One product using this standard is Bluetooth, a short-range wireless technology developed by Ericsson and found in a variety of devices, such as PDAs, cell phones, ear phones/headsets, GPS devices, MP3 players, and cars. It's the technology in cell phones that allows them to communicate with Bluetooth-equipped cars, as discussed, and Bluetooth headsets for hands-free phone calls.



Hands-On Project 1-8: Looking Up Computer and Networking Acronyms

Time Required: 15 minutes

Objective: Use Acronymfinder.com to learn the meaning of common computer and networking acronyms.

Required Tools/Equipment: Your classroom computer and Internet access

Description: This project requires access to the Internet. Half the battle of learning any new field or technology is learning the language used by professionals in the field. Computer and networking technology is well known for its heavy use of acronyms. In this project, you use Acronymfinder.com to look up acronyms.

1. Start your Web browser, and go to www.acronymfinder.com. Figure 1-29 shows the Acronym Finder home page.



Figure 1-29 The Acronym Finder home page

Courtesy of Course Technology/Cengage Learning



Web locations and Web pages change constantly. Don't worry if the Acronym Finder home page looks a little different from the example—you should still be able to navigate it in much the same way.

2. You can look up acronyms by typing them in the “abbreviation to define” text box at the top and clicking the **find** button. If there’s more than one common definition for an acronym, Acronymfinder.com lists them by popularity ranking. Look up the following acronyms; you’ll need some of them later:
 - TCP/IP
 - Wi-Fi
 - SSID
 - WEP
 - OSI
 - Ping
 - UTP
 - Cat6
 - EMI
 - RJ-45
3. Bookmark Acronymfinder.com for future use, and exit your browser. Shut down your computer, unless you’re continuing to the case projects.

Chapter Summary

- All computers perform three basic tasks: input, processing, and output. Some components of today’s computers are designed to perform only one of these three functions; others are designed to perform two or all three functions.
- Storage is a major part of a computer’s configuration. Storage types include short-term storage (RAM) and long-term storage (disk drives and flash drives).
- PC hardware consists of four major components: motherboard, hard drive, RAM, and BIOS/CMOS. The motherboard is the nerve center of the computer and contains the CPU, expansion slots, and RAM slots.
- The operating system (OS) and device drivers control access to hardware and provide a user interface, memory management, and multitasking.
- The components needed to make a stand-alone computer a networked computer include a NIC, network medium, and usually a device to interconnect with other computers. In addition, network software consisting of client and server software, protocols, and the NIC driver are needed to enable the computer to communicate on the network.
- The layers of the network communication process can be summarized as user application, network software, network protocol, and network interface.
- The terms used to describe networks of different scope are LAN, internetwork, WAN, and MAN. A LAN is a single collection of devices operating in a small geographic area. An internetwork is a collection of LANs tied together by routers, and a WAN and MAN are geographically dispersed internetworks.
- Packets and frames are the units of data handled by different network components. Packets, which are processed by the network protocol, are units of data with the

source and destination IP addresses added. Frames, which are processed by the network interface, have MAC addresses and an error code added to the packet.

- A client is the computer or network software that requests network data, and a server is the computer or network software that makes network data available to requesting clients.
- A peer-to-peer network model has no centralized authority over resources; a server-based network usually uses a directory service for centralized resource management.
- Network servers can perform a number of specialized roles, including directory service, file and print server, application server, communication server, e-mail/fax server, and Web server.
- Specialized networks can include storage area networks (SANs) and wireless personal area networks (WPANs).

Key Terms

application servers Computers that supply the server side of client/server applications, and often the data that goes along with them, to network clients.

bus A collection of wires that carry data from one place to another on a computer's motherboard.

client Term used to describe an OS designed mainly to access network resources, a computer's primary role in a network (running user applications and accessing network resources), and software that requests network resources from servers.

communication servers Computers that provide a mechanism for users to access a network's resources remotely.

core An instance of a processor inside a single CPU chip. *See also* multicore CPU.

credentials A username and password or another form of identity used to access a computer.

device driver Software that provides the interface between the OS and computer hardware.

directory service The software that manages centralized access and security in a server-based network.

domain A collection of users and computers in a server-based network whose accounts are managed by Windows servers called domain controllers. *See also* domain controller.

domain controller A computer running Windows Server with Active Directory installed; maintains a database of user and computer accounts as well as network access policies in a Windows domain. *See also* directory service.

encapsulation The process of adding header and trailer information to chunks of data.

file and print servers Computers that provide secure centralized file storage, sharing, and access to networked printers.



frame A packet with source and destination MAC addresses added and an error-checking code added to the back end. Frames are generated by and processed by the network interface. *See also* packet.

header Information added to the front end of a chunk of data so that the data can be correctly interpreted and processed by network protocols.

internetwork A networked collection of LANs tied together by devices such as routers. *See also* local area network (LAN).

local area network (LAN) A small network, limited to a single collection of machines and linked by interconnecting devices in a small geographic area.

mail servers Computers that handle sending and receiving e-mail messages for network users.

metropolitan area network (MAN) An internetwork confined to a geographic region, such as a city or county; uses third-party communication providers to provide connectivity between locations. *See also* internetwork.

multicore CPU A CPU containing two or more processing cores. *See also* core.

multitasking An operating system's capability to run more than one application or process at the same time.

name server A computer that stores names and addresses of computers on a network, allowing other computers to use computer names rather than addresses to communicate with one another.

network Two or more computers connected by a transmission medium that enables them to communicate.

network client software The application or OS service that can request information stored on another computer.

Network Information Service (NIS) A Linux directory service that supports centralized logon.

network model A model defining how and where resources are shared and how access to these resources is regulated.

network protocols The software defining the rules and formats a computer must use when sending information across the network.

network server software The software that allows a computer to share its resources by fielding requests generated by network clients.

packet A chunk of data with source and destination IP addresses (as well as other IP information) added to it. Packets are generated by and processed by network protocols.

peer-to-peer network A network model in which all computers can function as clients or servers as needed, and there's no centralized control over network resources.

server Term used to describe an OS designed mainly to share network resources, a computer with the primary role of giving client computers access to network resources, and the software that responds to requests for network resources from client computers.

server-based network A network model in which servers taken on specialized roles to provide client computers with network services and to provide centralized control over network resources.

stand-alone computer A computer that doesn't have the necessary hardware or software to communicate on a network.

storage area network (SAN) A specialized network that uses high-speed networking technologies to give servers fast access to large amounts of disk storage.

trailer Information added to the back end of a chunk of data so that the data can be correctly interpreted and processed by network protocols.

Web server A computer running software that allows users to access HTML and other document types with a Web browser.

wide area networks (WANs) Internetworks that are geographically dispersed and use third-party communication providers to provide connectivity between locations. *See also* internetwork.

wireless personal area network (WPAN) A short-range networking technology designed to connect personal devices to exchange information.



Review Questions


1. Which of the following is one of the three basic functions a computer performs? (Choose all that apply.)
 - a. Processing
 - b. Internet access
 - c. Input
 - d. Graphics
 - e. Output
 - f. E-mail
2. The _____ executes instructions provided by computer programs.
 - a. CPU
 - b. NIC
 - c. Hard drive
 - d. USB
3. When a CPU is composed of two or more processors, each one is referred to as a(n) _____.
 - a. I/O
 - b. Core
 - c. OS
 - d. Flash

4. Which of the following is considered long-term storage? (Choose all that apply.)
 - a. Flash drive
 - b. RAM
 - c. Working storage
 - d. Hard drive
5. Which motherboard component controls data transfers between memory, expansion slots, I/O devices, and the CPU?
 - a. RAM slots
 - b. IDE connectors
 - c. Chipset
 - d. PCI-Express
6. You want to purchase a new high-performance graphics card for your computer. Which type of connector should it have?
 - a. PCI
 - b. SATA
 - c. IDE
 - d. PCI-Express
7. The time it takes for read/write heads to move to the correct spot on the platter is the _____.
 - a. Rotational delay
 - b. Seek time
 - c. Transfer time
 - d. Access time
8. Which of the following is a task usually performed by the BIOS? (Choose all that apply.)
 - a. Perform a POST.
 - b. Create an interrupt.
 - c. Store the operating system.
 - d. Begin the boot procedure.
9. Place the following steps of the boot procedure in order.
 - a. The OS is loaded into RAM.
 - b. CPU starts.
 - c. OS services are started.
 - d. Power is applied.
 - e. The POST is executed.
 - f. Boot devices are searched.

10. Which of the following is a critical service provided by the OS? (Choose all that apply.)
- Power-on self test
 - Memory management
 - Web browsing
 - File system
 - Storage
11. An OS's capability to run more than one application or process at the same time is referred to which of the following?
- Multicore
 - Doubletime
 - Multitasking
 - Multiprocessor
 - Interrupt processing
12. You have just installed a new NIC in your PC to replace the old one that had started malfunctioning. What additional software must be installed to allow the OS to communicate with the new NIC?
- Network application
 - Device driver
 - BIOS
 - Protocol
13. Which of the following requests information stored on another computer?
- NIC
 - Network client
 - Network server
 - Network protocol
 - Device driver
14. Choose the correct order for the process of a user attempting to access network resources:
- Network protocol
 - Application
 - Network client
 - NIC driver
- 4, 2, 1, 3
 - 3, 2, 1, 4
 - 1, 4, 2, 3
 - 2, 3, 1, 4
 - 3, 1, 2, 4



15. TCP/IP is an example of which of the following?
- NIC
 - Network client
 - Network server
 - Network protocol
 - Device driver
16. In network communication, the _____ address is used to deliver a frame to the correct computer on the network. (Choose all that apply.)
- MAC
 - Logical
 - IP
 - Physical
17. A(n) _____ message is used to determine whether a computer is listening on the network.
- MAC
 - Ping
 - IP
 - TCP
18. TCP/IP uses _____ to look up a computer's IP address, given its name.
- DNS
 - Ping
 - MAC
 - TCP
19. The unit of information containing MAC addresses and an error-checking code that's processed by the network interface layer is referred to as a _____.
- Packet
 - Ping
 - Frame
 - Chunk
20. Data is processed from the time an application creates it to the time it reaches the network medium. This process includes adding information such as addresses and is called which of the following?
- Packetization
 - Encapsulation
 - Deencapsulation
 - Layering

- 
21. You're the network administrator for a company that has just expanded from one floor to two floors of a large building, and the number of workstations you need has doubled from 50 to 100. You're concerned that network performance will suffer if you add computers to your existing LAN. In addition, new users will be working in a separate business unit, and there are reasons to logically separate the two groups of computers. What type of network should you configure?
- WAN
 - MAN
 - Internetwork
 - Extended LAN
22. Which of the following best describes a client?
- A computer's primary role in the network is to give other computers access to network resources and services.
 - A computer's primary role in the network is to run user applications and access network resources.
 - It's the software that responds to requests for network resources.
 - The OS installed on a computer is designed mainly to share network resources.
23. You work for a small company with four users who need to share information on their computers. The budget is tight, so the network must be as inexpensive as possible. What type of network should you install?
- Server-based network
 - Peer-to-peer network
 - WPAN
 - Storage area network
24. Which of the following characteristics is associated with a peer-to-peer network? (Choose all that apply.)
- Decentralized data storage
 - Inexpensive
 - User-managed resources
 - Centralized control
 - Uses a directory service
25. A device interconnects five computers and a printer in a single office so that users can share the printer. This configuration is an example of which of the following?
- LAN
 - MAN
 - WAN
 - Internetwork

26. At Yavapai College, the Prescott and Prescott Valley campuses (8 miles apart) have LANs connected via the local phone company. This configuration is an example of which of the following? (Choose the best answer.)
- MAN
 - WPAN
 - WAN
 - SAN
27. You have installed Windows Server 2008 on a new server and want to centralize user logons and security policies. What type of software should you install and configure on this server?
- Naming services
 - Application services
 - Communication services
 - Directory services
28. Peer-to-peer networks aren't suitable in which of the following situations?
- Tight security is required.
 - Five or fewer users need network access.
 - Budget is the primary consideration.
 - No one uses the network heavily.
29. Which of the following best describes a storage area network?
- Provides a mechanism for users to access a network's storage resources remotely
 - Uses high-speed networking technologies to give servers fast access to large amounts of disk storage
 - Is a short-range networking technology designed to connect personal devices to exchange information
 - Provides secure centralized file storage and sharing and access to networked printers
30. Why might Windows 7 or Windows Vista require more RAM or disk space than Windows Server 2008?
- They need to accommodate handling centralized logon.
 - They include a directory service and a naming service.
 - They run many background networking services.
 - They support a graphics-intensive user interface.

Case Projects



Case Project 1-1

Networking Gadgets, Inc. currently employs eight people but plans to hire 10 more in the next four months. Users will work on multiple projects, and only those users assigned to a project should have access to the project files. You're instructed to set up the network to make it easy to manage and back up yet still provide centralized storage for project files. Would you choose a peer-to-peer network, a server-based network, or a combination of both? Why?

Case Project 1-2

CNT Books hired you as a productivity consultant. Currently, it employs six people who still use floppy disks to get files from one computer to the next. You're to bring them into the 21st century by configuring a network that allows them to share files through the network. Employees must also be able to control resources on their own machines. The company wants the most inexpensive solution and only minimal training for employees. Would you choose a peer-to-peer network or a server-based network? Write a list of supplies you might need to purchase to accomplish this task. What computer configuration tasks might you need to perform?

Case Project 1-3

CNT Books has expanded considerably since you first got the network up and running three years ago. It now occupies an entire floor in the building, and its LAN has grown to include several servers and more than 60 workstations. CNT Books has recently purchased another book company and needs more space and computers. Expansion plans include leasing another floor four stories up in the same building and adding 35 workstations and at least one more server immediately, with additional equipment purchases expected. What type of network is called for—LAN, WAN, MAN, or internetwork? What additional devices might be needed to ensure efficient network communication?

Case Project 1-4

You want your GPS to upload data to your PDA wirelessly so that you can plot your travels on the advanced mapping software installed on your PDA. What networking technology might you want as a feature of both your GPS and PDA to accomplish this task?

Case Project 1-5

Chapter 2 discusses network hardware. To prepare for this topic, go to *www.about.com* and look up the following terms. Read at least one article about each term and be prepared to discuss these terms in class:

- Network interface card
- Hub
- Switch



This page intentionally left blank

Network Hardware Essentials

After reading this chapter and completing the exercises, you will be able to:

- Describe the basic operation of network repeaters and hubs
- Explain the purpose of network switches
- Summarize the operation of wireless access points
- Describe the basic operation of network interface cards
- Explain the function of routers

LANs, WANs, and internetworks are built with a variety of network hardware. Your understanding of how the most common network hardware works is crucial to your success in building reliable, high-performance networks.

This chapter begins by discussing the simplest of network devices, the hub, a device that's nearly obsolete but is still found in older installations. Switches have largely supplanted hubs in networks large and small and are the main network building block today. Wireless networking has exploded in popularity in the past several years and can be found everywhere from small home networks to coffee shops and bookstores to large corporate networks. Wireless access points are the foundation of wireless networks, and you learn about their operation and basic configuration later in the chapter. Network interface cards have become such an essential component of computers that they're now built into most motherboards. Whether they're built in or installed as an expansion card, however, your understanding of NIC configuration options and properties will help you build a better network. The last section of this chapter covers the most complex network devices: routers, the gateway to the Internet that make it possible for large companies to build vast internetworks and WANs.



Because Ethernet is the dominant network technology used in LANs today, the network hardware components discussed in this chapter are Ethernet devices, unless otherwise stated.

Network hardware devices can be complex. This chapter serves as an introduction to the most common devices so that you have a basic understanding of their function when they're discussed with other topics in later chapters. The function of these devices is intertwined with network topologies and technologies (discussed in Chapter 3) and network protocols (discussed in Chapter 5). Chapter 7 includes a more thorough examination of some features and functions of network devices.

Network Repeaters and Hubs

Early networks didn't use interconnecting devices. Computers were connected in daisy-chain fashion by lengths of cable (see Figure 2-1). The problem with this arrangement was that you were limited in the total length of the cabling and the number of computers that could be connected together.

Some problems associated with the type of network shown in Figure 2-1 were solved with a device called a repeater. A **repeater** has the rather straightforward job of receiving bit signals generated by NICs and other devices, strengthening them, and then sending them along or repeating them to other parts of the network. Think of a repeater as a microphone for network signals. When people speak, their voices carry only so far until people in the back of the room can no longer hear what's being said. Network signals, too, carry only so far on their medium before receiving computers can no longer interpret them correctly. A repeater enables you to connect computers whose distance from one another would otherwise make communication impossible.




Figure 2-1 Older networks didn't use interconnecting devices

Courtesy of Course Technology/Cengage Learning



Repeaters don't strengthen signals in the sense that the original signal is amplified; instead, a repeater takes a weakened signal and repeats it at its original strength.

A traditional repeater has two ports or connections that you can use to extend the distance your network can cover, as shown in Figure 2-2. Assuming the two groups of computers in this figure are separated by several hundred feet, the repeater is needed to allow them to communicate with one another.




Figure 2-2 A repeater extends the distance a network can cover

Courtesy of Course Technology/Cengage Learning

Multiport Repeaters and Hubs

A multiport repeater is just a repeater with several ports to which you can connect cabling. Most multiport repeaters have at least four ports, and some have 24 or more. A multiport repeater is commonly referred to as a **hub**, and although it performs the same function as a traditional repeater, it's used as a central connecting device for computers instead of merely a way to extend the network. So instead of daisy-chaining computers together, all computers are connected to the

central hub (see Figure 2-3), as you saw in Chapter 1. Because “hub” is the more common term and is much easier to write and say, a multiport repeater is often referred to as a hub in this book.




Figure 2-3 A multiport repeater or hub

Courtesy of Course Technology/Cengage Learning

A hub performs the same function as a repeater but with more outgoing ports to which bit signals are repeated, so its function is as follows:

- Receives bit signals generated from a connected computer on one of its ports (with all ports capable of receiving signals)
- Cleans the signal by filtering out electrical noise
- Regenerates the signal to full strength
- Transmits the regenerated signal to all other ports a computer (or other network device) is connected to



Like repeaters, hubs require power to operate, so they're sometimes referred to as “active hubs.” However, this term isn't common because unpowered devices known as “passive hubs” aren't used for the same purposes.

Simulation 3 on the book's CD shows basic hub operation.



Simulation 3: Basic operation of a hub

Hubs and Network Bandwidth Network bandwidth is the amount of data that can be transferred on a network during a specific interval. It's usually measured in bits per second,

and networks operate at speeds from 10 million bits per second (10 Mbps) up to 10 gigabits per second (Gbps). This bandwidth is determined by how fast network devices can send bits of data to the medium. A 10 Mbps hub, for example, transmits bits at the rate of 10 million per second. To put this rate into perspective, two computers connected to a 10 Mbps hub can copy one minute of MP3 music to each other in about 1.25 seconds, but a 100 Mbps hub can transfer the same amount of information in about one eighth of a second.

One drawback of using hubs as the central connecting device on a network is that only one computer can successfully transmit data at a time. On a busy network with dozens of computers transferring large files and accessing network applications and databases, this limitation is serious. This setup is referred to as **bandwidth sharing** because all computers connected to the hub must share the amount of bandwidth the hub provides. For example, say a network has 10 computers connected to a 10 Mbps hub, and all 10 computers are trying to send and receive files frequently. Because the computers must share the bandwidth, the average effective bandwidth for each computer is only 1 Mbps. Transferring that one minute of MP3 music in this example would take more than 12 seconds.

In the early days of networking, bandwidth sharing wasn't a big problem because the number and frequency of data transfers in a typical LAN were low and files tended to be small, making the actual effective bandwidth in the preceding example much higher than 1 Mbps. However, in today's LANs, large multimedia data files are transferred often, so the need for additional dedicated bandwidth is paramount. In fact, this need has become so critical that network administrators stopped including hubs in their network designs, and finding a hub to buy from major computer parts retailers is difficult now.



There are more details involved in the concept of bandwidth sharing and how computers transmit data to the medium. The details vary for different network technologies, such as Ethernet, token ring, and Wi-Fi, and are hammered out in Chapter 3.

Hub Indicator Lights Most hubs have indicator lights for power, link status, network activity, and collisions. Each port has a link status indicator (link light) that glows (usually green) when a cable has been plugged in and a valid network connection, or link, has been made to a device on the other end of the cable. Some hubs have a separate indicator, or the indicator might vary in color for different connection speeds. For example, the link light might glow green for a 100 Mbps connection and amber for a 10 Mbps connection.

Another indicator light you're likely to find on a hub is for network activity. When the hub receives bit signals on any of its ports, this indicator flashes. Some hubs combine the link status indicator with the network activity indicator so that when the light is on solidly, a valid link is detected, and when the light is blinking, a valid link and network activity are detected.

A third type of indicator is for collisions. A collision occurs on a hub when two stations try to transmit at the same time, which isn't allowed on a hub-based network. When a collision occurs, the stations that were transmitting must retransmit their data. Collisions are discussed in more detail in Chapter 3.

Figure 2-4 shows a typical hub with indicator lights. This hub also has a series of indicator lights showing the utilization percentage for the network. In addition, the rightmost port has





Figure 2-4 A typical hub with indicator lights

Courtesy of Course Technology/Cengage Learning

a button next to it for changing the port's configuration, depending on whether it's connected to a computer's NIC or another hub. This port is referred to as the **uplink port**. The term "uplink" is used when multiple hubs are connected. When this button is pressed in, you can connect the hub to another hub with a standard cable rather than a crossover cable. Cable types are discussed more in Chapter 4.

Network hubs were the mainstay for connecting computers in a LAN for several years, and although you can still find hubs in the workplace, they're becoming obsolete. Because of their disadvantages, mainly bandwidth sharing, they're being replaced with switches, discussed in the next section.

Network Switches

A network **switch**, like a hub, is used to interconnect multiple computers so that they can communicate with one another. A switch looks just like a hub, with several ports for plugging in network cables. However, instead of simply regenerating incoming bit signals and repeating them to all other ports, a switch actually reads data in the message, determines which port the destination device is connected to, and forwards the message to only that port. So the first important difference between hubs and switches is that hubs work only with electrical signals and the bits these signals represent, whereas switches work with the actual information these bits combine to make frames.

Basic Switch Operation

Data is sent to the medium one frame at a time, and the beginning of each frame contains the destination computer's MAC address and the source computer's MAC address. When the frame reaches a switch, the switch reads both addresses. By reading the source MAC address, the switch keeps a record of which port the sending computer is on. This function is referred to as "learning" because the switch is learning to which port each MAC address in the network corresponds. By reading the destination MAC address, the switch can forward the frame to the port the destination computer is on. A switch maintains a **switching table** (see Figure 2-5) of MAC addresses that have been learned and their associated port numbers.



MAC addresses are 12 hexadecimal digits. Figure 2-5 uses shorter addresses only as an example.

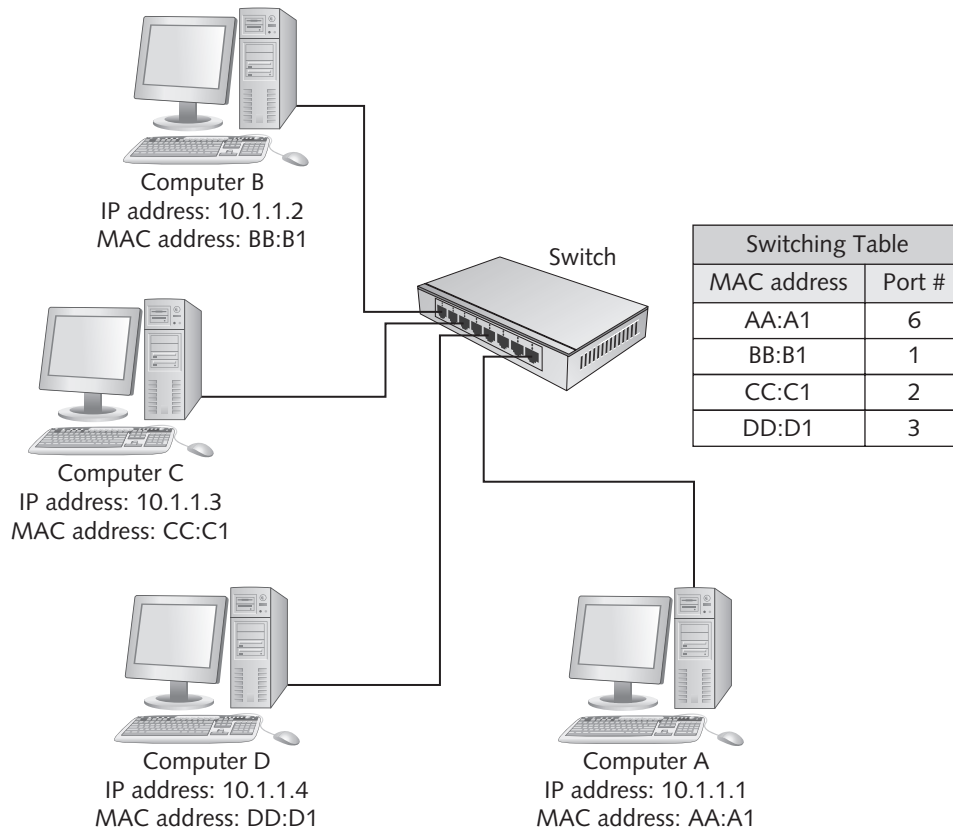


Figure 2-5 Switches maintain a switching table

Courtesy of Course Technology/Cengage Learning

A switch's operation can be summarized in these steps:

1. The switch receives a frame.
2. The switch reads the source and destination MAC addresses.
3. The switch looks up the destination MAC address in its switching table.
4. The switch forwards the frame to the port where the computer owning the MAC address is found.
5. The switching table is updated with the source MAC address and port information.

Simulation 4 on the book's CD shows basic switch operation.



Simulation 4: Basic operation of a switch

These steps raise some questions. For example, what happens if the switch doesn't find the destination MAC address in its switching table? In this case, the switch does the most reasonable thing: It forwards the frame to all ports. You can think of a switch as acting like a switchboard operator. When a call comes in for a person the operator knows, the operator can forward the call to the correct phone extension. If the call is for a person the operator doesn't know, the person can be paged via an intercom system.

You might also be wondering what happens if the source address is already in the switching table and how long each MAC address stays in the switching table. The answers to these questions are related. MAC addresses can't stay in the switching table forever because computers might be shut down or moved to other locations, and their MAC addresses can change. Leaving MAC addresses in the switching table for a long time is akin to having an out-of-date employee phone directory that still lists people who have left the company and others who have changed locations. To ensure that the switching table doesn't become out of date, a timestamp is included in each entry, and each entry can stay in the table for only a certain amount of time unless the timestamp is updated. So when a switch first sees a source MAC address, it creates the switching table entry that includes the MAC address, the port from which the frame arrived, and a timestamp. If the same MAC address is seen again coming from the same port, the timestamp is updated. If the entry remains in the table beyond the maximum allowed time (which varies between switches but is often about 5 minutes) without being updated, it's deleted.

Switches and Network Bandwidth Because a switch is capable of forwarding frames to only a single port instead of all ports, as a hub does, it can handle several computer conversations at one time, thereby allowing each device the full network bandwidth, or **dedicated bandwidth**, instead of requiring bandwidth sharing. In other words, if the switch in Figure 2-5 is a 10 Mbps switch, Computer A could communicate with Computer C at an uninterrupted 10 Mbps, and Computer B could communicate with Computer D at 10 Mbps simultaneously. Furthermore, each computer can receive data at 10 Mbps at the same time it's sending data at 10 Mbps, making each conversation between computers effectively 20 Mbps (10 Mbps in both directions). When a device can send data and receive data simultaneously, it's called **full-duplex mode**. When a device can send or receive (but not both) at one time, it's called **half-duplex mode**. Hubs operate only in half-duplex mode, but switches can operate in both half-duplex and full-duplex modes. Chapter 3 describes these modes of communication in more detail.

The performance advantage of switches has made them the device of choice in networks of all sizes. In addition, although they used to cost more than hubs because of their higher complexity, this is no longer the case. As mentioned, you can still find hubs in the workplace, but new installations rarely specify them, and the tables have been turned—hubs are now usually more expensive than switches because manufacturers simply aren't making them in large quantities.

Switch Indicator Lights Like hubs, switches have indicator lights so that you can see the basic operating status ports with a quick glance. Aside from the requisite power indicator, switches have link status indicators and activity indicators. They might also have indicators to show whether a port is operating in full-duplex or half-duplex mode. Switches, like hubs, can be connected to one another so that your LAN can grow beyond the limitations of the number of ports on a single switch. Some switches also have a dedicated port for uplinking to another switch.

Switches are complex devices, and this section introduces their basic operation. You can find a more detailed examination of switches in Chapter 7.



Hands-On Project 2-1: Downloading and Installing a Protocol Analyzer



If students can't install applications on classroom computers, Wireshark should be installed ahead of time, and this project can be started at Step 6.

Time Required: 15 minutes

Objective: Install the Wireshark protocol analyzer for use in upcoming projects.

Required Tools/Equipment: Your classroom computer with a connection to the Internet or the Wireshark setup program the instructor has supplied on a network share.

Description: In this project, you download and install the Wireshark protocol analyzer. A protocol analyzer, sometimes called a packet sniffer or packet capture program, captures network data the NIC receives and displays it in a user-friendly format. For this procedure, the NIC is placed in promiscuous mode. (A NIC normally receives and processes only broadcast frames and frames addressed to it, but a NIC operating in promiscuous mode processes all frames it receives.) Wireshark is used in projects throughout this book so that you can see how certain devices and protocols work.



Not all NICs can operate in promiscuous mode. This is particularly true of wireless NICs. If promiscuous mode isn't supported, your NIC will capture only broadcast frames and frames addressed to it.

1. Log on to your computer as an administrator. Start a Web browser, go to www.wireshark.org, and click **Download Wireshark**. In the Download Wireshark box, click **Windows Installer (32-bit)** or **Windows Installer (64-bit)**, depending on the Windows version you have installed.
2. When prompted to run or save the file, click **Run**. If the User Account Control (UAC) message box opens, click **Continue**.
3. When the setup wizard begins, click **Next**. In the License Agreement window, click **I Agree**. In the Choose Components window, accept the defaults and click **Next**. Click

- Next in the Select Additional Tasks window, and click **Next** to accept the default install location. Click **Install** in the Install WinPcap window.
4. When you see the WinPcap window, click **Next**, and then click **Next** again in the next window. Click **I Agree** in the WinPcap License Agreement window. Click **Install** in the Installation options window. In the next three windows, click **Finish**, click **Next**, and then click **Finish**.
 5. Exit your Web browser.
 6. To run Wireshark, click **Start**, point to **All Programs**, and click **Wireshark**. You see a window similar to Figure 2-6.

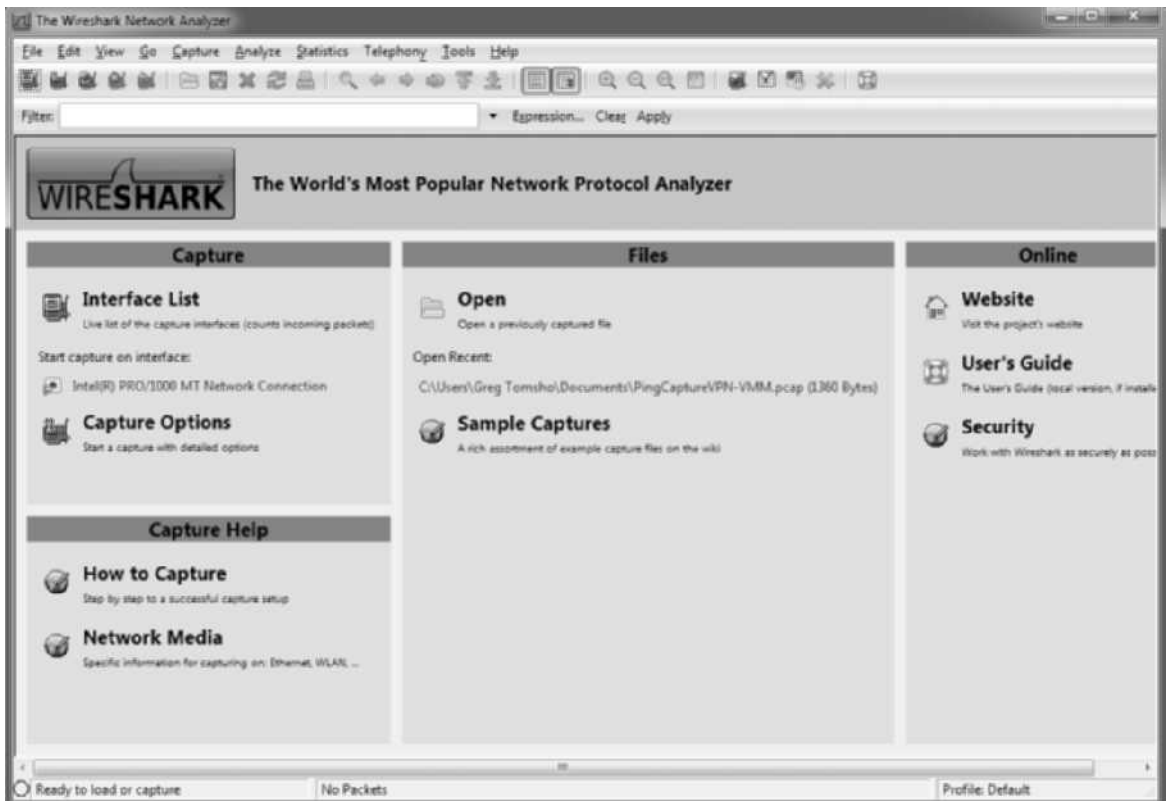


Figure 2-6 The Wireshark main window

Courtesy of Course Technology/Cengage Learning

7. To begin a capture, click the interface name in the Interface List section. In Figure 2-6, the interface name is Intel(R) PRO/1000 MT Network Connection, but your interface name will be different.
8. The number of packets you see in the Wireshark capture window depends on how busy your network is and the type of device used to connect your computers (a hub or

switch). If you see few or no packets, start a Web browser and go to some Web sites to create network activity, and then exit the browser. After you see some packets, click the **Stop the running live capture** toolbar icon shown in Figure 2-7.

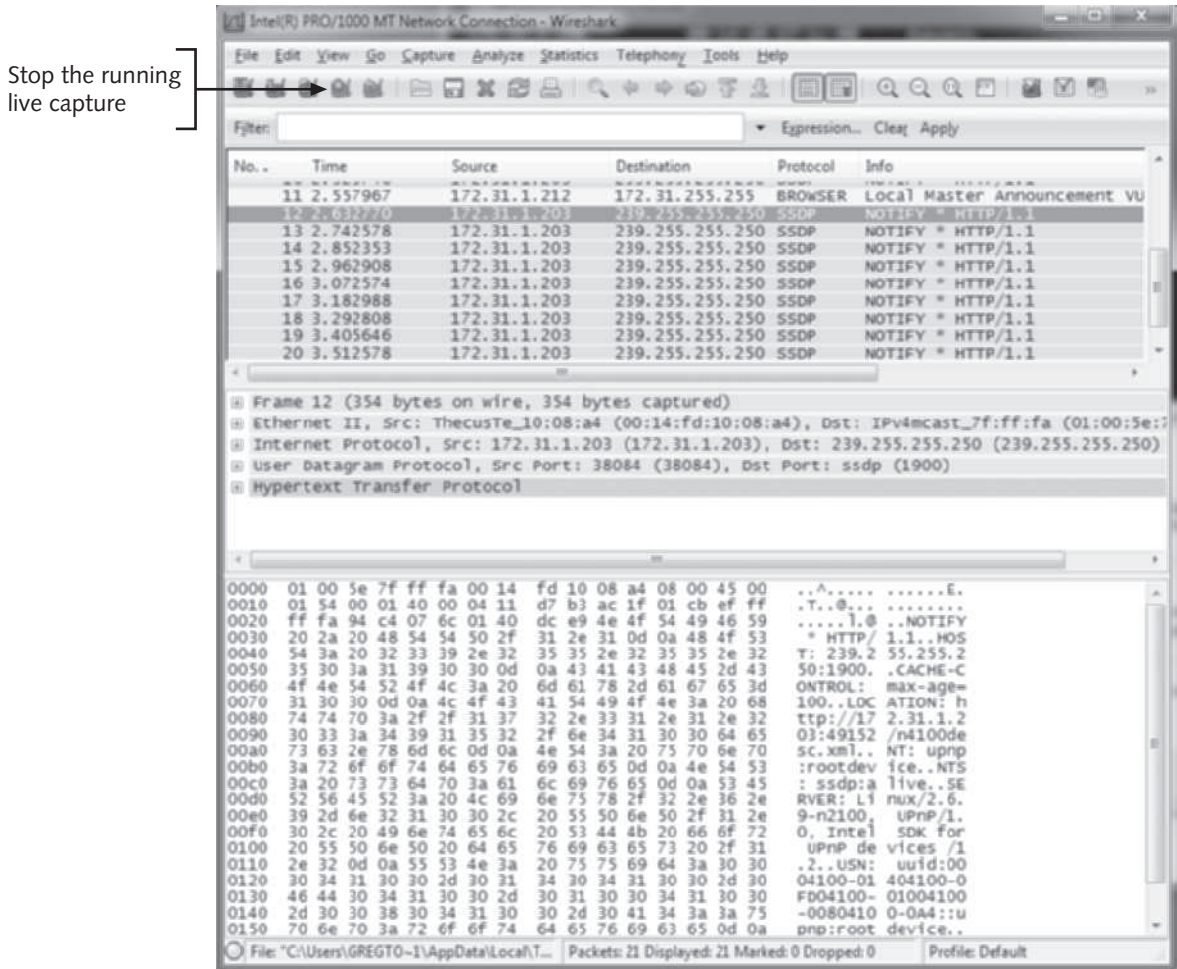


Figure 2-7 Wireshark packet capture window

Courtesy of Course Technology/Cengage Learning

- The Wireshark packet capture window is divided into three panes. The top pane lists the captured packets in summary form, showing the arrival time, source and destination address, protocol, and a description of the packet contents. Scroll through the top pane to see the different packet types that were captured, and then click a packet. The middle pane shows the details of what's in the packet in a more readable format. Click the plus sign next to any line to see more information on that part of the packet. The bottom pane shows the packet's contents in hexadecimal form on the left, and on the right is any data that can be shown as printable alphanumeric characters.

- Exit Wireshark. When prompted to save the capture file, click **Quit without Saving**. Shut down your computer for the next project.



Hands-On Project 2-2: Using Wireshark with a Hub

Time Required: 20 minutes

Objective: Use Wireshark on a computer connected to other computers via a hub to see that all data is repeated to all stations.

Required Tools/Equipment: Three computers (minimum) with Ethernet NICs installed; 10/100 or 10/100/1000 NICs are preferable, but 10 Mbps NICs will also work. Computers can be configured with static IP addresses or can use Automatic Private IP Addressing (APIPA). Three patch cables and a 10/100 hub, although a single-speed hub will also work.

Description: In this project, you run Wireshark on a group of computers connected via a hub. This project shows that a hub repeats all data to all stations so that Wireshark can capture packets generated by all stations. In the next project, you compare this behavior with a switch.



This project requires at least three computers connected to a hub, with at least one computer running Wireshark. It's probably best done in groups. If necessary, the instructor can perform it as a demonstration. The steps in this project assume three computers are connected to the hub and labeled Computer1, Computer2, and Computer3. Wireshark is assumed to be installed on Computer1, but it can be installed on all computers. It's preferable that the computers aren't attached to the classroom network and don't have access to the Internet.

- Connect three computers running Windows XP or later to a central hub with patch cables. Make sure the device is a hub, not a switch.
- Turn on the computers and log on with an administrator account. Open a command prompt window on each computer. On each computer, type **ipconfig** and press **Enter** to display its IP address configuration. Write down these IPv4 addresses so that you know each computer's address:
 - Computer1: _____
 - Computer2: _____
 - Computer3: _____
- If you're using Windows 7, disable Windows Firewall on each machine. By default, Windows 7 blocks ping packets if the network is categorized as "Public," which yours is because you don't have a default gateway configured. To disable the firewall, click **Start**, **Control Panel**, and under Network and Internet, click **View network status and tasks**. In the Network and Sharing Center, click **Windows Firewall** at the bottom left, and then click **Turn Windows Firewall on or off**. Under Public network location settings, click **Turn off Windows Firewall**, and then click **OK**. Close the Windows Firewall window.

4. To verify connectivity, from each computer, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with another computer's IP address). Repeat this step until you have successfully pinged each computer from the other computers. Leave the command prompt window open.
5. On Computer1, start Wireshark, and click **Capture Options** on the left. In the Capture Filter text box, type **icmp**. You must use lowercase letters. The capture filter tells Wireshark to capture only certain types of packets. Internet Control Message Protocol (ICMP) packets are created by the Ping program, so you're going to capture only ping packets. Click **Start**.
6. On Computer2 at the command prompt, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with the IP address of Computer1).
7. On Computer1, click the **Stop the running live capture** toolbar icon to stop the capture. You should see a window similar to Figure 2-8.

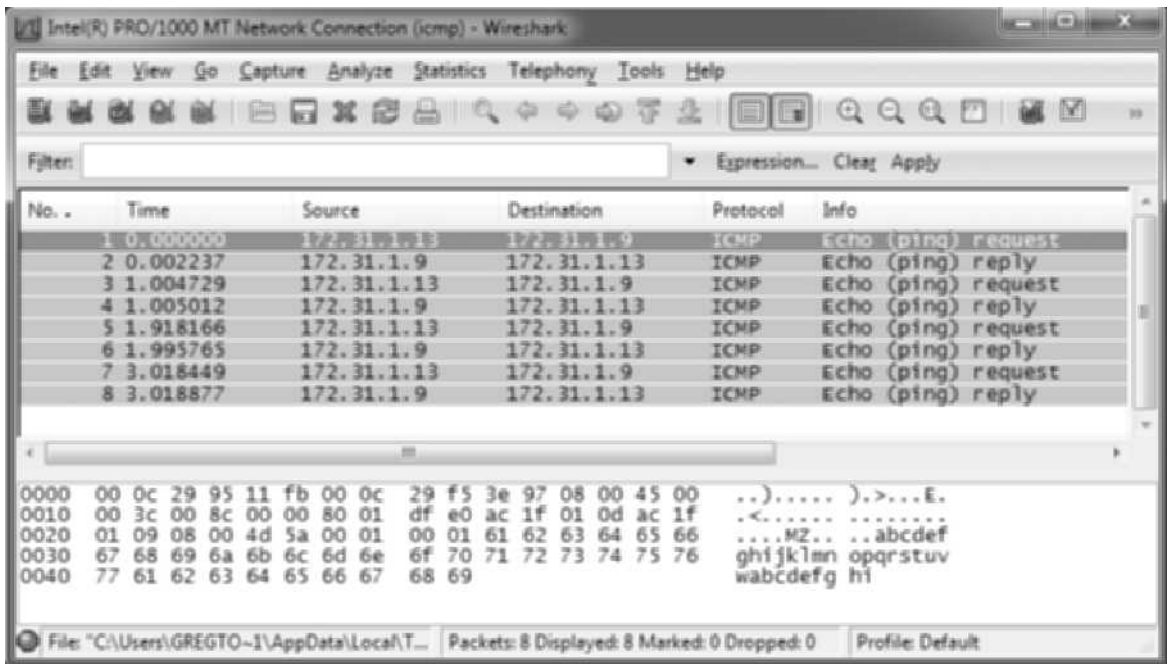


Figure 2-8 Ping packets captured on Computer1

Courtesy of Course Technology/Cengage Learning

8. On Computer1, in Wireshark, click **Capture, Start** from the menu. Click **Continue without Saving**.
9. On Computer2 at the command prompt, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with the IP address of Computer3). Wireshark running on Computer1 should have captured the ping packets.
10. Stop the capture and exit Wireshark. If you're continuing to the next project, stay logged on and leave the command prompt window open on all computers. Otherwise, close all open windows and shut down the computers.



Hands-On Project 2-3: Using Wireshark with a Switch

Time Required: 20 minutes

Objective: Use Wireshark on a computer connected to other computers via a switch to see that all data isn't repeated to all stations.

Required Tools/Equipment: Three computers (minimum) with Ethernet NICs installed; 10/100 or 10/100/1000 NICs are preferable, but 10 Mbps NICs will also work. Computers can be configured with static IP addresses or can use APIPA addresses. Three patch cables and a 10/100 switch, although a single-speed switch or Gigabit switch will also work.

Description: In this project, you run Wireshark on a group of computers connected via a switch. This project shows that a switch only forwards data to the station to which the frame is addressed. In Hands-On Project 2-2, you configured Wireshark to capture only ICMP packets. In this project, you configure Wireshark to also capture Address Resolution Protocol (ARP) broadcast packets to show that switches forward broadcasts to all stations.

1. Connect the computers you used in Hand-On Project 2-2 to a switch instead of a hub, using the same patch cables you used previously.
2. Turn on the computers and log on with an administrator account, if necessary. Open a command prompt window on each computer, if necessary. If the computers were shut down or restarted, their IP addresses might have changed. If so, write down each computer's IP address again.
3. To make sure you have connectivity with the switch, at each computer, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with the IP address of another computer). Repeat this step until you have successfully pinged each computer from all other computers. Leave the command prompt window open.
4. At each computer, type **arp -d** and press **Enter**. As you learned in Chapter 1, ARP manages the MAC addresses your computer has learned. The `arp -d` command deletes the entries created from the pings you did in Step 3. You're deleting these entries so that the computers have to learn the MAC addresses of other computers again. Leave the command prompt window open.



If you're running Windows Vista or Windows 7, when you use `arp -d` you might get an error stating "The ARP entry deletion failed: The requested operation requires elevation." If so, close the command prompt window and open a new one as administrator. To do this, click **Start, All Programs, Accessories**, and then right-click **Command Prompt** and click **Run as administrator**.

5. On Computer1, start Wireshark, and click **Capture Options**. In the Capture Filter text box, type **icmp or arp**. (You must use lowercase letters.) This capture filter tells Wireshark to capture only ICMP or ARP packets. Click **Start**.
6. On Computer2 at the command prompt, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with the IP address of Computer1).
7. On Computer1, click the **Stop the running live capture** toolbar icon. You should see a window similar to Figure 2-9. Notice that the first ARP packet you see has the destination

address “Broadcast.” Click this packet, and the middle pane displays the MAC address ff:ff:ff:ff:ff:ff.

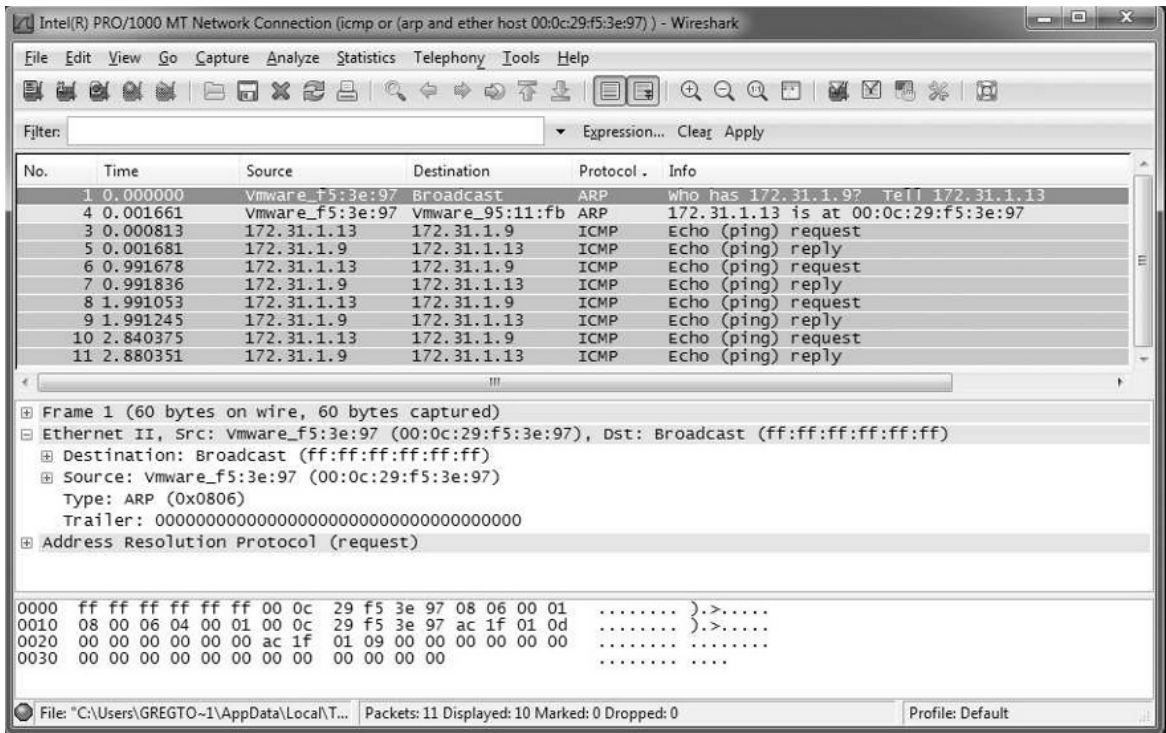


Figure 2-9 Ping and ARP packets

Courtesy of Course Technology/Cengage Learning

8. On Computer1, click **Capture, Start** from the Wireshark menu, and then click **Continue without Saving**.
9. On Computer2 at the command prompt, type **ping IPaddress** and press **Enter** (replacing *IPaddress* with the IP address of Computer3). Wireshark running on Computer1 should have captured only the ARP broadcast packet, not the actual ping ICMP packets, which are unicast packets (explained later in “NIC Basics”). Because the switch doesn’t forward unicast packets except to the intended destination, Computer1 never received the ping packets between Computer2 and Computer3.
10. Stop the capture in Wireshark. Close all open windows and shut down the computers.



Hands-On Project 2-4: Examining Hub and Switch Indicator Lights and Uplink Ports

Time Required: 30 minutes

Objective: Examine the indicator lights of a hub or switch and understand the purpose of the uplink port.

Required Tools/Equipment: Three computers (minimum) with Ethernet NICs installed; 10/100 or 10/100/1000 NICs are preferable, but 10 Mbps NICs will also work. Computers can be configured with static IP addresses or can use APIPA addresses. Four patch cables and one crossover cable, two 10/100 hubs with uplink switch, and a 10/100 switch.

Description: In this project, you view the indicator lights of hubs and switches. Ideally, your hub has indicators for link status, activity, and collisions. In addition, if your hub has an uplink port, you test its function. Like the previous two projects, this project can be done in groups or as a class demonstration.

1. The computers should be shut down and one hub should be plugged in and turned on, if necessary. Connect all three computers to the hub with patch cables, but don't use the uplink port on the hub. Turn on the computers.
2. Examine the hub's indicator lights. A link status light should be glowing for each port a computer is connected to. Next, examine the indicator lights on the NIC, which should also be glowing to indicate a good connection. See whether the hub's indicator lights vary for different connection speeds. Write the link status light's color and the connection speed, if available, in the following chart:

Computer	Link status light's color	Connection speed
Computer1		
Computer2		
Computer3		

3. Generate some traffic by using ping commands on each computer. At each computer, open a command prompt window and ping another computer (Computer1 ping Computer2, Computer2 ping Computer3, and Computer3 ping Computer1, for example) by typing `ping -n 20 IPaddress` and pressing **Enter**. Examine the activity indicator lights, which should blink as data is received. (On hubs combining the activity indicator with the link status, network activity just causes the link status indicator to blink.) The `-n 20` option in the ping command specifies sending 20 ICMP packets instead of just 4.
4. Next, if your hub has collision indicators, try to get them to glow. Note that the pings must be sent from each computer at the same time for a collision to occur. At each computer, perform the same ping, but this time, type `ping -n 20 -l 60000 IPaddress` and press **Enter**. The `-l 60000` (lowercase "L") option makes each ping packet 60,000 bytes in length. Even with these large amounts of data being transferred, you might not see a collision. Remember that a collision occurs when two or more computers send data simultaneously, which isn't permitted when using a hub. However, if your hub and NICs are operating at 100 Mbps, data is transferred so quickly that producing a collision might be difficult.
5. While leaving the first hub powered on, power on the second hub. With a regular patch cable, connect the first hub to the second hub, but don't use the uplink port. In most cases, you won't see the link lights glow at the ports where the two devices are connected. To fix this problem, plug one end of the patch cable into the uplink port on one hub (not on both hubs) and set the switch to the uplink position. You should now have connectivity between the hubs, and the link lights should be on.

6. If your hubs don't have an uplink port, you can connect two hubs with a crossover cable. To do this, disconnect the two hubs, and using a crossover cable from your instructor, connect each end of it to regular (not uplink) ports on the two hubs. The link lights should glow. (You learn more about patch and crossover cables in Chapter 4.)
7. List any other indicator lights you find on the hub and what these lights tell you:

8. Disconnect the computers from the hubs and put the hubs away. Connect the computers to the switch, and then power on the switch.
9. Along with link status lights, most switches have lights on each port to indicate whether the port is operating in full-duplex or half-duplex mode. If your switch has these indicators, find them and try to determine in which mode your NIC and switch are communicating. Most NICs and switches support full-duplex communication, and this mode is chosen automatically.
10. List any other indicator lights you find on the switch and what these lights tell you:

11. Close all open windows, and shut down the computers.



Wireless Access Points

As you probably know, not all networks require a cable tethering the computer to a switch or hub. Wireless networks have become ubiquitous in college and corporate campuses and in many public locations, such as airports and libraries. At the heart of a wireless LAN is the wireless **access point (AP)**. An AP is a lot like a hub, in that all computers send signals through it to communicate with other computers. The obvious difference is that signals don't travel through a physical medium; they travel through the airwaves as radio signals.

Most small business and home networks with wireless networks use a device typically called a wireless router that combines the functions of an AP, a switch, and a router (see Figure 2-10). Wireless routers can usually be identified by the two or more antennae on the device. These devices are usually used with a cable or DSL modem to provide wireless access to the Internet. Large businesses use dedicated APs to give users wireless access to the corporate network as well as the Internet.

Wireless networks rarely stand by themselves. They're almost always connected to a wired network at some point. APs typically have one or more connectors for connecting to a wired Ethernet network.

Basic AP Operation

An AP is much like a wired hub, in that all stations hear all network data transmitted by all other wireless devices in the network. All communication goes to the AP, which then retransmits or repeats the transmission to the destination station. However, unlike hubs, communication



Figure 2-10 A wireless router combines an access point, a switch, and a router

Courtesy of Hyperline Systems

between two stations requires an extra step. The destination device sends an acknowledgement back to the sending device to indicate the frame was received. When the sending device receives the acknowledgement, it knows that no error or collision has occurred.

Some wireless configurations require additional handshaking between two communicating devices. Before a computer can transmit data to the AP, it must first send a short **request to send (RTS)** message to let the AP know it intends to transmit data. If no other stations are trying to send data, the AP responds with a **clear to send (CTS)** message letting the requesting station (and all other stations on the network) know that it can send data. The RTS and CTS messages are sent in addition to the acknowledgement the receiving computer sends. Imagine if you had to communicate in this fashion while speaking. Before each sentence you wanted to speak, you would have to ask a moderator whether you could speak, and the moderator would have to answer affirmatively. Then, after each sentence, the moderator would have to acknowledge that you were heard before you could speak the next sentence. Conveying any real information would take much longer because so much time would be wasted on the overhead required by the communication rules. Fortunately, most wireless networks don't use the RTS/CTS configuration, but it's available as an option on some higher-end APs and wireless NICs.

Wireless APs and Network Bandwidth All the extra chatter required to send data in a wireless network slows communication quite a bit. In fact, the effective bandwidth (that is, the bandwidth used for actual data transmission) is about half the physical bandwidth. Keep in mind, too, that wireless network bandwidth is shared, as with a hub.

Most APs operate at anywhere from 11 Mbps to several hundred Mbps. So a wireless AP operating at 11 Mbps shares this 11 Mbps with all computers in the wireless network. Therefore, if 11 stations are connected to an 11 Mbps wireless network, each station has 1 Mbps of effective bandwidth; with all the extra network traffic (acknowledgements and possible RTS/CTS messages), however, you must halve this amount, leaving only about 500 Kbps of effective bandwidth. That's why developers are constantly striving to get more

bandwidth out of wireless networks. In recent years, the performance of basic 11 Mbps wireless networks has increased to more than 100 Mbps, and APs that can operate at speeds of 300 Mbps and higher are becoming common, with speeds over 600 Mbps on the horizon.



Wireless networking is a big subject to tackle, and you learn more about wireless networking standards and technologies in Chapter 3.



Network Interface Cards

As a networking professional, you must understand what a network interface card does and how it works as well as what's involved in configuring a NIC for special network situations in which the default configuration is inadequate. Although most NICs are built into a computer's motherboard, they occasionally fail or additional NICs are needed for your application, so you should know how to install a new NIC, too. The following sections discuss the basic operation of a NIC along with its device driver, its most common features, and some configuration options.

NIC Basics

Attaching a computer to a network requires a **network interface card (NIC)** to create and mediate the connection between the computer and the networking medium. The networking medium might be copper wire, fiber-optic cable, or the airwaves, but in all cases, data is represented as bit signals that the NIC transmits or receives.

For incoming data, the NIC must be able to interpret the signals used for the network medium, which are electrical for copper wire, light for fiber-optic cable, or radio waves for wireless networks. These signals are then converted to bits and assembled into frames. For outgoing data, the NIC converts frame data into bits and transmits these bits to the medium in the correct signal format.

The following list summarizes the tasks a NIC and its driver perform:

- Provide a connection from the computer to the network medium.
- For incoming messages, receive bit signals and assemble them into frames, verify the frame's destination address, remove the frame header and trailer, and transfer the packet to the network protocol.
- For outgoing messages, receive packets from the network protocol and create frames by adding source and destination MAC addresses and error-checking data.
- Convert the frame data into bit signals in a format suitable for the network medium and transmit the signals.

Figure 2-11 shows a NIC handling incoming data, and Figure 2-12 shows a NIC handling outgoing data.

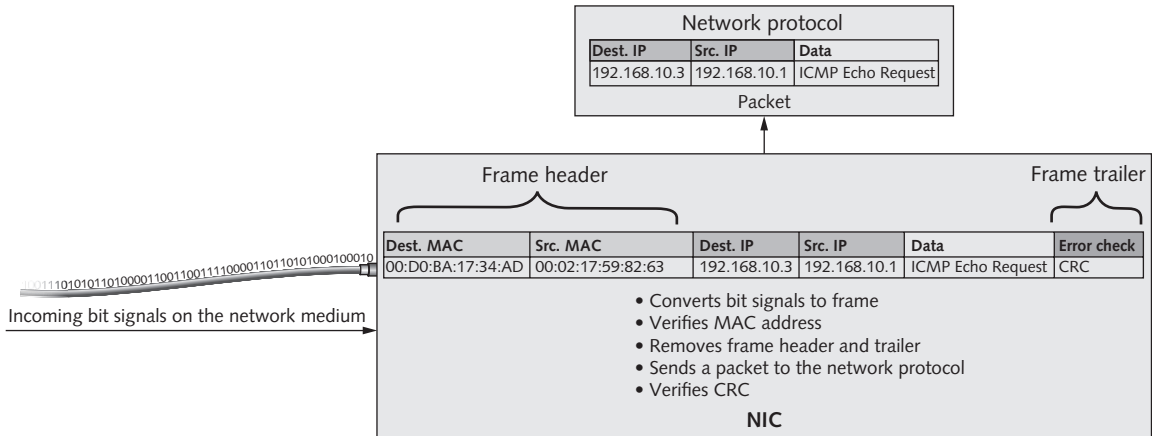


Figure 2-11 A NIC handles incoming data from the network medium

Courtesy of Course Technology/Cengage Learning

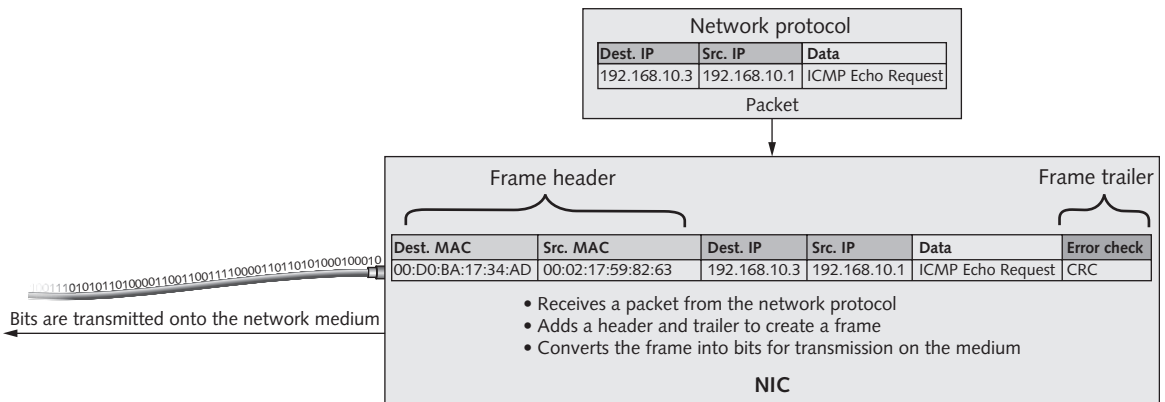


Figure 2-12 A NIC handles outgoing data to be sent to the network medium

Courtesy of Course Technology/Cengage Learning

Simulation 5 on the book’s CD shows an animated version of Figures 2-11 and 2-12.



Simulation 5: How a NIC works

NICs and MAC Addresses Aside from the tasks described previously, a NIC has the important function of giving a computer a MAC address, an integral part of each NIC. NIC manufacturers ensure that every NIC has a unique address because networks won’t function correctly if duplicate MAC addresses exist. The MAC address is stored in read-only memory (ROM) on the NIC. Because the address is said to be burned into

memory, it's sometimes referred to as the burned-in address (BIA). The MAC address is composed of two 24-bit numbers:

- A 24-bit manufacturer ID called an organizationally unique identifier (OUI)
- A 24-bit serial number assigned by the manufacturer

The 48-bit MAC address is expressed in hexadecimal notation, usually as six two-digit alphanumeric characters separated by dashes or colons, such as 04-40-31-5B-1A-C4. The first three two-digit groups represent the OUI, and the last three are the unique serial number.

**TIP**

You can find the manufacturer of a NIC by its MAC address. Go to <http://standards.ieee.org/regauth/oui/index.shtml> and enter the first three numbers (six digits) of a MAC address, separated by dashes.

The NIC as Gatekeeper When a frame arrives at a NIC, the NIC doesn't simply read the frame and send a packet to the network protocol. It examines incoming network frames and checks the frame's destination MAC address to see whether it matches its built-in MAC address. The NIC acts as a gatekeeper and permits inbound communications to pass through the interface only if the destination MAC address meets these criteria:

- The destination MAC address in the frame matches the NIC's built-in MAC address.
- The destination MAC address in the frame is the broadcast address.
- The NIC is operating in promiscuous mode.

A frame with a destination MAC address composed of all binary 1s or FF-FF-FF-FF-FF-FF in hexadecimal is a **broadcast frame**. Broadcast frames are intended to be processed by all computers on the network. Destination MAC addresses intended for a single computer are called **unicast frames**. Most NICs can operate in what's called **promiscuous mode**—essentially, this mode turns off the gatekeeper functions and enables the NIC to process all frames it sees. This mode is used by software called a protocol analyzer or packet sniffer (such as the Wireshark program you installed in Hands-On Project 2-1) that captures frames and displays their contents for the purposes of troubleshooting and learning.

**NOTE**

A third type of MAC address, called a multicast address, is intended to be processed by a group of computers running a particular application or service. These MAC addresses are identified by a value of 1 in the rightmost bit of the first two digits, such as 01-22-33-44-55-66.

NIC Indicator Lights Like hubs and switches, NICs have indicator lights to show status information. Although the details vary across NIC models, NICs usually have a link status indicator and an activity indicator. The link status light is usually green when the NIC has a valid connection between the network medium and another device, such as a hub or switch. NICs usually also have an indicator light that flashes when the NIC detects network activity. As with hubs and switches, the link light and activity indicators are sometimes combined.



Some NICs supporting multiple speeds, such as 100 Mbps and 1000 Mbps, have a separate link light for each speed so that you can determine at what speed the NIC is connected to the hub or switch. In other cases, the link light indicates the connection speed by using a different color, such as amber for 100 Mbps and green for 1000 Mbps. There's no standard for NIC indicator lights, so you should consult the NIC's documentation to determine their purposes.

Selecting a NIC

The average user might never have to install a NIC because most NICs are built into the motherboard. However, onboard interfaces can fail or prove inadequate for how the computer is to be used. For example, the built-in NIC might operate at only 100 Mbps, and you want the interface to operate at 1000 Mbps, or there might be only one built-in NIC, and you need two or more NICs for a server. In these cases, you need to select a NIC with the correct bus interface to connect to your computer.

The connection the NIC makes to the motherboard is the bus connection, and when a NIC receives data, the data must be transferred to the bus and then to the CPU so that it can be processed. The bus speed determines how fast data can be transferred between components. When data is to be transmitted to the network, it goes from the CPU through the bus and to the NIC before being sent to the network medium.

Several bus types are in common use on PC motherboards. Chapter 7 delves into specifics of the bus architectures commonly used for NICs, but for now, you just need to know that Peripheral Component Interconnect (PCI) or PCI Express (PCIe) are the ones you're most likely to encounter when installing an internal NIC. To make installation easier, you might want to choose a NIC that connects to your computer via an external USB connector.

What's most important in selecting a NIC to install is that you choose one your system supports, both in bus type and availability of device drivers for your computer's OS. The NIC's specifications tell you the bus type, and the packaging or manufacturer's Web site lists the OSs for which device drivers are available.

A close second in importance is selecting a NIC that's suitable for the role your computer will play in the network. If the computer is a typical desktop system, a standard \$10 PCI NIC that operates at speeds of 10/100/1000 Mbps is probably sufficient. For servers or high-performance workstations, consider a NIC that has onboard memory and multiple ports and connects with the faster PCIe bus.

NIC Drivers

Installing a driver for a NIC is usually easy. Most OSs ship with drivers for a wide range of NIC manufacturers and models. Also, most NICs include drivers for the most common OSs, including current Windows and Linux versions. In most cases, you simply need to shut down your computer, install the NIC, and restart the computer. If the OS has a suitable driver available, it's installed automatically. If not, you're usually prompted to insert a CD/DVD containing the driver files.

After the drivers are installed, the NIC is usually ready to function without further configuration. In Windows 7, you can verify that your NIC is installed in the Network Connections window, which you access by clicking "Manage network connections" in the Network and Sharing Center (see Figure 2-13).

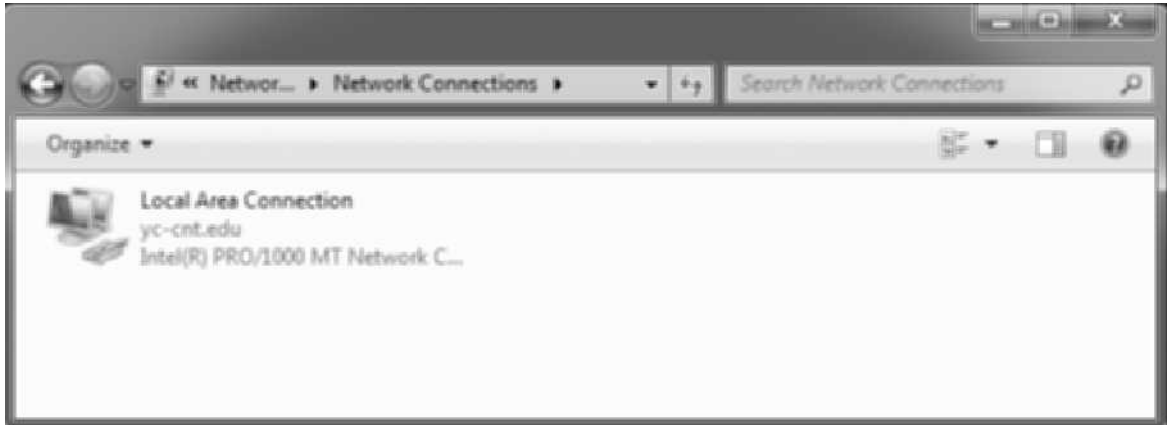


Figure 2-13 The Network Connections window in Windows 7

Courtesy of Course Technology/Cengage Learning

In Windows, each NIC is assigned a connection name. The first NIC in the system is assigned “Local Area Connection.” If you have a second NIC, it’s assigned “Local Area Connection 2,” and so forth, but you can change the name to be more descriptive. To view a connection’s settings, right-click the connection and click Properties to open the dialog box shown in Figure 2-14. The Connect using text box shows the type of NIC that’s installed. To change the NIC’s settings and its driver, click the Configure button. Common NIC configuration options are discussed in Chapter 7.



Figure 2-14 The Local Area Connection Properties dialog box

Courtesy of Course Technology/Cengage Learning

Wireless NICs

The selection process for a wireless NIC differs somewhat from selecting a wired NIC. Wireless NICs are most often built into laptops and other portable computers, but you still might want to install one on a desktop computer, particularly in a small business that uses wireless networking exclusively.

Wireless NICs must be chosen according to the type of wireless AP you have installed. Most are described in terms such as Wireless-g or Wireless-n or perhaps 802.11 b/g. The letters g, n, and b refer to the wireless networking standard the device supports. These standards support increasing speeds and features in this order from slowest to fastest: b, g, and n. Wireless-b, or 802.11b, is among the earliest wireless standards and supports up to 11 Mbps transfer rates. Wireless-g, or 802.11g, came next and supports up to 54 Mbps transfer rates. Wireless-n, or 802.11n, is a new standard supporting speeds from 54 Mbps to more than 300 Mbps. These standards are backward-compatible, meaning that a Wireless-n NIC works with a Wireless-b AP, albeit at the slower 11 Mbps speed, and a Wireless-b NIC works with a Wireless-n AP. Chapter 3 covers these standards in more detail.

Unlike a wired NIC, a wireless NIC often requires a few more steps before a successful connection can be made. Figure 2-15 shows the Network and Sharing Center in Windows 7. You can click the “Connect to a network” link to open a dialog box listing all the wireless APs in range of your wireless NIC. To connect to a wireless network, click the network name and click Connect. The name assigned to a wireless network, called the **service set identifier (SSID)**, is configured on the AP. You might also be prompted for a security key or a username and password, depending on the network’s security configuration. When security is enabled on a WLAN, communication is encrypted so that unauthorized parties can’t connect or easily interpret the data traveling through airwaves. The security key serves as a decryption key, allowing a client to access the wireless network. You learn more about wireless networks and how to configure them in Chapter 3.

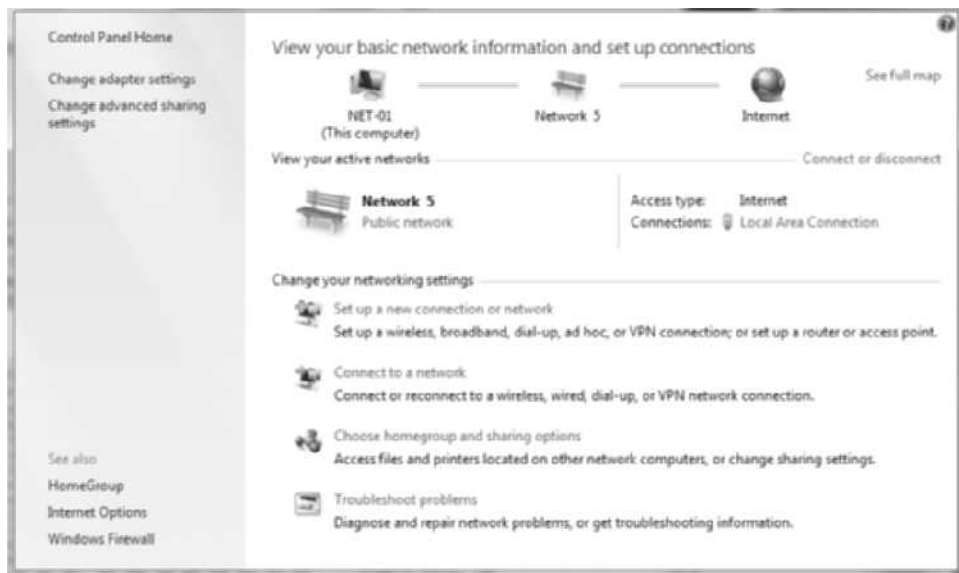


Figure 2-15 Viewing network connections in the Network and Sharing Center

Courtesy of Course Technology/Cengage Learning



Hands-On Project 2-5: Connecting to a Wireless Access Point

Time Required: 15 minutes

Objective: Install a wireless NIC and connect to an access point.

Required Tools/Equipment: Two or more computers with 802.11 wireless NICs installed. USB wireless NICs work well, as they don't require opening the computer case. Laptops with built-in wireless NICs will also do. One wireless AP or wireless router configured with the SSID "NetEss." The 802.11 standard supported doesn't matter as long as the AP is compatible with the NICs. Windows 7 is the preferred OS, but some steps can be changed to accommodate other OSs. The computers shouldn't be connected to a hub or switch.

Description: In this project, you connect to a wireless AP and test the connection by pinging another computer connected to the same AP.

1. Start your computer and log on as an administrator. If the wireless NIC isn't installed yet, install it according to your instructor's instructions.
2. After the wireless NIC has been installed, click the network connection icon in the notification area to display a list of available wireless networks (see Figure 2-16).



Figure 2-16 List of available wireless networks in Windows 7

Courtesy of Course Technology/Cengage Learning

3. Click the **NetEss** wireless network. (Remember that the wireless network name is called the SSID.) A message is displayed, stating that information sent over the network might be visible to others because it's not secured with encryption. You secure the network later, so click the **Connect** button.
4. After a short time, you should see the Set Network Location window. The network location can be Home, Work, or Public and is used to set up firewall rules for the connection. Click **Work network**, and then click **Close**.



5. You're now connected to the NetEss wireless network. To test the connection, get the IP address of another computer connected to the wireless network and ping this address. Alternatively, you can ping the router, which should be at address 192.168.1.1. Ask your instructor for the correct address if pinging 192.168.1.1 doesn't work.
6. Click the wireless network connection icon in the notification area and click **Open Network and Sharing Center**. You'll see a window similar to Figure 2-15, shown previously.
7. Click **See full map** at the upper right. Figure 2-17 shows the network map for the wireless connection. If other computers are attached to the wireless network, they're displayed in the map. The dotted lines indicate a wireless network connection; a wired connection is shown with solid lines. If the network map can't be created, it might be because the NIC driver doesn't support the protocol used to create the map.

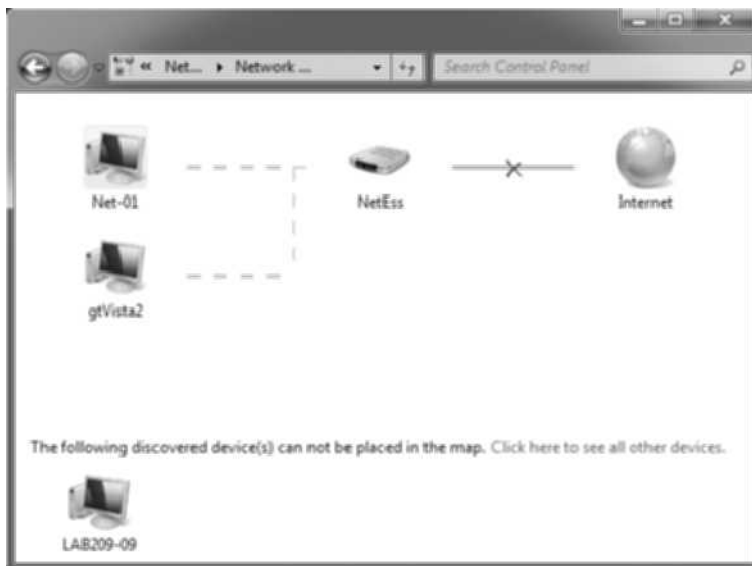


Figure 2-17 Viewing the network map

Courtesy of Course Technology/Cengage Learning

8. Close all open windows.



Hands-On Project 2-6: Examining Properties of Your NIC

Time Required: 15 minutes

Objective: View the properties of your NIC in Windows and look up its vendor by using the MAC address.

Required Tools/Equipment: Your classroom computer; no other tools or equipment are required. Windows 7 is the assumed OS, but this project can also be done in Windows Vista or Windows XP.

Description: When describing a NIC as a component of networking, it means both the hardware NIC and its driver. NIC drivers are configured in the OS in which they're installed and control certain operational aspects of the network interface as a whole. In this project, you examine the properties of your installed NIC. You also use your NIC's MAC address to

look up the vendor. Not all NICs or NIC drivers are equivalent in features, so your NIC might have more or fewer features than are described here.

1. Turn on your computer and log on as **NetAdmin**, if necessary.
2. Click the network connection icon in the notification area and click **Open Network and Sharing Center**.
3. In the Network and Sharing Center, click the **Change adapter settings** link on the left. Right-click **Local Area Connection** and click **Status** if necessary.
4. The Local Area Connection Status window shows a summary of information of your network connection. To see more information about the connection, click **Details**.



The `ipconfig/all` command shows the same information as the Local Area Connection Status window.

5. The Network Connection Details window shows information about your connection, including the NIC model, the physical (MAC) address, and your IP address configuration. Write down your MAC address, which you use later to look up the NIC vendor. Review the remaining information, and then click **Close**.

- MAC address: _____

6. In the Local Area Connection Status window, click **Properties**. In the Local Area Connection Properties dialog box, click the **Configure** button under the Connect using text box. In the Network Connection Properties dialog box, click the **Advanced** tab (see Figure 2-18). Your NIC might have fewer, more, or different options.



Figure 2-18 Viewing advanced settings in the Network Connection Properties dialog box

Courtesy of Course Technology/Cengage Learning

7. Review the available properties for your NIC. When you select a property, you can see its possible values in the Value drop-down list.
8. Click **Link Speed & Duplex** (if this property is available), and then click the **Value** down arrow to see the possible values. On most NICs, the default value is Auto Negotiation, which means the NIC and hub or switch exchange signals to determine the optimum operational mode. Other modes usually include combinations of 10, 100, and 1000 Mbps and full- and half-duplex. Normally, you don't need to change these values unless auto negotiation fails to work. If this happens, you'll probably see the link status light change from on to off repeatedly or never turn on at all.
9. Click the **Locally Administered Address** property. (It might also be referred to as network address, physical address, or MAC address.) In most cases, this property's value is set to Not Present. You can use this property to override the NIC's burned-in MAC address by entering a new address in the Value text box. Normally, however, you shouldn't override the burned-in MAC address because if you duplicate an existing address accidentally, it can cause a loss of communication. Click **Cancel** to close Network Connection Properties.
10. Close the Local Area Connection Status window and the Network Connections window.
11. Start a Web browser, and go to www.coffer.com/mac_find.
12. In the MAC Address or Vendor to look for text box, type the first six digits of the MAC address you wrote down in Step 5. You don't need to enter the hyphen between each pair of digits, but you do need to enter the leading zeros. Click **string** to find the vendor of the MAC address. Knowing the vendor can help you track down devices that might be causing problems on your network.
13. Close all open windows. If you aren't going on to the next project, shut down your computer; otherwise, stay logged on.

Routers

Routers are the most complex devices discussed in this chapter. Hubs and switches connect computers to the LAN; routers connect LANs to one another. Routers typically have two or more network ports to which switches or hubs are connected to form an internetwork. Figure 2-19 is a diagram of an internetwork, with two LANs connected via a router. Each LAN in this example uses switches to connect workstations and a router port to the LAN. LAN 2 has two switches that are connected.




Figure 2-19 Two LANs connected by a router to make an internetwork

Courtesy of Course Technology/Cengage Learning

Routers enable multiple LANs to communicate with one another by forwarding packets from one LAN to another. They also forward packets from one router to another when LANs are separated by multiple routers. The Internet is built on a vast collection of LANs, all interconnected via routers. Figure 2-20 shows a small business network connected to its Internet service provider (ISP), followed by connections to several other Internet routers and ultimately to a Web server on the Course.com network.




Figure 2-20 Routers interconnect LANs to form the Internet

Courtesy of Course Technology/Cengage Learning



Recall from Chapter 1 that the Internet and its complex arrangement of routers is usually depicted as a cloud in most diagrams to hide the complex web of routers and devices that make up the Internet.

On the surface, it might seem as though switches and routers perform a similar function, but in reality, they have very different jobs to do and how they work with network data differs substantially. The following points summarize the key properties and features of a router versus a switch:

- Routers connect LANs, and switches connect computers.
- Routers work with logical (IP) addresses rather than physical (MAC) addresses, as switches do.
- Routers work with packets rather than the frames that switches work with.
- Routers don't forward broadcast packets, but switches do.
- Routers use routing tables, and switches use switching tables.

The following sections discuss how and why routers are used to connect LANs and how routers use routing tables. Simulation 6 on the book's CD shows basic router operation.



Simulation 6: Router operation in a simple internetwork

Routers Connect LANs

Switches are the device of choice to connect computers to create a LAN. However, if you look at a LAN as a group of people with similar interests getting together to converse and share information, there's a point at which the group can become too large for effective communication. For example, in most group discussions, several conversations often occur at once, but periodically, someone wants to speak to the entire group. For small groups with tightly coupled interests, this method works well, but as the group gets larger, the frequency of group announcements can affect the flow of communication adversely. This is particularly true when only a small subset of the group is interested in the announcement, yet the whole group must stop to listen. In this case, communication can be enhanced by dividing the large group into smaller groups of similar interests in different locations. By doing so, announcements to the entire group are contained in the small group and need not interrupt other groups' conversations. You can look at these announcements as network broadcast frames that switches (and hubs) are obliged to forward to all connected stations.

Breaking a large group into smaller groups works well until a member of one group must communicate with a member of another group. A messenger could be used to get a message from one group to another and would normally forward only messages directed to a person in another group, not announcements to the entire group. This messenger is analogous to a router in an internetwork, and like the messenger, the router doesn't forward announcements (broadcasts); it forwards only messages destined for a particular address.

Review Figure 2-21, which shows a large LAN with all workstations and servers connected via switches. All these switches are connected through the switch the servers are on. This arrangement works fine if the number of workstations on each switch doesn't exceed about 20. However, if each group of computers represents as many as 25 computers (making 150 total workstations), announcement messages (broadcasts) will probably start affecting communication efficiency. Remember that each time a computer sends a broadcast frame, the switch forwards it out all connected ports so that all computers eventually receive the broadcast. One deleterious effect of broadcast frames is that when a computer receives a broadcast, a CPU interrupt occurs, causing the computer to stop what it's doing to service the interrupt. If enough interrupts occur in a short period because of many broadcast frames, the computer's overall performance can suffer as a result of the CPU having to service the interrupt and process the broadcast.




Figure 2-21 A large LAN connected by switches

Courtesy of Course Technology/Cengage Learning

Now review Figure 2-22, in which the network has been redesigned for efficiency. Workstations have been organized so that each department's users are grouped together, and servers have been configured so that each department's frequently accessed documents and applications reside on the departmental server. In this arrangement, the switches for each LAN allow all computers in the LAN to communicate with one another and forward important broadcast frames so that all computers receive the announcement, but broadcasts aren't forwarded to other LANs because the router doesn't forward broadcast frames. However, the router does allow communication

between LANs so that if a management computer needs to access data on the marketing server, it can do so.




Figure 2-22 Three smaller LANs connected by a router
Courtesy of Course Technology/Cengage Learning

Routers Create Broadcast Domains

The scope of devices to which broadcast frames are forwarded is called a **broadcast domain**. Because routers don't forward broadcasts, router interfaces are the delimiter for broadcast domains. In other words, each router interface in a network creates another broadcast domain. Figure 2-23 shows the same network as Figure 2-22, with circles around each broadcast domain. (Note that Figure 2-21, with no routers at all, is a single broadcast domain.) Chapter 7 describes broadcast domains and how to create them with advanced switch features.

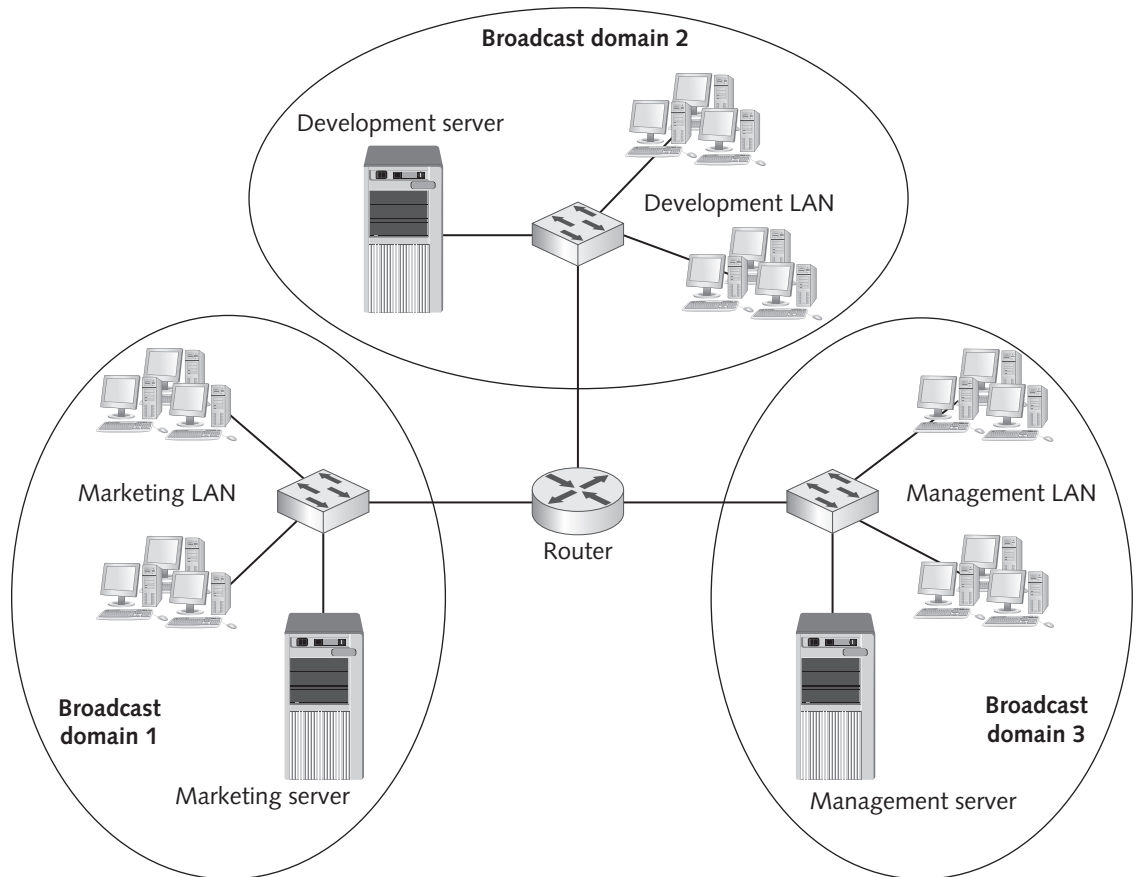


Figure 2-23 Each router interface creates a broadcast domain

Courtesy of Course Technology/Cengage Learning

Routers Work with IP Addresses and Routing Tables

Switches, as you know, maintain a switching table of MAC address/switch port pairs to determine where to forward frames in a LAN. Routers maintain routing tables composed of IP network addresses and interface pairs to determine where to forward packets in an internetwork.

Routers have two or more interfaces, with each interface connected to a different network. When a router receives a packet on one interface, it looks at the destination IP address in the packet to determine which network the packet is addressed to. Then it forwards the packet out of the interface that its routing table indicates is the best way to get the packet to its destination. Figure 2-24 shows the same internetwork as Figure 2-23, with each LAN assigned a network number, and an example of what the routing table might look like. The router's three interfaces are labeled EthA, EthB, and EthC.




Figure 2-24 An internetwork with a routing table

Courtesy of Course Technology/Cengage Learning

When the router receives a packet from a computer in Network 1.0 that has a destination address of a computer in Network 3.0, the router looks in its routing table and discovers that Network 3.0 can be found via the EthC interface. The router then forwards the packet out its EthC interface to reach the intended computer.

This routing table has been simplified for demonstration purposes; routing tables have more information than simply the network number and interface name. In addition, network numbers are derived from IP addresses and contain more numbers than shown. Chapter 7 has additional details about how routers work, and Chapter 5 discusses IP addresses and network addresses in more depth.

You might wonder what happens when a router isn't connected to the network the packet is addressed to. Figure 2-25 illustrates this situation and shows what the routing table would look like on each router between the source and destination networks. In this example, if a computer on Network 1.0 sends a packet to a computer on Network 5.0, router R1 receives the packet and looks up Network 5.0 in its routing table. According to its routing table, it forwards the packet out its WAN A interface. Router R2 receives the packet and forwards it out its WAN B interface, as specified by its routing table, and finally, router R3 receives the packet and forwards it out its EthA interface to the destination computer.




Figure 2-25 Packets are forwarded through multiple routers

Courtesy of Course Technology/Cengage Learning

Default Routes Routers on a corporate network might have a routing table entry for every network in the organization, but what about routers connected to the Internet? The Internet is composed of hundreds of thousands of networks, and routers on the Internet are responsible for getting packets from any network to any other network. Although it might be technically possible for routers to have a record of every network in the Internet, having such large routing tables isn't practical. To solve this dilemma, routers can have a special routing table entry called a **default route**, which tells a router where to send a packet with a destination network that can't be found in the routing table. The default route usually leads to another router with the network address in its table or results in the packet being sent to another default route and so on, until it reaches a router that has the network address in its routing table.

Network Unreachable Most routers are configured with a default route, but not always. If a router receives a packet with a destination network address that isn't in its routing table and no default route is configured, the router simply discards the packet. The router also sends a message to the sending station informing it that the network is unreachable. By doing so, the sender is made aware that the destination network doesn't exist or the routers must be configured differently to access the destination network.

Default Gateway Just as a router must know where to forward a packet it receives, a workstation must know when to send a packet to the router instead of simply addressing the packet and sending it to the local LAN. When a workstation has a packet ready to send, it compares its own IP address with the destination IP address in the packet. If the two addresses are on the same network, the workstation gets the destination computer's MAC address and sends the frame to the local LAN to be delivered to the destination. If the two addresses are on separate networks, the workstation must instead get the router's MAC address and send the frame to the router, which then tries to get the packet to the destination network. In this case, the workstation must know the address of a router. The **default gateway** in a computer's IP address settings must be set to the address of a router to which the computer can send all packets destined for other networks. If the default gateway doesn't have a valid address of a router, the computer can communicate only with computers on the same LAN. In Chapter 5, you learn more about how a computer determines its network address.

This chapter has explained the basic operation of the most common network hardware components. There's more to learn about all these components, but before you delve deeper into network hardware, examining other aspects of networking is helpful. The next several chapters discuss network topologies and technologies, network media, protocols, and networking standards, among other topics.



Hands-On Project 2-7: Communicating over a Router

Time Required: 20 minutes

Objective: Configure workstations to communicate with one another through a router.

Required Tools/Equipment: Three workstations, two hubs or switches, a router, and five patch cables are required. The router can be the same router/AP you used for Hands-On Project 2-5. Assuming you're using a typical home network router/AP, such as a Linksys WRT54GL, the router should be set up so that the WAN interface is assigned the address 192.168.2.1 with subnet mask 255.255.255.0, and the LAN interface is left as the default 192.168.1.1 address.

Description: This project requires some setup by your instructor. You should verify with your instructor that it's complete before starting this project. In this project, you configure workstations to communicate with each other through a router. The router is configured to support two networks: 192.168.1.0 and 192.168.2.0. Computer1 and Computer2 are configured to operate in the 192.168.1.0 network, and Computer3 is configured to work in the 192.168.2.0 network. Figure 2-26 shows the initial network setup, in which all three computers are connected to the same hub or switch. Cable the network as shown. The router should already be configured.

Network 192.168.1.0

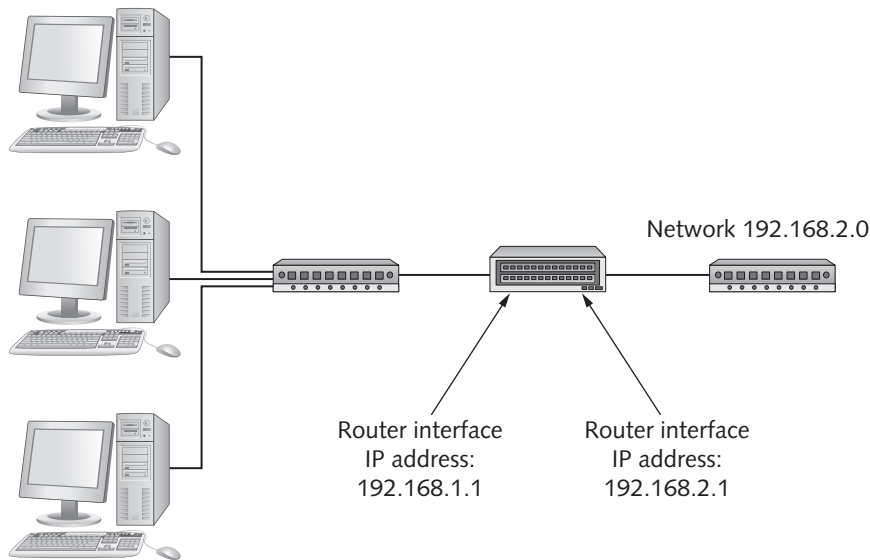


Figure 2-26 Initial network setup for Hands-On Project 2-7

Courtesy of Course Technology/Cengage Learning

1. Start all three computers, if necessary. To configure the IP address of each computer, click the network connections icon in the notification area and click **Open Network and Sharing Center**. Click the **Change adapter settings** link on the left. Right-click **Local Area Connection** and click **Properties**. Double-click **Internet Protocol Version 4**, and then click **Use the following IP address**. For now, just set the IP address and subnet mask, using the following values:
 - Computer1: IP address 192.168.1.11, subnet mask 255.255.255.0
 - Computer2: IP address 192.168.1.12, subnet mask 255.255.255.0
 - Computer3: IP address 192.168.2.21 subnet mask 255.255.255.0
2. After you have entered these values, click **OK** twice and close all windows.
3. To test your configuration, open a command prompt window on Computer1 and Computer2, and ping each other's IP address. The ping should be successful. If it's not, verify that the IP address settings are correct by typing **ipconfig** and pressing **Enter** and comparing the values you see with the ones listed in Step 1. From both computers, type **ping 192.168.1.1** and press **Enter** to verify that they can communicate with the router.
4. On Computer1, ping Computer3 by typing **ping 192.168.2.21** and pressing **Enter**. You should get a message that the ping failed or timed out. The reason the ping between Computer1 and Computer3 failed is that the two computers are configured to be on different networks. In this case, Computer1 is configured to be on network 192.168.1.0, and Computer2 is configured to be on network 192.168.2.0. When two computers are configured to be on separate networks, their connection to each other must be separated

by a router. Move Computer3 to the other network by plugging the cable from Computer3 into the other hub or switch so that your network configuration now looks like Figure 2-27.




Figure 2-27 Corrected network setup

Courtesy of Course Technology/Cengage Learning

5. Try the ping again from Computer1 to Computer3. Again, you get an error because one piece of the IP address configuration has been omitted. When a computer needs to send a packet to a device on another network, it must know the address of the router to which to send the packet. This address is called the default gateway. To set the default gateway for all three computers, follow the instructions in Step 1 to get to the IP settings. In the Default gateway text box, enter the following values:
 - Computer1: **192.168.1.1**
 - Computer2: **192.168.1.1**
 - Computer3: **192.168.2.1**
6. After you have finished configuring the default gateway for all three computers, you should be able to ping from Computer1 to Computer3 and from Computer2 to Computer3 and vice versa. Try it now, and write down your results.

7. Next, try another command that shows the route your packet travels to get to the other computer. From Computer1 and Computer2, type **tracert 192.168.2.21** and press **Enter**. From Computer3, type **tracert 192.168.1.11** and press **Enter**. You'll see a few lines of output showing that the packet had to travel through the router to get to its destination. This command is used again in the next project.
8. Close all open windows on all three computers.



Hands-On Project 2-8: Using Trace Route to See How Packets Travel Through the Internet

Time Required: 10 minutes

Objective: Use the Trace Route program to see the routers packets must travel through to get from your computer to a destination on the Internet.

Required Tools/Equipment: Your classroom computer, Internet access, and a valid DNS server; no other tools or equipment are required

Description: The importance of routers is made clear when you need to access servers on the Internet. The Trace Route program (tracert.exe) lists each router your packets must travel through between your computer and an Internet server.

1. Log on to your computer as **NetAdmin**, if necessary, and open a command prompt window.
2. Type **tracert www.yahoo.com** and press **Enter**. You should see output that's similar to Figure 2-28, but the details will vary, depending on your location. In this output, there are five columns of information. The first column is just a count of how many routers the packet traversed. The second, third, and fourth columns show the amount of time in milliseconds (ms) the router took to respond. Three packets are sent, so three times are listed. The last column is the router's IP address or name and IP address.

```

C:\Windows\system32\cmd.exe
Tracing route to any-fp.ua1.h.yahoo.com [67.195.145.138]
over a maximum of 30 hops:
  0  3  ns    1  ns    1  ns    172.31.1.258
  1  2  ns    1  ns    2  ns    172.16.0.2
  2  2  ns    2  ns    2  ns    ycexpress.video.yc.edu [198.68.121.28]
  3  5  ns    5  ns    5  ns    phn-edge-06.inet.qwest.net [67.135.198.249]
  4  5  ns    5  ns    5  ns    phn-core-01.inet.qwest.net [205.171.12.77]
  5  15  ns   15  ns   15  ns    lap-brdr-03.inet.qwest.net [67.14.22.78]
  6  26  ns   16  ns   16  ns    if-2-5.icore1.eql-losangeles.as6453.net [206.82.129.33]
  7  38  ns   34  ns   43  ns    vlan1114.icore1.pdi-paloalto.as6453.net [209.58.17.5]
  8  26  ns   26  ns   26  ns    if-15-0-0-59.ncore3.pdi-paloalto.as6453.net [216.6.29.49]
  9  26  ns   26  ns   26  ns    ix-11-0-4.ncore3.pdi-paloalto.as6453.net [64.86.84.158]
 10  27  ns   61  ns   51  ns    unknown-216-115-107-73.yahoo.com [216.115.107.73]
 11  38  ns   26  ns   26  ns    ir3.fp.vip.spl.yahoo.com [67.195.145.138]
Trace complete.
C:\Users\Greg Toncho>

```

Figure 2-28 Output of the Trace Route program

Courtesy of Course Technology/Cengage Learning

3. You can garner some information about the geography of the path your packet took by looking at the router's name. For example, in Figure 2-28, the domain name of the third router is `yc.edu`, which is a router at Yavapai College in Prescott, Arizona, where this book has been written. The fourth and fifth routers have the domain name `qwest.net`, and the router's name begins with "phn," which tells you that the router is on Qwest's network in Phoenix. You get the idea. However, looking up router names can sometimes make the trace run slowly. To do the same trace without looking up names, type `tracert -d www.yahoo.com` and press **Enter**. This time, you should see only the IP address of each router.
4. Try using Trace Route to determine the path packets take to other destinations. Try `books.tomsho.com`, and for a destination on the East Coast, try `www.course.com`. For a destination in Germany, try `www.kontron.de`. If the trace repeatedly times out (indicated by an asterisk, *, in the output), press **Ctrl+C** to stop the trace.
5. Close the command prompt window.
6. You can also find tools that show you the route on a map. Start your Web browser, and go to `www.yougetsignal.com/tools/visual-tracert`. In the Remote Address text box, type any of the destinations in Step 4 or any other address you like. This online tool attempts to map out the path your packets take to get to their destination.
7. Exit your Web browser and shut down your computer.

Chapter Summary

- Network repeaters and hubs take incoming bit signals and repeat them at their original strength out all connected ports. A hub is just a multiport repeater. Hubs are a central connecting device for multiple computers, but because hubs allow only one device to communicate at a time, the bandwidth of each port must be shared between all connected computers.
- Network switches interconnect multiple computers, just as hubs do. However, instead of simply regenerating incoming bit signals and repeating them to all other ports, a switch reads the destination MAC address in the frame to determine which port the destination device is connected to and forwards the frame to only that port.
- Switches use switching tables to determine which MAC address can be found on which port. Switches can operate in full-duplex mode, allowing connected devices to both transmit and receive data simultaneously. Hubs operate only in half-duplex mode.
- Access points are a central device in a wireless network and perform a similar function to hubs. An AP requires devices to use an RTS signal when they want to transmit data, and the AP responds with a CTS signal when it's okay to transmit. This extra network traffic reduces the effective bandwidth of wireless networks.
- Network interface cards create and mediate the connection between the computer and network medium. A computer's MAC address is defined on the NIC as a burned-in address. The NIC reads each frame arriving on the network medium and determines whether the frame's destination address matches its MAC address. If it matches or is a broadcast frame, the NIC processes the frame; otherwise, it's discarded.

- Wireless NICs perform the same function as wired NICs. Wireless NICs must be selected to match the wireless standard supported on the AP. When a wireless client connects to an AP, it uses the SSID to identify the wireless network's name.
- Routers connect LANs to one another and forward packets from one LAN to another, according to the destination IP address specified in the packet. Routers use routing tables to determine where to forward packets.
- Unlike hubs and switches, routers don't forward broadcast frames. Each interface on a router is the delimiter for a broadcast domain. When a router receives a unicast frame, it reads the destination IP address and compares it with the list of networks in its routing table. If a match is found, the router forwards the packet to the destination network or to another router that gets the packet to its destination. If no match is found, the router discards the frame. If a router has a default route defined, it forwards any packets that don't match networks in its routing table to the default route.



Key Terms

access point (AP) A wireless device that serves as the central connection point of a wireless LAN and mediates communication between wireless computers.

bandwidth sharing A network design in which interconnecting devices allow only one connected device to transmit data at a time, thus requiring devices to share available bandwidth.

broadcast domain The scope of devices to which broadcast frames are forwarded. Router interfaces delimit broadcast domains because they don't forward broadcasts, whereas switches and hubs do.

broadcast frame A network message intended to be processed by all devices on a LAN; has the destination address FF:FF:FF:FF:FF:FF.

clear to send (CTS) A signal an AP generates in response to a request-to-send signal. A CTS signal indicates that the computer that sent an RTS can transmit data. *See also* access point (AP) *and* request to send (RTS).

dedicated bandwidth A property of switches in which each port's bandwidth is dedicated to the devices connected to the port; on a hub, each port's bandwidth is shared between all devices connected to the hub.

default gateway The address configured in a computer's IP address settings specifying the address of a router to which the computer can send all packets destined for other networks.

default route A routing table entry that tells a router where to send a packet with a destination network address that can't be found in the routing table.

full-duplex mode A communication mode in which a device can simultaneously transmit and receive data on the same cable connection. Switches can operate in full-duplex mode, but hubs can't.

half-duplex mode A communication mode in which a device can send or receive data but can't do both simultaneously. Hubs operate only in half-duplex mode; switches can operate in both half-duplex and full-duplex modes.

hub A network device that performs the same function as a repeater but has several ports to connect a number of devices; sometimes called a multiport repeater. *See also* repeater.

network bandwidth The amount of data that can be transferred on a network during a specific interval; usually measured in bits per second.

network interface card (NIC) A device that creates and mediates the connection between a computer and the network medium.

promiscuous mode An operational mode of a NIC in which all frames are read and processed rather than only broadcast and unicast frames addressed to the NIC. Protocol analyzer software sets a NIC to promiscuous mode so that all network frames can be read and analyzed.

repeater A network device that takes incoming signals and regenerates, or repeats them to other parts of the network.

request to send (RTS) A signal used in wireless networks indicating that a computer has data ready to send on the network. *See also* access point *and* clear to send (CTS).

router A device that enables multiple LANs to communicate with one another by forwarding packets from one LAN to another. Routers also forward packets from one router to another when LANs are separated by multiple routers; they have multiple interfaces, and each interface communicates with a LAN.

service set identifier (SSID) The name assigned to a wireless network so that wireless clients can distinguish between them when more than one is detected.

switch A network device that reads the destination MAC addresses of incoming frames to determine which ports should forward the frames.

switching table A table containing MAC address and port pairs that a switch uses to determine which port to forward frames it receives.

unicast frame A network message addressed to only one computer on the LAN.

uplink port A designated port on a hub or switch used to connect to another hub or switch without using a crossover cable.


Review Questions

1. Which of the following is a limitation of early networks that used a daisy-chain method of connecting computers? (Choose all that apply.)
 - a. Total number of computers that could be connected
 - b. The processing speed of the computers connected
 - c. Cable length
 - d. No Internet access
2. Which of the following is true of a repeater?
 - a. Receives frames and forwards them
 - b. Determines which network to send a packet
 - c. Receives bit signals and strengthens them
 - d. Has a burned-in MAC address for each port

3. Which of the following is true of a hub? (Choose all that apply.)
 - a. Usually has just two ports
 - b. Transmits regenerated signals to all connected ports
 - c. Usually has four or more ports
 - d. Works with MAC addresses
4. Which of the following is the unit of measurement by which a hub's bandwidth is usually specified?
 - a. Bytes per second
 - b. Bits per second
 - c. Packets per second
 - d. Bytes per minute
5. Which of the following describes how devices connected to a hub use the speed at which the hub can transmit data?
 - a. Bandwidth optimization
 - b. Bandwidth dedication
 - c. Bandwidth sharing
 - d. Bandwidth multiplier
6. Which of the following is a likely indicator light on a hub? (Choose all that apply.)
 - a. CRC error
 - b. Link status
 - c. Connection speed
 - d. Activity
 - e. Signal strength
7. Which of the following describes how devices connected to a switch use the speed at which the switch can transmit data?
 - a. Dedicated bandwidth
 - b. Half-duplex bandwidth
 - c. Half-scale bandwidth
 - d. Shared bandwidth
8. What does a switch use to create its switching table?
 - a. Source IP addresses
 - b. Destination logical addresses
 - c. Destination physical addresses
 - d. Source MAC addresses



9. What purpose does the timestamp serve in a switching table?
 - a. Tells the switch when to forward a frame
 - b. Tells the switch how long to wait for a response
 - c. Tells the switch when to delete an entry
 - d. Tells the switch how long it has been running
10. What feature of a switch allows devices to effectively communicate at 200 Mbps on a 100 Mbps switch?
 - a. Uplink port
 - b. Full-duplex mode
 - c. Shared bandwidth
 - d. Bit strengthening
 - e. Frame doubling
 - f. Signal regeneration
11. To which device is a wireless access point most similar in how it operates?
 - a. Hub
 - b. Switch
 - c. NIC
 - d. Router
12. What's the purpose of an RTS signal in wireless networking?
 - a. It allows the AP to request which device is the transmitting station.
 - b. It allows the AP to tell all stations that it's ready to transmit data.
 - c. It allows a client to notify the AP that it's ready to send data.
 - d. It allows a client to request data from the AP.
13. Which of the following is a common operational speed of a wireless network?
 - a. 10 Kbps
 - b. 110 Gbps
 - c. 600 Kbps
 - d. 11 Mbps
14. Which of the following is a task performed by a NIC and its driver? (Choose all that apply.)
 - a. Provides a connection to the network medium
 - b. Converts bit signals into frames for transmission on the medium
 - c. Receives packets from the network protocol and creates frames
 - d. Adds a header before sending a frame to the network protocol
 - e. Adds error-checking data to the frame

- 
15. Which of the following best describes a MAC address?
 - a. A 24-bit number expressed as 12 decimal digits
 - b. Two 24-bit numbers, in which one is the OUI
 - c. A 48-bit number composed of 12 octal digits
 - d. A dotted decimal number burned into the NIC
 16. Under which circumstances does a NIC allow inbound communications to pass through the interface? (Choose all that apply.)
 - a. The source MAC address is the broadcast address.
 - b. The destination MAC address matches the built-in MAC address.
 - c. The destination MAC address is all binary 1s.
 - d. The NIC is operating in exclusive mode.
 17. How does a protocol analyzer capture all frames?
 - a. It configures the NIC to capture only unicast frames.
 - b. It sets all incoming destination addresses to be broadcasts.
 - c. It configures the NIC to operate in promiscuous mode.
 - d. It sets the exclusive mode option on the NIC.
 - e. It captures only multicast frames.
 18. In Windows 7, which of the following displays information about currently installed NICs?
 - a. Network Connections
 - b. NICs and Drivers
 - c. Local Area Networks
 - d. Computers and Devices
 19. Which of the following is the purpose of an SSID?
 - a. Assigns an address to a wireless NIC
 - b. Acts as a unique name for a local area connection
 - c. Acts as a security key for securing a network
 - d. Identifies a wireless network
 20. Which of the following describe the function of routers? (Choose all that apply.)
 - a. Forward frames from one network to another
 - b. Connect LANS
 - c. Attach computers to the internetwork
 - d. Work with packets and IP addresses
 21. What information is found in a routing table?
 - a. Computer names and IP addresses
 - b. Network addresses and interfaces

- c. MAC addresses and ports
 - d. IP addresses and MAC addresses
22. You currently have 15 switches with an average of 20 stations connected to each switch. The switches are connected to one another so that all 300 computers can communicate with each other in a single LAN. You have been detecting a high percentage of broadcast frames on this LAN. You think the number of broadcasts might be having an impact on network performance. What should you do?
- a. Connect the switches in groups of five, and connect each group of switches to a central hub.
 - b. Upgrade the switches to a faster speed.
 - c. Reorganize the network into smaller groups and connect each group to a router.
 - d. Disable broadcast forwarding on the switches.
23. Review the routing table in Figure 2-29. Based on this figure, where will the router send a packet with the source network number 1.0 and the destination network number 3.0?
- a. EthA
 - b. WAN A
 - c. WAN B
 - d. None of the above

Routing Table	
Network	Interface
1.0	EthA
2.0	WAN A
3.0	WAN B

Figure 2-29 Routing table

Courtesy of Course Technology/Cengage Learning

24. If a router receives a packet with a destination network address unknown to the router, what will the router do?
- a. Send the packet out all interfaces.
 - b. Discard the packet.
 - c. Add the destination network to its routing table.
 - d. Query the network for the destination network.
25. Which of the following is true about routers? (Choose all that apply.)
- a. Forward broadcasts
 - b. Use default routes for unknown network addresses
 - c. Forward unicasts
 - d. Used primarily to connect workstations

Challenge Labs



Challenge Lab 2-1: Determining Whether Your Computer Is Connected to a Hub or Switch

Time Required: 15 minutes

Objective: Use packet information captured with Wireshark to determine whether your computer is attached to the rest of the classroom with a hub or switch.

Required Tools/Equipment: Your classroom computer with Wireshark installed and the IP address of another classroom computer or device

Description: You saw the difference between hubs and switches in earlier projects. Specifically, you saw which packets Wireshark captured when your computer was connected to a hub versus a switch. In this challenge lab, set up a test, working with your classmates, to determine whether classroom computers are connected to a hub or switch. Write a short memo to your instructor with the following information:

- What filter options (if any) did you configure in Wireshark?
- What commands did you use to generate packets on the network?
- What IP addresses did you attempt to communicate with?
- What was your result? Is your computer attached to a hub or switch? Why did you come to this conclusion?



Challenge Lab 2-2: Capturing Trace Route Packets

Time Required: 15 minutes

Objective: Use Wireshark to capture Trace Route packets.

Required tools/equipment: Your classroom computer with Wireshark installed and Internet access

Description: In this challenge lab, you capture packets generated by the Trace Route program. You need to determine what types of packets are generated so that you know which types of packets to capture and inspect. Run Trace Route (use any Web sites you like) and capture the packets your computer generates and the router responses. After you have finished this lab, write a short memo discussing the following points:

- What type of packets does Trace Route use?
- What is the response each router sends back to your computer?
- How does your computer get a response from each router between your computer and the destination?



TIP

You can look up information on Trace Route at www.ehow.com/how-does_5164102_traceroute-work.html or do a Google search for it.



Case Projects



Case Project 2-1

You have been hired to upgrade a network of 50 computers currently connected to 10 Mbps hubs. This long-overdue upgrade is necessary because of poor network response time caused by a lot of collisions occurring during long file transfers between clients and servers. How do you recommend upgrading this network? What interconnecting devices will you use, and what benefit will you get from using these devices? Write a short memo describing the upgrade and, if possible, include a drawing of the new network.

Case Project 2-2

Two hundred workstations and four servers on a single LAN are connected by a number of switches. You're seeing an excessive number of broadcast packets throughout the LAN and want to decrease the effect this broadcast traffic has on your network. What steps must you take to achieve this goal?

Case Project 2-3

In Chapter 3, you learn about network topologies and technologies. As preparation, do

Internet research on the following topics:

- Physical versus logical topology
- Bus topology
- Star topology
- Ring topology
- Ethernet and CSMA/CD

Write a short explanation (two to three sentences) of each concept and be prepared to discuss it with the class.

Network Topologies and Technologies

After reading this chapter and completing the exercises, you will be able to:

- Describe the primary physical networking topologies in common use
- Describe the primary logical networking topologies in common use
- Describe major LAN networking technologies

Not so long ago, there was a real choice to be made between available network topologies and technologies when designing and building a new internetwork. Thankfully, this area of networking has gotten simpler rather than more complex, mainly because the choices have narrowed, with inferior or costly solutions becoming obsolete.

This chapter discusses network topologies, which describe both the physical arrangement of cabling or pathways between network devices and the logical manner in which data is transferred from device to device. Next, you learn about network technologies or architectures that describe the methods computers use to transmit data to the networking medium in an orderly fashion. As you'll see, the topology and technology are often tightly coupled, as certain technologies can be used only with certain topologies. The choices have been limited because only a few technologies and topologies remain as viable options. As is often the case, however, it helps to know where networking started to get an idea of where it might be heading. So even though some information covered in this chapter is obsolete or nearly so, your understanding of these older technologies will help you better understand current and future technologies.

Physical Topologies

The word “topology,” for most people, describes the lay of the land. A topographic map, for example, shows the hills and valleys in a region, whereas a street map shows only the roads. A network topology describes how a network is physically laid out and how signals travel from one device to another. However, because the physical layout of devices and cables doesn't necessarily describe how signals travel from one device to another, network topologies are categorized as physical and logical.

The arrangement of cabling and how cables connect one device to another in a network are considered the network's **physical topology**, and the path data travels between computers on a network is considered the network's **logical topology**. You can look at the physical topology as a topographic map that shows just the lay of the land along with towns, with only simple lines showing which towns have pathways to one another. The logical topology can be seen as a street map that shows how people actually have to travel from one place to another. As you'll see, a network can be wired with one physical topology but pass data from machine to machine by using a different logical topology.

All network designs today are based on these basic physical topologies: bus, star, ring, and point-to-point. A bus consists of a series of computers connected along a single cable segment. Computers connected via a central device, such as a hub or switch, are arranged in a star topology. Devices connected to form a loop create a ring. Two devices connected directly to one another make a point-to-point topology. Keep in mind that these topologies describe the physical arrangement of cables. How the data travels along these cables might represent a different logical topology. The dominant logical topologies in LANs include switching, bus, and ring, all of which are usually implemented as a physical star (discussed later in “Logical Topologies”).

Physical Bus Topology

The **physical bus topology**, shown in Figure 3-1, is by far the simplest and at one time was the most common method for connecting computers. It's a continuous length of cable




Figure 3-1 A physical bus topology network

Courtesy of Course Technology/Cengage Learning

connecting one computer to another in daisy-chain fashion. One of this topology's strengths is that you can add a new computer to the network simply by stringing a new length of cable from the last computer in the bus to the new machine. However, this strength is countered by a number of weaknesses:

- There's a limit of 30 computers per cable segment.
- The maximum total length of cabling is 185 meters.
- Both ends of the bus must be terminated.
- Any break in the bus brings down the entire network.
- Adding or removing a machine brings down the entire network temporarily.
- Technologies using this topology are limited to 10 Mbps half-duplex communication because they use coaxial cabling, discussed in Chapter 4.

Because of the preceding limitations, a physical bus topology is no longer a practical choice, and technology has moved past this obsolete method of connecting computers. However, the original Ethernet technology was based on this topology, and the basis of current LAN technology has its roots in the physical bus. So your understanding of bus communication aids your general understanding of how computers communicate with each other across a network.

How Data Travels in a Physical Bus Two properties inherent in a physical bus are signal propagation and signal bounce. In any network topology, computers communicate with each other by sending information across the media as a series of signals. When copper wire is the medium, as in a typical physical bus, these signals are sent as a series of electrical pulses that travel along the cable's length in all directions. The signals continue traveling along the cable and through any connecting devices until they weaken enough that they can't be detected or until they encounter a device that absorbs them. This traveling across the medium is called **signal propagation**. However, even if a signal encounters the end of a cable, it bounces back and travels in the other direction until it weakens or is otherwise impeded.

When a signal hits the end of a cable and bounces back up the cable's length, it interferes with signals following it, much like an echo. Imagine if you were trying to communicate

in an empty room with hard walls that caused your voice to echo continuously. The echo from the first words out of your mouth would garble the sound of words that followed, and your message would be unintelligible. The term used when electricity bounces off the end of a cable and back in the other direction is called **signal bounce** or **reflection**. To keep signal bounce from occurring, you do what you would to keep excessive echo from occurring; you install some type of material at both ends of the medium to absorb the signal. In a physical bus, you install a **terminator**, which is an electrical component called a resistor that absorbs the signal instead of allowing it to bounce back up the wire.

Physical Bus Limitations Now that you know more about how a physical bus works, the previous list of weaknesses needs some additional explanation. The limitation of 30 stations per cable segment means only 30 computers can be daisy-chained together before the signal becomes too weak to be passed along to another computer. As an electrical signal encounters each connected workstation, some of its strength is absorbed by both the cabling and the connectors until the signal is finally too weak for a computer's NIC to interpret. For the same reason, the total length of cabling is limited to 185 meters, whether there's 1 connected station or 30 connected stations. The network can be extended in cable length and number of workstations by adding a repeater to the network, which, as you know, regenerates the signal before sending it out.

At all times, both ends of the bus must be terminated. An unterminated bus results in signal bounce and data corruption. When a computer is added or removed from the network, both ends are no longer terminated, resulting in an interruption to network communication.

For a small network of only a few computers, you might think a bus topology is fine, until you consider the last weakness listed: maximum bandwidth of 10 Mbps half-duplex communication. A physical bus uses coaxial cable (a cabling type discussed in Chapter 4, similar to what's used in cable TV connections), which is limited to a top speed of 10 Mbps and communication in only half-duplex mode. Most of today's networks use twisted-pair cabling, which can operate at 100 Mbps or faster and run in full-duplex mode, so communication between devices is much faster.

For all these reasons, the physical bus topology has long since fallen out of favor and been replaced largely by the star topology, discussed next.

Physical Star Topology

The **physical star topology** uses a central device, such as a hub or switch, to interconnect computers in a LAN (see Figure 3-2). Each computer has a single length of cable going from its NIC to the central device.

Some advantages of a physical star topology are the following:

- Much faster technologies are used than in a bus topology.
- Centralized monitoring and management of network traffic is possible.
- Network upgrades are easier.

A physical star is the topology of choice for these reasons and more. With a central device, communication options are available that simply aren't possible with a physical bus. For example, the central device can be a 100 Mbps hub, which increases a physical bus's top




Figure 3-2 A physical star topology network

Courtesy of Course Technology/Cengage Learning

speed tenfold, or a switch, making it possible for multiple communication sessions to occur simultaneously and in full-duplex mode.

As a budding network administrator, being able to monitor and manage your network with a central device is a big advantage over what was possible with a physical bus topology. Today's hubs and switches can include software that collects statistics about your network traffic patterns and even alerts you when excessive errors or unusually high traffic rates are occurring on your network. You don't get these features in a \$19.99 hub or switch, but enterprise-level devices can be equipped with several network management tools.

As long as your current cabling and installed NICs support it, your network can be upgraded quickly and easily from a ponderous 10 Mbps hub-based LAN to a blazing fast 100 Mbps or even 1000 Mbps switched network simply by replacing the central device. In addition, if your NICs must also be upgraded, you can upgrade in steps because most devices support multiple speeds. So if you want to upgrade from 100 Mbps to 1000 Mbps, you can replace the central device with a switch that supports both speeds, and then upgrade NICs as time and money allow. The switch transmits and receives on each port at the speed supported by the NIC connected to that port.

What happens if the number of workstations you need to connect exceed the number of ports on the central device? In this case, you can connect hubs or switches, as you learned in Chapter 2. When several hubs or switches must be connected, usually one device is used as the central connecting point, forming an extended star.

Extended Star The extended star topology, shown in Figure 3-3, is the most widely used in networks containing more than just a few computers. As the name implies, this topology is a star of stars. A central device, usually a switch, sits in the middle. Instead of attached computers forming the star's arms, other switches (or hubs) are connected to the central switch's ports. Computers and peripherals are then attached to these switches or






Figure 3-3 An extended star topology network

Courtesy of Course Technology/Cengage Learning

hubs, forming additional stars. The extended star is sometimes referred to as a “hierarchical star” because there are two or more layers of stars, all connecting back to the central star.

The extended star can be used to connect many computers, with the central device running at a very fast speed to shuttle data between the LAN’s outer stars. This topology is most effective when the center of the star is running at a much faster speed than other devices; for example, the central device can run at 1000 Mbps while other devices run at 100 Mbps.

How Data Travels in a Physical Star The details of how data travels from computer to computer in a physical star depend on the type of central device. Data transmission starts at a device at the end of one of the central device’s arms. From there, it travels along the network medium’s length until it arrives at the central device. As you know from learning how hubs and switches work, the transmission path differs, depending on the device. Other devices, such as multistation access units (MAUs) used in token ring networks, move data differently. The type of central device, therefore, determines the logical topology, discussed later in this chapter.

Physical Star Disadvantages With all the clear advantages of a physical star, you might wonder whether there are any disadvantages. None outweigh the advantages, but it’s worth mentioning that the central device represents a single point of failure. In other words, if the hub or switch fails or someone kicks the power cord out of the outlet, down goes the entire network. Thankfully, these devices tend to be reliable and are usually placed out of the way of everyday foot traffic. That being said, they do fail from time to time, and having a spare on hand is a good idea.

When a physical bus was still the norm and the physical star was just coming on the networking scene in the late 1980s, it was often argued that because each computer must be

cabled directly to the central device, instead of a bus's daisy-chain arrangement, more cable was required to connect computers. This point is indeed true, and at the time, the amount of cabling needed was a factor in designing a network with a bus or star arrangement. By the time the star network's advantages were fully realized in the mid-1990s, however, the cabling cost difference had diminished substantially, and the advantages clearly outweighed the minor cost disadvantage.

Physical Ring Topology

A **physical ring topology** is like a bus, in that devices are daisy-chained one to another, but instead of terminating each end, the cabling is brought around from the last device back to the first device to form a ring. This topology had little to no following in LANs as a way to connect computers. It was used, however, to connect LANs with a technology called Fiber Distributed Data Interface (FDDI). FDDI was most often used as a reliable and fast **network backbone**, which is cabling used to communicate between LANs or between hubs or switches. In Figure 3-4, the devices used to connect buildings form a ring, but computers on each LAN are connected with a physical star topology.




Figure 3-4 A physical ring topology is usually used to connect LANs

Courtesy of Course Technology/Cengage Learning

The physical ring also had reliability issues because data had to be forwarded from one station to the next. Unlike a bus, in which data travels in all directions and is terminated at both ends, a ring doesn't have any beginning or end. So each station must reproduce data and pass it along to the next station until it reaches the destination or the originator of the data. In other words, data always travels in one direction. If any station in the ring fails, data can no longer be passed along, and the ring is broken.

Technologies such as FDDI overcome some problems with a physical ring network by creating a dual ring, in which data can travel in both directions so that a single device failure doesn't break the entire ring. However, this technology is costly, and although it was used extensively in the 1990s and early 2000s because it was fast (100 Mbps) and reliable, 100 Mbps and 1000 Mbps Ethernet have largely supplanted it with an extended star technology.



Point-to-Point Topology

As its name implies, a **point-to-point topology** is a direct link between two devices. It's most often used in WANs, in which a device on a business's network has a dedicated link to a telecommunication provider, such as the local phone company. The connection then hooks into the phone company's network to provide Internet access or a WAN or MAN link to a branch office. The advantage of this type of topology is that data travels on a dedicated link, and its bandwidth isn't shared with other networks. The disadvantage is that this topology tends to be quite expensive, particularly when used as a WAN link to a distant branch office.

Point-to-point topologies are also used with wireless networks in what's called a **wireless bridge**. This setup can be used to connect two buildings without using a wired network (see Figure 3-5) or to extend an existing wireless network.



Figure 3-5 A point-to-point wireless topology

Courtesy of Course Technology/Cengage Learning

A rudimentary LAN can also be set up with a point-to-point topology by connecting a cable between the NICs on two computers. Of course, this method allows only two computers on the network, but it can be used effectively for transferring files from one computer to another in the absence of a hub or switch.

So as you can see, point-to-point topologies are used for specialized purposes. They aren't commonly used in LANs; they're used more often in WANs and large internetworks.

Mesh Topology A **mesh topology** connects each device to every other device in a network. You can look at a mesh topology as multiple point-to-point connections for the purposes of redundancy and fault tolerance. Figure 3-6 shows a full mesh topology between four locations, with the switch in each location providing connectivity to multiple computers. Each switch is connected to every other switch, which is called a “full mesh.” If each switch were connected to only two other switches, it would be called a “partial mesh.” In either case, the purpose of creating a mesh topology is to ensure that if one or more connections fail, there's another path for reaching all devices on the network. For example, in Figure 3-6, two connections could fail, but all devices could still communicate with one another. This type of topology is used mostly commonly in large internetworks and WANs, where routers or switches in multiple buildings or towns are connected in a partial or full mesh. Parts of the Internet are also designed with a partial mesh topology, in which major ISPs are connected so that even if one ISP's network fails, data can bypass this part of the network to get to its destination.

Mesh topologies, although reliable, are also expensive because of the additional cabling and ports required. In most cases, the ports used to connect devices are the highest speed available, such as 1 Gbps or 10 Gbps, and they often use expensive fiber-optic cabling for connecting buildings.

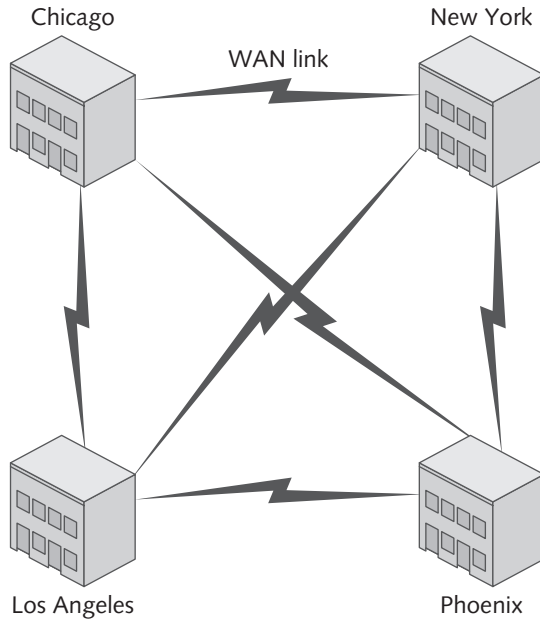


Figure 3-6 Switches in each building are connected in a full mesh topology

Courtesy of Course Technology/Cengage Learning

Logical Topologies

As mentioned, a network’s logical topology describes how data travels from computer to computer. In some cases, as with a physical bus and physical ring, the logical topology mimics the physical arrangement of cables. In other cases, as with a physical star, the electronics in the central device determine the logical topology.

A network’s logical topology reflects the underlying network technology (covered later in “Network Technologies”) used to transfer frames from one device to one another. Table 3-1 summarizes the main logical topologies, the technologies using them, and the physical topologies for implementing them.

Table 3-1 Logical topologies and associated network technologies and physical topologies

Logical topology	Network technology	Physical topology	Description
Bus	Ethernet	Bus or star	A logical bus topology can be implemented as a physical bus (although this topology is now obsolete). When a logical bus is implemented as a physical star using wired Ethernet, the center of the star is an Ethernet hub. Whatever the physical topology is, data transmitted from a computer is received by all other computers.
	Wireless LANs	Star	Wireless LANs use a physical star topology because they connect through a central access point. However, only one device can

Table 3-1 Logical topologies and associated network technologies and physical topologies (continued)

Logical topology	Network technology	Physical topology	Description
			transmit at a time and all devices hear the transmission, so a wireless LAN can be considered a logical bus topology.
Ring	Token ring	Star	Token ring networks use a central device called a multistation access unit (MAU or MSAU). Its electronics form a logical ring, so data is passed from computer to computer in order, until it reaches the destination device.
	FDDI	Ring	As discussed, FDDI devices are connected in a physical ring, and data passes from device to device until it reaches the destination.
Switched	Ethernet	Star	A switched logical topology using a physical star topology running Ethernet is by far the most common topology/technology combination now and likely will be well into the future. A switched topology creates dynamic connections or circuits between two devices whenever data is sent. This topology is sometimes considered a switched point-to-point topology because a circuit is established between two points as needed to transfer data (like turning on a switch), and then the circuit is broken when it's no longer needed (like turning off a switch).

You have seen what a logical bus looks like when implemented as a physical bus. All computers are daisy-chained to one another, and network signals travel along the cable's length in all directions, much like water flowing through interconnected pipes. When a logical bus is implemented as a physical star, the same process occurs, but the pathways are hidden inside the central hub. Figure 3-7 shows what a logical bus might look like when implemented with a hub.




Figure 3-7 A logical bus implemented as a physical star

Courtesy of Course Technology/Cengage Learning



A logical bus is sometimes called a “shared media topology” because all stations must share the bandwidth the media provides.

A logical ring using a physical star implements the ring inside the central device’s electronics, which is an MAU in the token ring technology. Data is passed from one node or computer to another until it reaches the destination device (see Figure 3-8). When a port has no device connected to it, it’s simply bypassed, and data is sent out the next connected port.




Figure 3-8 A logical ring implemented as a physical star

Courtesy of Course Technology/Cengage Learning

A switched topology works something like what’s shown in Figure 3-9. Although there’s always an electrical connection between the computer and switch, when no data is being transferred, there’s no logical connection or circuit between devices. However, when the switch receives a frame, a logical circuit is made between the source and destination devices until the frame is transferred.




Figure 3-9 The logical functioning of a switch

Courtesy of Course Technology/Cengage Learning

To better understand how these logical topologies work, it helps to know the network technology that drives each topology (discussed later in “Network Technologies”).



Hands-On Project 3-1: Building a Physical Star Topology Network

Time Required: 20 minutes

Objective: Build a physical star topology network.

Required Tools/Equipment: Three workstations named Computer1, Computer2, and Computer3; a hub; and three patch cables. Workstations should be configured with an IP address or automatic IP address assignment. Each station should have Wireshark installed.

Description: In this project, you build a small physical star topology; this task can be done in groups of three or more or as an instructor demonstration. After each station is connected to the hub, you ping another station to verify connectivity. Next, you use Wireshark to capture ping packets so that you can determine the network’s logical topology.

1. Power on the hub.
2. Connect each workstation to the hub with the supplied cables.
3. Inspect the hub and the workstation NIC to verify that you have a good connection with the hub. Write down how you determined whether the connection with the hub is good:

4. On each workstation, open a command prompt window, and then type **ipconfig** and press **Enter** to determine your IP address. Write down the IP address of each computer:
 - IP address of Computer1:

 - IP address of Computer2:

 - IP address of Computer3:

5. Ping each computer to verify that you can communicate with it. If the pings aren’t successful, check that the IP addresses you wrote down are correct and the connection with the hub is good, and then try again.
6. Make sure you coordinate the rest of the project, starting with this step, with students at the other computers. Start Wireshark, and start a capture session by clicking the interface name listed in the Interface List section.
7. At the command prompt, ping the next computer. For example, if you’re at Computer1, ping Computer2; if you’re at Computer2, ping Computer3; and if you’re at Computer3, ping Computer1.

8. Based on which packets Wireshark captured, what's your logical topology?
-
9. Exit Wireshark, close all open windows, and leave the computers running if you're continuing to the next project.

Network Technologies

A network technology, as the phrase is used here, can best be described as the method a network interface uses to access the medium and send data frames and the structure of these frames. Other terms include network interface layer technologies, network architectures, and Data Link layer technologies. What it comes down to is whether your network uses Ethernet, 802.11 wireless, token ring, or some combination of these and other technologies to move data from device to device in your network. Most LANs are now based on a combination of Ethernet and 802.11 wireless. WANs use technologies specifically designed to carry data over longer distances, such as frame relay, FDDI, Asynchronous Transfer Mode (ATM), and others.

The network technology sometimes, but not always, defines frame format and which media types can be used to transfer frames. For example, different Ethernet speeds specify a minimum grade of copper or fiber-optic cabling that must be used as well as the connectors attached to the ends of cables. FDDI requires fiber-optic cabling, but other technologies, such as frame relay, can run on a variety of media types.

This book focuses on LAN technologies with particular emphasis on Ethernet and 802.11 wireless because they're the most commonly used. Some WAN technologies are also described briefly in this chapter and in more detail in Chapter 12.

Network Technologies and Media

Because some of the network technologies discussed in this chapter specify the types of media they require to operate, the following sections summarize the most common media types. However, you can find more details on network media in Chapter 4.

Unshielded Twisted Pair Unshielded twisted pair (UTP) is the most common media type in LANs. It consists of four pairs of copper wire, with each pair tightly twisted together and contained in a plastic sheath or jacket (Figure 3-10).




Figure 3-10 UTP cabling

Courtesy of Course Technology/Cengage Learning



UTP comes in numbered categories, up to Category 7 as of this writing. The higher the category, the higher the cable's bandwidth potential. Category 5 Enhanced (Cat 5E) and Category 6 (Cat 6) are the most common in wired LANs, allowing speeds up to 10 Gbps. UTP cabling is used in physical star networks, and the maximum cable length from NIC to hub or switch is 100 meters in LAN applications. UTP cabling is susceptible to electrical interference, which can cause data corruption, so it shouldn't be used in electrically noisy environments.

Fiber-Optic Cabling Fiber-optic cabling uses extremely thin strands of glass to carry pulses of light long distances and at high data rates. It's usually used in large internetworks to connect switches and routers and sometimes to connect high-speed servers to the network. Because of its capability to carry data over long distances (several hundred to several thousand meters), it's also used in WAN applications frequently. Fiber-optic cabling isn't susceptible to electrical interference, so unlike UTP, it can be used in electrically noisy environments. It requires two strands of fiber to make a network connection: one for transmitting and one for receiving.

Coaxial Cable Best known for its use in cable TV, coaxial cable is obsolete as a LAN medium, but it's used as the network medium for Internet access via cable modem. Coaxial cable was the original media used by Ethernet in physical bus topologies, but its limitation of 10 Mbps half-duplex communication made it obsolete for LAN applications after star topologies and 100 Mbps Ethernet became the dominant standard. Coaxial cable in LANs can have lengths of around 200 meters.

Baseband and Broadband Signaling Network technologies can use media to transmit signals in two main ways: baseband and broadband. The **baseband** transmission method sends digital signals in which each bit of data is represented by a pulse of electricity (on copper media) or light (on fiber-optic media). These signals are sent at a single fixed frequency, using the medium's entire bandwidth. In other words, when a frame is sent to the medium, it occupies the cable's entire bandwidth, and no other frames can be sent along with it—much like having cable TV that carries only a single channel. LAN technologies, such as Ethernet and token ring, use baseband transmission. If cable TV used baseband signaling, you would need one cable for each channel!

Thankfully, cable TV and cable modem Internet access use broadband transmission. Instead of digital pulses, **broadband** systems use analog techniques to encode binary 1s and 0s across a continuous range of values. Broadband signals move across the medium in the form of continuous electromagnetic or optical waves rather than discrete pulses. On broadband systems, signals flow at a particular frequency, and each frequency represents a channel of data. That's why broadband systems, such as cable TV and Internet, can carry dozens or hundreds of TV channels plus Internet access on a single cable wire: Each channel operates at a different frequency. In addition, incoming and outgoing Internet data use separate channels operating at different frequencies from TV channels.

Ethernet Networks

Ethernet, the most popular LAN technology, has many advantages, including ease of installation, scalability, media support, and low cost. It supports a broad range of transmission speeds, from 10 Mbps to 10 Gbps.

As discussed, Ethernet can operate in a bus or star physical topology and a bus or switched logical topology. It has been in use since the mid-1970s but didn't mature as a technology until the early to mid-1980s. Ethernet being around for almost 40 years is a testament to the

original designers, whose forethought enabled Ethernet to scale from a 3 Mbps technology in its early years to a 10 Gbps and beyond technology today.

Although there are many variations of Ethernet, all forms are similar in their basic operation and frame formatting. What differs in the variations are the cabling, speed of transmission, and method by which bits are encoded on the medium. Because the frame formatting is the same, however, Ethernet variations are compatible with one another. That's why you often see NICs and Ethernet hubs and switches described as 10/100 or 10/100/1000 devices. These devices can support multiple Ethernet speeds because the underlying technology remains the same, regardless of speed.

Ethernet Addressing Every Ethernet station must have a physical or MAC address. As you learned in Chapter 2, a MAC address is an integral part of network interface electronics and consists of 48 bits expressed as 12 hexadecimal digits. When a frame is sent to the network medium, it must contain both source and destination MAC addresses. When a network interface detects a frame on the media, the NIC reads the frame's destination address and compares it with its own MAC address. If they match or if the destination address is the broadcast MAC address (all binary 1s or FF:FF:FF:FF:FF:FF in hexadecimal), the NIC reads the frame and sends it to the network protocol for further processing.

Ethernet Frames A frame is the unit of network information NICs and switches work with. It's the NIC's responsibility to transmit and receive frames and a switch's responsibility to forward frames out the correct switch port to get the frame to its destination.

Ethernet frames come in four different formats, or **frame types**, depending on the network protocol used to send frames, and unfortunately, these frame types are incompatible with one another. They were developed during Ethernet's early days, before standards were solidified. If your network needed to support multiple protocols, such as TCP/IP, IPX/SPX, and AppleTalk, you had to make sure your computers were configured to support all these frame types. Thankfully, TCP/IP has become the dominant network protocol in LANs, so supporting multiple frame types is largely unnecessary, except for networks that still run older Novell NetWare servers. Given this reality, this section examines only the frame type used by TCP/IP: Ethernet II. The other frame types are Ethernet SNAP, Ethernet 802.3, and Ethernet 802.2. For information on these frame types, see Appendix B.



The four Ethernet frame types are incompatible in the same Ethernet standard (such as using both Ethernet II and Ethernet SNAP in 100 Mbps), but each frame type is compatible with the same frame type in different standards. For example, Ethernet II in 10 Mbps Ethernet is compatible with Ethernet II in 100 Mbps and 1000 Mbps Ethernet.

Regardless of frame type, Ethernet networks can accommodate frames between 64 bytes and 1518 bytes. Shorter or longer frames are considered errors. Each frame is composed of the following (see Figure 3-11):

- A 14-byte frame header composed of these three fields:
 - A 6-byte Destination MAC Address field
 - A 6-byte Source MAC Address field
 - A 2-byte Type field



- A Data field from 46 to 1500 bytes
- A frame trailer (frame check sequence [FCS]) of 4 bytes

Destination MAC Address (6 bytes)	Source MAC Address (6 bytes)	Type (2 bytes)	Data (46–1500 bytes)	FCS (4 bytes)
Frame header			Data (frame payload)	Frame trailer

Figure 3-11 Ethernet II frame format

Courtesy of Course Technology/Cengage Learning

You’ve already learned the purpose and format of destination and source MAC addresses. The Type field in the frame header indicates the network protocol in the data portion. For example, this field might indicate that the Data field contains an IP, IPv6, or ARP packet, to name just a few possibilities. The data portion, often referred to as the “frame payload,” contains network protocol header information as well as the actual data an application is transferring. The FCS in the frame trailer is an error-checking code (discussed later in “Ethernet Error Handling”).



There are exceptions to the 1518-byte maximum frame size. For example, a function of some switches requires an additional 4-byte field in the Ethernet frame, bringing the maximum size to 1522 bytes. In addition, Jumbo frames of up to 9000 bytes are supported by some NICs and switches but aren’t officially supported in the current Ethernet standards. To use Jumbo frames, the feature must be enabled on every device on the LAN and be implemented the same way by these devices.

Ethernet Media Access Before a NIC can transmit data to the network medium, it must adhere to some rules governing how and when the medium can be accessed for transmission. The rules ensure that data is transmitted and received in an orderly fashion and all stations have an opportunity to communicate. The set of rules for each networking technology is referred to as its **media access method** or **media access control**. Note that the acronym for “media access control” is MAC, which is where the term “MAC address” comes from.

The media access method Ethernet uses in half-duplex mode is **Carrier Sense Multiple Access with Collision Detection (CSMA/CD)**. To understand this method better, break this term down into parts. “Carrier sense” means to listen. The rules for half-duplex Ethernet state that a device can send or receive data but can’t do both simultaneously. So before a device can send, it must listen to see whether the medium is already busy, much like a group of people having a conversation. Each person listens for a pause in the conversation before speaking up. “Multiple access” simply means that multiple computers can be listening and waiting to transmit at the same time, which brings you to “collision detection.” A **collision** occurs if two or more devices on the same medium transmit simultaneously. For example, if two people are waiting to chime in on a group conversation, they both hear a lull in the conversation at the same time and speak up simultaneously, causing a “collision” in the conversation. Ethernet’s collision detection method is much like a person’s; Ethernet detects, or “hears,” the other station transmit, so it knows a collision has occurred. The NIC then waits for a random period before attempting to transmit again. Ethernet repeats the “listen before transmitting” process until it transmits the frame without a collision. Simulation 7 on the book’s CD shows a simulation of the CSMA/CD process.



Simulation 7: Ethernet operation using CSMA/CD

As you determined in Hand-On Project 2-4, when you attempted to create enough traffic to generate a collision, the CSMA/CD access method is efficient. It takes quite a bit of traffic to generate collisions, especially on a 100 Mbps network. However, the more devices on a logical bus topology and the more data they transmit, the greater the chance of a collision. So although CSMA/CD works well, today’s multimedia-heavy networks have somewhat outgrown it, and Ethernet has adapted to this development.



CSMA/CD is considered a contention-based access method, which means computers are allowed to send whenever they have data ready to send. Obviously, CSMA/CD modifies this rule somewhat by stipulating that the computer must listen first to ensure that no other station is in the process of transmitting.

Collisions and Collision Domains Remember that collisions can occur only in an Ethernet shared-media environment, which means a logical bus topology is in use. In this environment, all devices interconnected by one or more hubs hear all signals generated by all other devices. The signals are propagated from hub to hub until there are no more devices or until a device is encountered that doesn’t use a logical bus topology, such as a switch or a router. The extent to which signals in an Ethernet bus topology network are propagated is called a **collision domain**. Figure 3-12 shows a network diagram with two




Figure 3-12 A network diagram showing two collision domains delimited by a switch

Courtesy of Course Technology/Cengage Learning

collision domains enclosed in circles. All devices in a collision domain are subject to the possibility that whenever a device sends a frame, a collision might occur with another device sending a frame at the same time. This fact has serious implications for the number of computers that can reasonably be installed in a single collision domain. The more computers, the more likely it is that collisions occur. The more collisions, the slower network performance is.

Notice in Figure 3-12 that all computers connected to Hubs 1 to 3 are in the same collision domain, and computers connected to Hubs 4 to 6 are in a different collision domain. This is because a switch port delimits the collision domain, which means collisions occurring in one collision domain don't propagate through the switch.

Although collisions in an Ethernet network are usually associated with hubs, technically it's possible for a collision to occur with a computer connected to a switch. A collision with a switch can occur only if the NIC connected to the switch port is operating in half-duplex mode. In addition, the collision domain is limited to only the devices connected to a single switch port. The same is true of routers. However, given that an Ethernet frame of maximum size is transmitted on a 10 Mbps switch in just over a millisecond and just over a microsecond on a 100 Mbps switch, the likelihood of a collision with a switch is low.



If a hub is connected to a switch port in an extended star topology, collisions can occur between devices connected to the hub and the switch port. To avoid collisions altogether, use only switches in your network design with computers that have NICs operating in full-duplex mode.

Ethernet Error Handling One reason for Ethernet's low cost and scalability is its simplicity. It's considered a best-effort delivery system, meaning that when a frame is sent, there's no acknowledgement or verification that the frame arrived at its intended destination. Ethernet relies on network protocols, such as TCP/IP, to ensure reliable delivery of data. It's similar to the package delivery guy at a corporation. His job is to take what he's given to its intended destination; it's the package receiver's job to verify its contents and let the sender know it was received.

Ethernet can also detect whether a frame has been damaged in transit. The error-checking code in an Ethernet frame's trailer is called a **Cyclic Redundancy Check (CRC)**, which is the result of a mathematical algorithm computed on the frame data. The CRC is calculated and placed in the frame trailer before the frame is transmitted. When the frame is received, the calculation is repeated. If the results of this calculation don't match the CRC in the frame, it indicates that the data was altered in some way, usually from electrical interference. If a frame is detected as damaged, because Ethernet is a best-effort delivery system, it simply discards the frame but doesn't inform the sending station that an error occurred. Again, it's the network protocol's job to ensure that all expected data was actually received. The network protocol or, in some cases, the application sending the data is responsible for resending damaged or missing data, not Ethernet.



A collision is the exception to Ethernet's lack of action when an error occurs. When frames are involved in a collision, Ethernet resends them automatically because all stations detect that a collision has occurred.

Half-Duplex Versus Full-Duplex Communication As discussed in Chapter 2, half-duplex communication means a station can transmit and receive data but not at the same time, much like a two-way radio. When Ethernet is implemented as a logical bus topology (using hubs), NICs can operate only in half-duplex mode and must use the CSMA/CD access method.

However, a network switch allows half-duplex or full-duplex communication. If a NIC is operating in half-duplex mode while connected to a switch, it must use CSMA/CD. However, the only time a collision can occur in this circumstance is if the switch happens to transmit a frame to the NIC at the same time the NIC is attempting to transmit.

Full-duplex mode, by definition, means a NIC can transmit and receive simultaneously. Therefore, when an Ethernet NIC is operating in full-duplex mode connected to a switch, CSMA/CD isn't used because a collision can't occur in full-duplex mode. Because full-duplex mode eliminates the delays caused by CSMA/CD and allows double the network bandwidth, most Ethernet LANs now operate in this mode using switches.

Ethernet Standards

Ethernet can operate at different speeds over different types of media, and each variation is associated with an IEEE standard. The following sections discuss many of these standards, some of which are obsolete or had limited use.

Standards Terminology Ethernet standards are generally expressed in one of two ways. One way is using the IEEE document number defining the standard. For example, IEEE 802.3 is the parent document specification for 10 Mbps Ethernet using thick coaxial cable, which was ratified in 1983. All other variations and speeds of Ethernet are subdocuments of the original 802.3 specification.

The second way of expressing an Ethernet standard is to use the XBaseY terminology. Most IEEE 802.3 documents describe the transmission speed, type of transmission, and length or type of cabling and are designated with terms such as 100BaseT. In 100BaseT, for example, the "100" designates the speed of transmission (100 Mbps), the "Base" indicates a base-band signaling method, and the "T" specifies twisted-pair cabling. All the BaseT Ethernet standards use a physical star topology. The following sections discuss the major standards and their designations.

10BaseT Ethernet 10BaseT Ethernet, defined by IEEE 802.3i, has been the mainstay of Ethernet networks since the early 1990s. It runs over Category 3 or higher UTP cabling and uses two of the four wire pairs. Because of its slower transmission speed, 10BaseT networks using a logical bus topology (with hubs) are more susceptible to collisions than faster 100BaseT networks are. In addition, the amount of data sent and received by a typical user makes 10BaseT seem slow in typical media-heavy environments compared with the more common 100BaseT and 1000BaseT standards.

If you work for an organization still using hubs, you need to know that there are limits to how many hubs you can string together to connect all computers. The rule for expanding a 10BaseT network with hubs is that no more than four hubs can be placed between two communicating workstations. This rule ensures that all stations on the network can detect a collision. Because of the limited time for signals to propagate through a network, if more



than four hubs exist between end stations, a collision on one end of the network might not be detected by stations on the other side of the network in time for them to react properly. If switches rather than hubs are used, there's no such limitation because a collision on a switch can take place only between the switch and a single workstation.

A business network still using 10BaseT should upgrade to 100 or 1000BaseT to take full advantage of current technology. A home or small-office network that uses the network mainly for sharing Internet access and transferring documents can still use 10BaseT effectively if its Internet connection is considerably slower than 10 Mbps. However, 10BaseT is essentially an obsolete technology, and networks using 10BaseT should upgrade as soon as circumstances permit.

100BaseTX Ethernet 100BaseTX (often referred to as simply “100BaseT”), defined by IEEE 802.3u, is the most commonly used Ethernet variety today. It runs over Category 5 or higher UTP cable and uses two of the four wire pairs: one to transmit data and the other to receive data. There are other varieties of 100BaseT Ethernet (discussed later in this section), but 100BaseTX is the standard that's usually in mind when discussing 100 Mbps Ethernet. It's also sometimes called Fast Ethernet.

An important consideration when designing a 100BaseTX network with hubs is the total number of hubs allowed between end stations. There are two types of 100BaseTX hubs: class I and class II. Class I hubs can have only one hub between communicating devices; class II hubs can have a maximum of two hubs between devices. This limitation is designed to ensure that when a collision occurs on a hub-based network, all stations in the collision domain have enough time to hear the collision and respond appropriately. If a 100BaseTX network uses mainly hubs to connect computers, a switch is often used in the center of an extended star to interconnect multiple hubs, as shown in Figure 3-13, to avoid this limitation. If you're using only switches in your network, this limitation doesn't apply. You'll probably see hubs of any type only in older network installations; all new designs and upgrades use switches of at least 100 Mbps.




Figure 3-13 Using a switch to interconnect 100BaseTX hubs

Courtesy of Course Technology/Cengage Learning

100BaseFX Ethernet In environments that aren't conducive to using copper wiring to carry network data (such as electrically noisy settings) or where the cable run length exceeds the reach of twisted-pair wiring, the only real choice in a wired network is fiber optics. 100BaseFX (with the F indicating “fiber optic”), which uses two strands of fiber-optic cable, is often the best choice of network technology in these settings. Fiber-optic cable installation is

still far more expensive than twisted-pair cable, but its advantages of being impervious to electrical noise and supporting longer cable segment lengths are worth the cost if the network requires these properties. 100BaseFX is rarely used as a complete replacement for 100BaseTX; instead, it's typically used as backbone cabling between hubs or switches and to connect wiring closets between floors or buildings. It's also used to connect client or server computers to the network when immunity to noise and eavesdropping is required.

1000BaseT Ethernet 1000BaseT Ethernet, released as the IEEE 802.3ab standard, supports 1000 Mbps Ethernet (usually called **Gigabit Ethernet**) over Category 5 or higher UTP cable. The 1 Gbps data rate results from sending and receiving data simultaneously (in full-duplex mode) at 250 Mbps in both directions over each of the four wire pairs in Category 5 cable. Therefore, each wire pair can send and receive data at the same time at 250 Mbps, which results in a bandwidth of 1000 Mbps (or 1 Gbps) in each direction in full-duplex mode. To support full-duplex transmission over a single pair of wires, 1000BaseT uses equipment called hybrids and cancellers, which combine multiple signals and cancel interference. So if the link operates in half-duplex mode, the channel speed is 1000 Mbps (250 Mbps times four wire pairs). When operating in full-duplex mode, 1000BaseT actually delivers 2 Gbps total bandwidth. In most cases, it runs in full-duplex mode connected to switches.

Unlike 10BaseT and 100BaseT Ethernet, 1000BaseT Ethernet doesn't dedicate a wire pair to transmitting or receiving. Each wire pair is capable of transmitting and receiving data simultaneously, thereby making the 1000 Mbps data rate possible in both half-duplex and full-duplex modes. Similarly to 100BaseT, 1000BaseT allows only one hub or repeater between end stations when using half-duplex communication. Most installations use switches that detect the speed of the connected device automatically, whether it's 10 Mbps, 100 Mbps, or 1000 Mbps. In addition, you'll be hard-pressed to find a 1000BaseT hub/repeater to purchase, making the one-repeater limitation an unlikely problem in new network designs.

1000BaseT Ethernet has gained wide acceptance in corporate data centers to connect servers to central switches and connect power users' desktop computers. Because 1000BaseT runs over standard Category 5 cable, the upgrade path for companies currently running 100BaseT is fairly simple. NICs and switches that don't operate at 1000 Mbps must be replaced, but the cabling infrastructure doesn't need to be. Although Category 5 cable is the minimum requirement, most new installations planning a 1000BaseT network should opt for Cat 5e or Cat 6 cable for their better high-speed transmission properties.



TIP

A power user is a network user who often uses the most recent technologies and runs high-end software and hardware that require more network resources than what the average user runs.

10GBaseT Ethernet The 2006 IEEE 802.3an standard defines 10 Gigabit Ethernet as running over four pairs of Category 6A UTP cabling. Unlike the other BaseT Ethernet standards, 10GBaseT operates only in full-duplex mode, so you won't find any 10 Gbps hubs—only switches. 10GBaseT isn't likely to find its way into many desktop computers in the near future; a search for 10GBaseT NICs shows their cost at more than \$1000 as of this writing. However, as more desktop systems begin operating at 1 Gbps, you might need to equip your network's servers with 10 Gigabit Ethernet NICs so that they can keep up with desktop systems.



Additional Ethernet Standards

Although the standards discussed previously constitute the majority of Ethernet LANs, quite a few other standards exist; some are common, and others are uncommon or obsolete. The following sections describe these other standards and their use in current networks briefly.

100BaseT4 As the name implies, 100BaseT4 Ethernet uses all four pairs of wires bundled in a UTP cable. The one advantage that 100BaseT4 has over 100BaseTX is the capability to run over Category 3 cable. When 100 Mbps speeds became available, many companies wanted to take advantage of the higher bandwidth. However, if the cable plant consisted of only Category 3 cable, there were just two choices: Replace the cabling with higher-grade Category 5 cabling so that 100BaseTX could be used, or use 100BaseT4 Ethernet. One of the biggest expenses of building a network is cable installation, so many organizations chose to get higher speeds with the existing cable plant by using 100BaseT4. Although these differences from 100BaseTX might seem like a good idea, 100BaseT4 never caught on and is essentially obsolete.

1000BaseLX 1000BaseLX uses fiber-optic media; the “L” stands for “long wavelength,” the kind of laser used to send signals across the medium. These lasers operate at wavelengths between 1270 to 1355 nanometers and work with single-mode fiber (SMF) and multimode fiber (MMF). Long-wavelength lasers cost more than short-wavelength lasers but can transmit their signals over longer lengths of cable.

Although the 1000BaseLX standard specifies a maximum cable segment length of 5000 meters, some manufacturers have extended it by using specialized and proprietary optical transceivers. Cisco Systems, for example, offers 1000BaseLH (“LH” stands for “long haul”), which provides a maximum cable segment length of 10,000 meters over SMF cable. For extremely long-distance Gigabit Ethernet communication, 1000BaseZX, another Cisco product, is capable of distances up to 100,000 meters over SMF cable.

1000BaseSX 1000BaseSX uses fiber-optic media; the “S” stands for “short wavelength.” These lasers operate at wavelengths between 770 to 860 nanometers and work only with MMF cable. Short-wavelength lasers can’t cover as much distance as long-wavelength lasers, but they are less expensive (and use cheaper MMF cable).

1000BaseCX 1000BaseCX uses specially shielded, balanced, copper jumper cables; the “C” stands for “copper,” the kind of electrical signaling used. Jumper cables are normally used for interconnections between devices or to link virtual LANs (VLANs) on a switch; these jumper cables might also be called “twinax” (short for “twin-axial”) or “short-haul” copper cables. Segment lengths for 1000BaseCX cables top out at 25 meters, which means they’re used mostly in wiring closets or equipment racks.

10 Gigabit Ethernet: IEEE 802.3ae Standards The 802.3ae standard governing several varieties of 10 Gigabit Ethernet before 10GBaseT was adopted in June 2002. This Ethernet version is much like the others in frame formats and media access method. However, it does have some important technical differences. It’s defined to run only on fiber-optic cabling, but the 10 Gigabit Ethernet standard specifies a maximum distance of 40 kilometers, compared with just 5 kilometers for the 1000BaseLX Gigabit Ethernet. This distance has important implications for WANs and MANs because although most WAN

and MAN technologies can be measured in megabits, 10 Gigabit Ethernet provides bandwidth that can transform how WAN speeds are considered. Like 10GBaseT Ethernet, 802.3ae 10 Gigabit Ethernet technologies run in full-duplex mode only, so the CSMA/CD access method isn't necessary.

The primary use of 10 Gigabit Ethernet technologies is as the network backbone, interconnecting servers and network segments running 100 Mbps and 1000 Mbps Ethernet technologies. However, they also have their place in storage area networks (SANs) and, along with 10GBaseT, can be used as the interface for enterprise-level servers.

As this technology matured, a number of implementations were developed that are divided into two basic groups: 10GBaseR for LAN applications and 10GBaseW for WAN applications. The W group of standards uses SONET framing over OC-192 links. (SONET and OC standards are explained in Chapter 12). Both groups have (S)hort range, (L)ong range, and (E)xtended range versions. The short-range versions use MMF fiber-optic cabling, and the long-range and extended-range versions run over SMF fiber-optic cabling. (These fiber-optic types are discussed in Chapter 4.) The following list summarizes the 802.3ae technologies:

- *10GBaseSR*—Runs over short lengths (between 26 and 82 meters) on MMF cabling. Applications are likely to include connections to high-speed servers, interconnecting switches, and SANs.
- *10GBaseLR*—Runs up to 10 km on SMF cabling and is used for campus backbones and MANs.
- *10GBaseER*—Runs up to 40 km on SMF cabling; used primarily for MANs.
- *10GBaseSW*—Uses MMF cabling for distances up to 300 meters; used for SONET campus network applications.
- *10GBaseLW*—Uses SMF cabling for distances up to 10 km; used for SONET WAN applications.
- *10GBaseEW*—Uses SMF cabling for distances up to 40 km; used for SONET WAN applications.

40 Gigabit and 100 Gigabit Ethernet: The 802.3ba Standard IEEE 802.3ba was ratified on June 17, 2010, and vendors are already shipping or announcing products. Some goals for 40 Gigabit Ethernet include the capability to transmit up to 10 kilometers over SMF fiber-optic cabling and at least 7 meters over a special copper wire assembly; for 100 Gigabit Ethernet, goals include transmitting up to 40 km over SMF cabling and 7 meters over copper. The existing frame format is preserved, meaning these new high-speed versions are backward-compatible with slower standards.



TIP

Although the 802.3ba task force has completed its work, you can read about how this new standard came to be at www.ieee802.org/3/ba/index.html.

As you can see, Ethernet has come a long way since Xerox transmitted at 3 Mbps over coaxial cable, and the journey from 3 Mbps to 10 Gbps isn't over yet. Table 3-2 summarizes many features and properties of the Ethernet standards discussed in this section.



Table 3-2 Ethernet standards and properties

Ethernet standard	IEEE document #	Transmission speed	Cable type	Minimum cable grade	Maximum distance	Design notes
10BaseT	802.3i	10 Mbps	UTP	Cat 3	100 meters	Maximum four hubs between stations
100BaseT/TX	802.3u	100 Mbps	UTP	Cat 5	100 meters	Maximum two hubs between stations
100BaseFX	802.3u	100 Mbps	MMF or SMF	N/A	2 km over MMF, 10 km over SMF	
1000BaseT	802.3ab	1000 Mbps	UTP	Cat 5 (Cat 5e or 6 preferred)	100 meters	Maximum one hub between stations
10GBaseT	802.3an	10 Gbps	UTP	Cat 6A	100 meters	Full-duplex only; no hubs
100BaseT4	802.3u	100 Mbps	UTP	Cat 3	100 meters	Obsolete; saw little use
1000BaseLX	802.3z	1000 Mbps	MMF or SMF	N/A	550 meters over MMF, 5 km over SMF	
1000BaseSX	802.3z	1000 Mbps	MMF	N/A	550 meters	
1000BaseCX	802.3z	1000 Mbps	Twinax	N/A	25 meters	Succeeded by 1000BaseT
10GBaseSR 10GBaseLR 10GBaseER 10GBaseSW 10GBaseLW 10GBaseEW	802.3ae	10 Gbps	MMF or SMF	N/A	Varies from 82 meters up to 40 km	Choice of technology depends on application
40 Gigabit Ethernet and 100 Gigabit Ethernet	802.3ba	40 and 100 Gbps	MMF, SMF, and copper assembly	N/A	40 km over SMF, 7 meters over copper	Standard ratified June 17, 2010

What's Next for Ethernet? Estimations are that Ethernet speeds will continue to increase, with Terabit Ethernet (1000 Gbps) available by 2015. This kind of mind-boggling speed will allow networks to transfer data across a city faster than some CPUs can transfer data to memory. When Internet providers begin using this level of bandwidth to connect to the Internet backbone and when homes and businesses can tap into it, too, extraordinary amounts of information will be at your fingertips. This speed has major implications for the entertainment industry and many other fields. The Ethernet train is revving up, and it promises to be an exhilarating ride.



Hands-On Project 3-2: Determining and Changing Your Ethernet Standard

Time Required: 15 minutes

Objective: Determine your Ethernet standard and change your connection speed to use a different standard.

Required Tools/Equipment: Classroom computers connected to a classroom hub or switch. The hub or switch and NICs must be capable of connecting at multiple speeds. For example, if you're using a 10/100 Mbps switch, and your NICs are capable of 10/100 Mbps, you change the connection speed to the slower rate. If possible, each student should be assigned a separate partner. Separate lab computers and a hub or switch can also be used for this project.

Description: In this project, you view your network connection properties to see at what speed your NIC is operating. Then you send a large ping message and note how long the reply takes. Next, you change the speed, if your NIC driver allows, and perform the same ping to see whether you can detect a time difference.



This project works best with physical computers rather than virtual machines. Even if you change the connection speed on the virtual machine, it transmits bits at the host computer's connection speed.

1. Log on to your computer as **NetAdmin**.
 2. Open a command prompt window, and then type **ipconfig** and press **Enter**. Exchange your IP address with your partner and write down your partner's IP address. Leave the command prompt window open for later.
-
3. Click the network connection icon in the taskbar and click **Open Network and Sharing Center**.
 4. In the Network and Sharing Center, click **Local Area Connection** to open the Local Area Connection Status dialog box (see Figure 3-14).
 5. In the Connection section, find the connection speed. Write down this information and, based on the speed listed, the Ethernet variety your computer is running:
 - Connection speed: _____
 - Ethernet variety: _____
 6. At the command prompt, ping your partner by typing **ping -l 60000 IPaddress** and pressing **Enter**. The **-l 60000** option in the command specifies that the ping message should be 60000 bytes instead of the typical length of 32 bytes. Note the time = values in the ping replies and write them down. For example, yours might say "time = 2ms," meaning the reply took 2 milliseconds. Not all times might be the same. Sometimes the first time is slower than the rest. Try pinging a few times to get an idea of the average time.
 - Ping reply times: _____
 7. Click the **Properties** button, and in the Local Area Connection Properties dialog box, click **Configure**.





Figure 3-14 The Local Area Connection Status dialog box

Courtesy of Course Technology/Cengage Learning

8. Click the **Advanced** tab. In the Property list box, click **Link Speed & Duplex** (or a similar name). Figure 3-15 shows the connection options. Not all NICs have the same options, so you might see different options.
9. The default setting is usually Auto Negotiation. Click **10 Mbps Half Duplex** if this option is available, and then click **OK**. If you were able to set this option, what speed and variety of Ethernet is your computer running now?
 - Connection speed: _____
 - Ethernet variety: _____

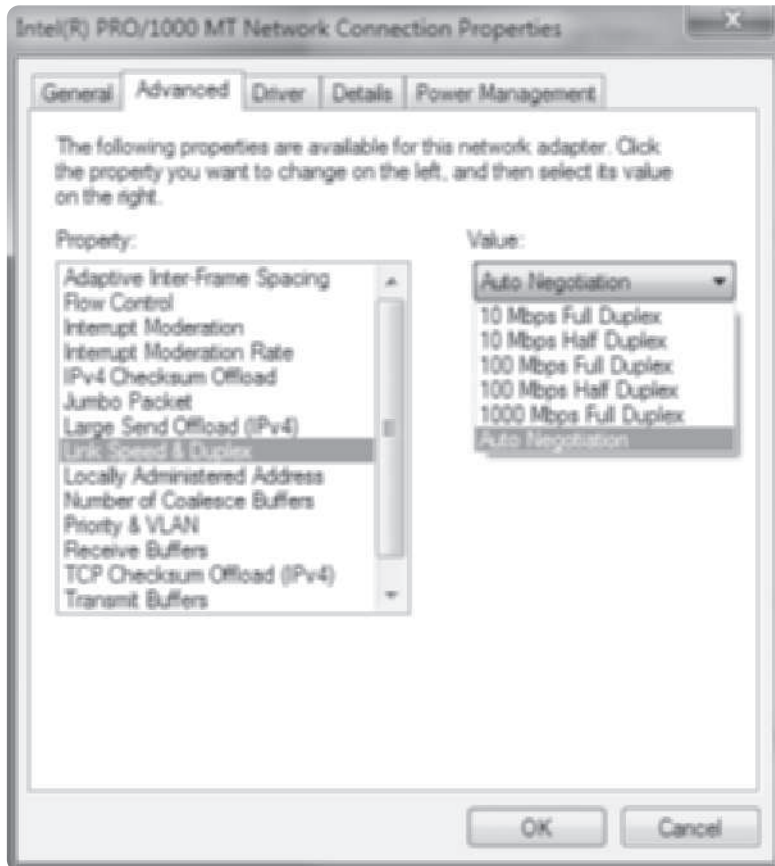
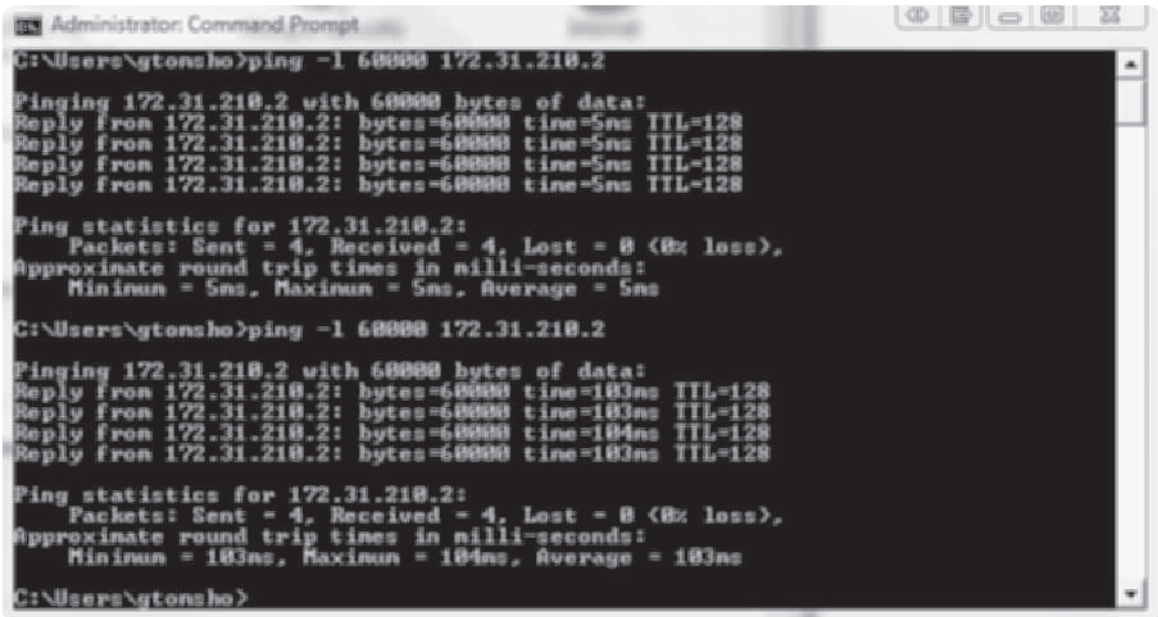


Figure 3-15 Settings for the Link Speed & Duplex property

Courtesy of Course Technology/Cengage Learning

10. After you and your partner have changed the connection speed to a lower value, repeat the ping command you used in Step 6. Write down the reply times, and state whether they were different:
 - Ping reply times at 10 Mbps: _____
11. Figure 3-16 shows two sets of ping results. The first result was from two computers connected at 1 Gbps (1000 Mbps) in full-duplex mode. The average reply took 5 ms. The second result was with the same computers connected at 10 Mbps half-duplex, and the average reply took 103 ms. Change your connection speed and duplex mode back to Auto Negotiation, and then close all open windows. Leave your computer running for the next project.



```

Administrator: Command Prompt
C:\Users\gtoncho>ping -l 60000 172.31.210.2

Pinging 172.31.210.2 with 60000 bytes of data:
Reply from 172.31.210.2: bytes=60000 time=5ms TTL=128
Reply from 172.31.210.2: bytes=60000 time=5ms TTL=128
Reply from 172.31.210.2: bytes=60000 time=5ms TTL=128
Reply from 172.31.210.2: bytes=60000 time=5ms TTL=128

Ping statistics for 172.31.210.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 5ms, Maximum = 5ms, Average = 5ms

C:\Users\gtoncho>ping -l 60000 172.31.210.2

Pinging 172.31.210.2 with 60000 bytes of data:
Reply from 172.31.210.2: bytes=60000 time=183ms TTL=128
Reply from 172.31.210.2: bytes=60000 time=183ms TTL=128
Reply from 172.31.210.2: bytes=60000 time=184ms TTL=128
Reply from 172.31.210.2: bytes=60000 time=183ms TTL=128

Ping statistics for 172.31.210.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 183ms, Maximum = 184ms, Average = 183ms

C:\Users\gtoncho>

```

Figure 3-16 Ping results at different connection speeds

Courtesy of Course Technology/Cengage Learning



Hands-On Project 3-3: Viewing an Ethernet Frame

Time Required: 20 minutes

Objective: Capture packets and examine details of the Ethernet II frame format.

Required Tools/Equipment: Classroom computers connected to a classroom hub or switch with Wireshark installed

Description: In this project, you capture some packets and then examine the frame and protocol headers.

1. If necessary, log on to your computer as NetAdmin.
2. Start Wireshark and click **Capture Options**. In the Capture Filter text box, type `icmp`, and then click **Start**.
3. Open a command prompt window, and then type `ping IPaddress` and press **Enter** (replacing `IPaddress` with the IP address of another student's computer or another device on your network).
4. In Wireshark, click the **Stop the running live capture** toolbar icon to stop the capture.
5. Click a packet summary in the top pane with ICMP listed in the protocol field.
6. In the middle pane, click to expand the **Ethernet II** row. It should look similar to Figure 3-17.




Figure 3-17 An Ethernet II frame in Wireshark

Courtesy of Course Technology/Cengage Learning

7. Notice the three fields in the Ethernet II frame: Destination, Source, and Type. The Destination and Source fields are the destination and source MAC addresses in the frame. In Figure 3-17, you see Cisco before the destination address and Vmware before the source address because Wireshark attempts to resolve the NIC manufacturer coded in the MAC address's first six digits. The full MAC address (without manufacturer name) is shown in parentheses. The Type field has the value 0x800, which indicates that the protocol in the frame is IP. Click to expand the **Internet Protocol** row.
8. Under Internet Protocol, you see details of the IP header, including the destination and source IP addresses. Click to expand the **Internet Control Message Protocol** row to view details of the ICMP protocol header. (You learn more about IP-related protocols in Chapter 5.)
9. Click to expand the **Data** portion of the frame, and then click the **Data** field to see the ICMP message data in hexadecimal in the bottom pane. The right side of this pane shows the translation from hexadecimal to ASCII (human-readable characters); as you can see, it's just portions of the alphabet repeated. Some Ping programs include more clever data, such as "Hello, are you there?" The actual data in a ping message doesn't matter; what matters is that the reply contains the same data as the ping request.
10. Exit Wireshark and click **Quit without Saving** when prompted. Close the command prompt window.
11. Stay logged on if you're going on to the next project; otherwise, shut down your computer.

802.11 Wi-Fi

The 1997 802.11 wireless networking standard, also referred to as **Wireless Fidelity (Wi-Fi)**, has continued to undergo development. With it, manufacturers of wireless networking devices have brought inexpensive, reliable wireless LANs to homes and businesses. The current standards include 802.11b, 802.11g, and 802.11n running at a 2.4 GHz frequency with speeds of 11 Mbps, 54 Mbps, and up to 600 Mbps, respectively. 802.11a isn't as prevalent as the other 802.11 standards; it specifies a bandwidth of 54 Mbps at a 5 GHz frequency.

Of these competing standards, 802.11b and 802.11g are the most widespread at this writing because they have been in use the longest. 802.11g is backward-compatible with 802.11b and, therefore, offers a convenient bandwidth upgrade path. 802.11n is the newest Wi-Fi standard, completed in October 2009. It can operate in the 2.4 GHz or 5 GHz frequency range and is backward-compatible with 802.11b and 802.11g. It's capable of speeds up to 600 Mbps by using the multiple-in, multiple-out (MIMO) technique. This technique uses multiple antennas and divides available frequency ranges into multiple channels to achieve

high speeds. However, 600 Mbps is the *maximum* bandwidth possible under 802.11n; actual speeds of common devices are typically 100 to 300 Mbps.

Essentially, 802.11 wireless is an extension to Ethernet, using airwaves as the medium. In fact, most 802.11 networks incorporate some wired Ethernet segments. The 802.11 networks can extend from several feet to several hundred feet, depending on environmental factors, such as obstructions and radio frequency interference. The prevalence of people owning 802.11-enabled laptops, iPads, and cell phones has spawned a new mode for accessing the Internet. Many businesses have set up Wi-Fi hot spots, which are localized wireless access areas. You can sit outside your favorite coffee shop, for example, and use a wireless Internet connection with your portable devices. College campuses, too, are using hot spots so that students can sit in a courtyard between classes and access the campus network and the Internet with their Wi-Fi-enabled laptops.



TIP

For more information on 802.11 standards, see www.wi-fiplanet.com.

Wi-Fi Modes Wi-Fi networks can operate in one of two modes: infrastructure and ad hoc. Most Wi-Fi networks operate in **infrastructure mode**, meaning wireless stations connect through a wireless AP before they can begin communicating with other devices. **Ad hoc mode**, sometimes called peer-to-peer mode, is a wireless mode of operation typically used only in small or temporary installations. There's no central device, and data travels from one device to another in a line (more or less). If you want to describe ad hoc mode in terms of a physical and logical topology, it most resembles a physical and logical bus. Most of this chapter's discussion of Wi-Fi focuses on infrastructure mode.



NOTE

Ad hoc mode should not be used in public environments because it's less secure than infrastructure mode.

Wi-Fi Communication Channels Wi-Fi networks operate at one of two radio frequencies: 2.4 GHz and 5.0 GHz. However, this frequency is not fixed. The 2.4 GHz Wi-Fi variety, which includes 802.11b, g, and n, operates from 2.412 GHz through 2.484 GHz, divided into 14 separate 5 MHz channels (although only the first 11 channels are used in North America, and other regions have channel use restrictions). The 5.0 GHz Wi-Fi variety divides frequencies between 4.915 GHz and 5.825 GHz into 42 channels of 10, 20, or 40 MHz each, depending on the region of the world where it's used. The remainder of the discussion about Wi-Fi channels pertains to 2.4 GHz Wi-Fi because it's the most popular, but most points also apply to the 5.0 GHz varieties.

A wireless channel works somewhat like a TV channel, in which each channel works at a different frequency and can, therefore, carry different streams of data. When you configure a wireless AP, you can choose the channel in which it operates (see Figure 3-18). By choosing a channel that's not in heavy use, you can improve reception and throughput rate. However, 2.4 GHz Wi-Fi actually requires 25 MHz to operate correctly, effectively spanning five channels. So if you're configuring several Wi-Fi networks, you should choose

channels that are five apart; for example, if you configure three Wi-Fi networks in close proximity, choose channels 1, 6, and 11.



Figure 3-18 Selecting a Wi-Fi channel on an access point

Courtesy of Course Technology/Cengage Learning

Several tools are available that scan channels to see how much activity is on each. You can then configure your AP to operate on a less frequently used channel. Figure 3-19 is an example of the output of the program inSSIDer and shows that a number of Wi-Fi networks were detected. Each one is labeled with its SSID.

Wi-Fi runs in the vast range of frequencies encompassed by microwave radio. For this reason, a microwave oven can cause interference in a Wi-Fi network. The result of this interference can vary from a slight loss in signal strength to disconnection from the network while the oven is running. A change in Wi-Fi channels can sometimes lessen the effects of microwave oven interference. In addition, some cordless phones use the same frequencies as Wi-Fi networks. If a cordless phone is causing interference, try changing the channel of the AP, the cordless phone (if possible), or both.

Wi-Fi Security Because the network signals and, therefore, network data, of a Wi-Fi network aren't constrained by physical media, access to a Wi-Fi network *must* be secure. The signals from a Wi-Fi network can travel several hundred feet, which means Wi-Fi devices outside your home or business can detect them. A person with a Wi-Fi-enabled device sitting

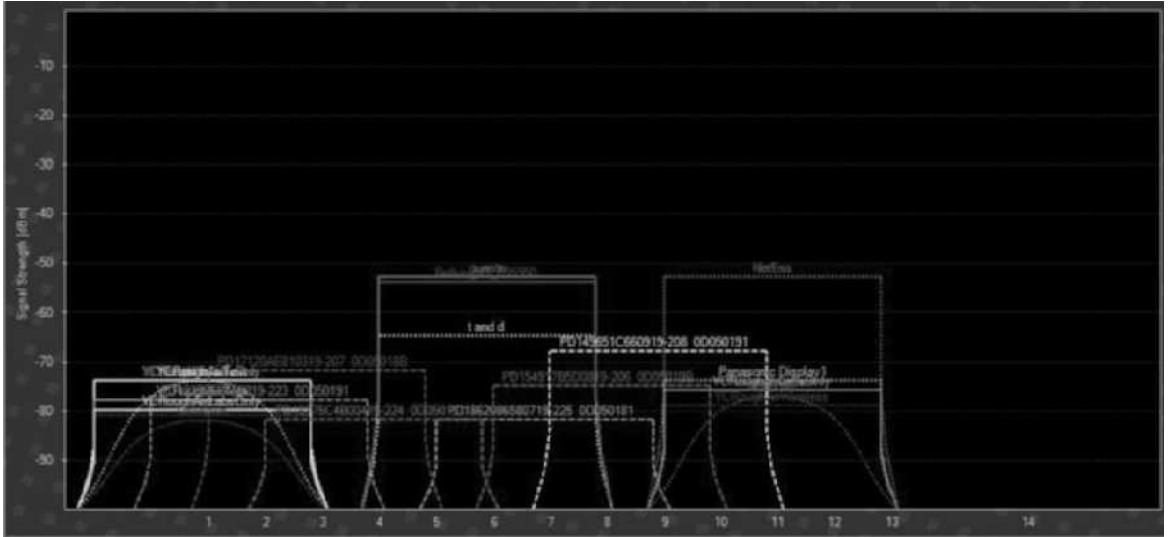


Figure 3-19 Wi-Fi network channel activity

Courtesy of Course Technology/Cengage Learning

outside your home or business can connect to an unsecured network and use your Internet access to capture packets with a program such as Wireshark—or worse, access files on your computers.

At the least, a Wi-Fi network should be protected by an encryption protocol that makes data unauthorized users capture extremely difficult to interpret. Wi-Fi devices typically support one of these encryption protocols, listed in order of effectiveness: Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), and Wi-Fi Protected Access 2 (WPA2). Not all devices support all three protocols; in particular, older devices might support only WEP and/or WPA. Wi-Fi encryption is configured on the AP, so to connect to the network, Wi-Fi devices connecting to the AP must be configured for the particular encryption protocol. Wi-Fi security is discussed in more depth in Chapters 7 and 11.



WEP should be used only when it's the only option available because its encryption protocols are easily broken.

Wi-Fi Access Method and Operation You have learned about CSMA/CD as the access method in wired forms of Ethernet, but wireless networks have a special problem with this access method. CSMA/CD requires that all stations be able to hear each other so that each station knows when another station is sending data. This requirement is reasonable, but if two stations try to send at the same time, a collision can occur. Fortunately, in a wired network, sending stations hear the collision and attempt to resend the data. If you've ever used a push-to-talk handheld radio, you know that when you're talking, you can't hear anybody else talking, and vice versa. 802.11 networks work the same way. If a station transmits data, it can't hear whether any other station is transmitting, so if a collision does occur, the sending station doesn't detect it. For this reason, 802.11 specifies the **Carrier Sense Multiple Access with**

Collision Avoidance (CSMA/CA) access method, in which an acknowledgement is required for every packet sent, as explained in Chapter 2. With this requirement, if a collision occurs, the sending station knows the packet didn't arrive safely because there's no acknowledgement. Simulation 8 on the book's CD shows a simulation of basic wireless LAN operation.



Simulation 8: Basic wireless LAN operation

Another problem exists in wireless networks that doesn't happen in wired networks. It's quite possible that in a three-station wireless network, all workstations can communicate with the AP: For example, workstation A can hear workstation B and workstation B can hear workstation C, but workstation A can't hear workstation C, perhaps because the two are out of range. This situation is called the "hidden node problem." CSMA/CA doesn't work because workstation A never knows whether workstation C is sending, and vice versa. To counteract this problem, the 802.11 standards specify another feature that uses handshaking before transmission. As described in Chapter 2, a station must send the AP a request-to-send (RTS) packet requesting transmission. If it's okay to transmit, the AP sends a clear-to-send (CTS) message, and the workstation starts its communication. All other devices communicating with the AP hear the exchange of RTS and CTS messages, thus informing them that another device has control of the medium.

The 802.11b standard specifies a transmission rate of 11 Mbps, but this value isn't absolute. Environmental conditions can prevent transmission at this speed. Therefore, transmission speeds might be dropped incrementally from 11 Mbps to 5.5 Mbps to 2 Mbps and, finally, to 1 Mbps to make a reliable connection. In addition, there's no fixed segment length for wireless networks because reliable communication relies heavily on the environment—for example, the number of walls between stations and the AP. The other three 802.11 standards behave similarly.

In general, an 802.11 network operating at 2.4 GHz has a maximum distance of 300 feet at full speed with no obstructions. However, this distance can be longer with 802.11n and large, high-quality antennas. Keep in mind that the data rate might suffer as the distance and number of obstructions increase.



For an excellent tutorial on wireless networking, visit www.networkcomputing.com/1115/1115ws22.html.



Hands-On Project 3-4: Configuring an Ad Hoc Wi-Fi Network

Time Required: 20 minutes

Objective: Configure an ad hoc Wi-Fi network.

Required Tools/Equipment: Two or more computers with 802.11 wireless NICs installed. USB wireless NICs work well because they don't require opening the computer case. Laptops with

built-in wireless NICs will also do. The 802.11 standard supported doesn't matter as long as the NICs are compatible with one another. These instructions are written for Windows 7, but some steps can be changed to accommodate Vista. Windows XP requires substantial changes to the steps.

Description: All students go through the steps of creating an ad hoc Wi-Fi network, but only one person in each group can actually create the network. You can work in groups of two to four or more, but each group must use a different ad hoc network name. For example, if each group is assigned a number, the network names can be AdHoc1, AdHoc2, and so forth. After the ad hoc network is created by the first person in each group who does the steps, the other members of the group cancel the creation of the ad hoc network and then connect to the ad hoc network the first group member created.

1. Start your computer and log on as **NetAdmin**. If the wireless NIC isn't installed yet, install it according to your instructor's instructions.
2. After the wireless NIC is installed, click the network connection icon in the taskbar and click **Open Network and Sharing Center**.
3. In the Network and Sharing Center, click **Set up a new connection or network**, and then click **Next** in the welcome window.
4. In the Choose a connection option window, click **Set up a wireless ad hoc (computer-to-computer) network**, and then click **Next**. Click **Next** again.
5. In the Network name text box, type **AdHocX** (replacing **X** with your group number). Click the **Security type** list arrow and click **No authentication (Open)**, as shown in Figure 3-20. Click **Next**.

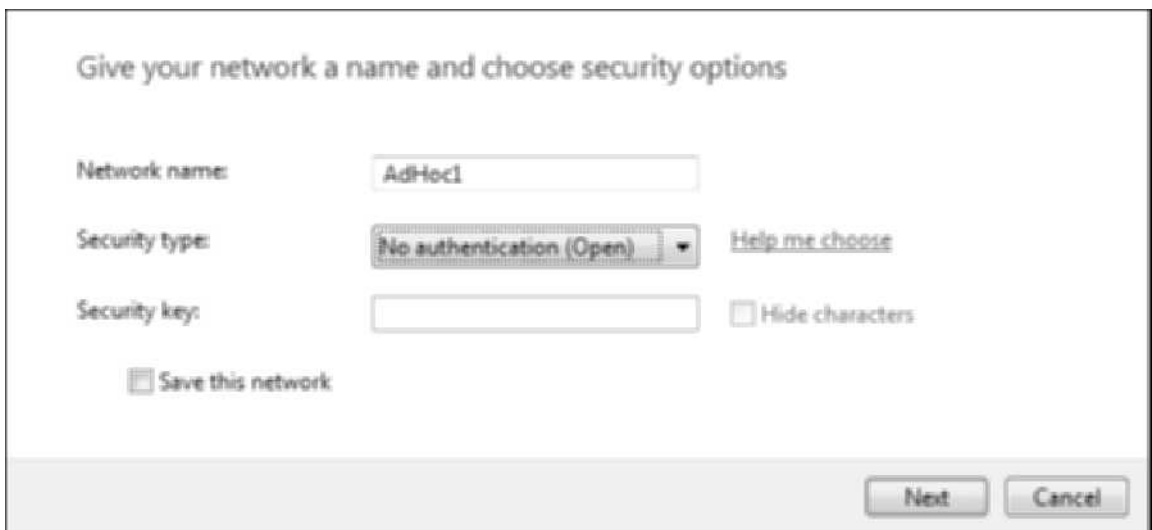


Figure 3-20 Creating an ad hoc wireless network

Courtesy of Course Technology/Cengage Learning

6. Only one person can create an ad hoc network with a particular name. After the network is created, anyone else can connect to it. If you're the first person to create the network, you see a window like Figure 3-21. If the network was already created, you see a window like Figure 3-22. If the network already exists, click **Cancel**; otherwise, click **Close**.

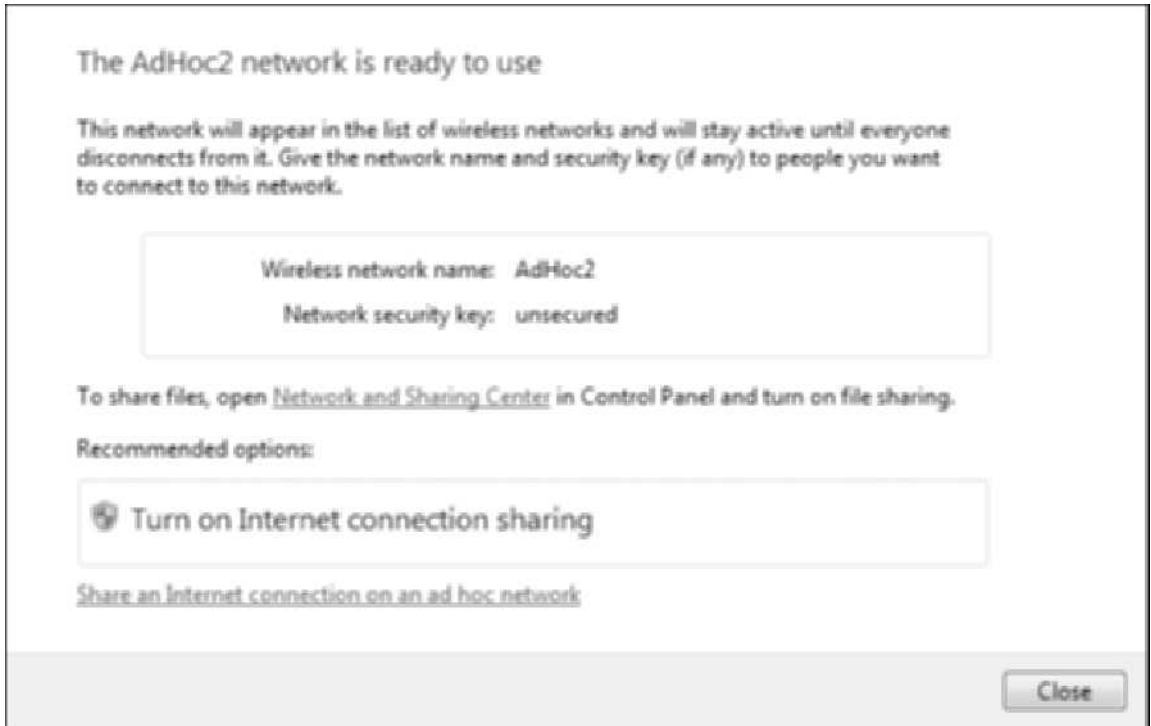


Figure 3-21 Successful creation of an ad hoc network

Courtesy of Course Technology/Cengage Learning

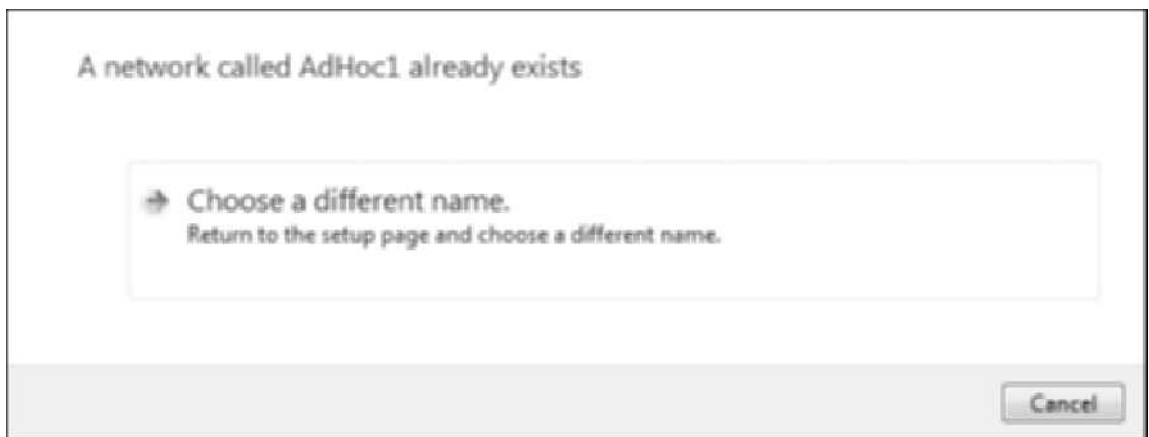


Figure 3-22 The ad hoc network already exists

Courtesy of Course Technology/Cengage Learning

7. In the Network and Sharing Center, click **Connect to a network**. If you're the one who created the ad hoc network, you are already connected. If not, click the **AdHocX** network you're assigned to, and then click **Connect**.

8. In the Network and Sharing Center, click **Change adapter settings**. Double-click **Wireless Network Connection** to view the connection's status, which includes the connection speed and the SSID (AdHocX). Click **Details**.
9. In the details window, note your IP address. Share this address with one or more people in your group, and then ping one of your group members to verify that you can communicate. You won't be able to communicate with members of other AdHocX networks. If the ping fails, it's probably because Windows 7 identified the network as a public network, and Windows Firewall blocked the ping packets. You can turn off the Public profile in the firewall (ask your instructor how to do this) if you want the ping to succeed. Be sure to enable the firewall when you're finished.
10. Close all open windows. Click the wireless network connection icon in the taskbar. Click the AdHocX network to which you're connected and click **Disconnect**. Shortly after all computers are disconnected from an ad hoc network, it's no longer displayed as an available wireless network.
11. If your wireless NIC is a USB device, you can remove it now. Stay logged on if you're going on to the next project; otherwise, shut down your computer.

Token Ring Networks

Developed by IBM in the mid-1980s, the **token ring** network technology provides reliable, albeit slow by today's standards, transport of data. Based on the IEEE 802.5 standard, token ring networks are cabled in a physical star topology but function as a logical ring, as shown earlier in Figure 3-8. Token ring originally operated at 4 Mbps, but this speed increased to 16 Mbps and later to 100 Mbps. A 1000 Mbps standard was approved in 2001, but by that time, the token ring technology had clearly lost out to 100 Mbps Ethernet, and no 1000 Mbps products were ever manufactured in quantity. Most token ring networks used Category 4 or higher UTP.

Token Ring Media Access Token ring uses the token-passing media access method, which is where the technology gets its name. Using this method, a special frame called the "token" passes from one computer to the next. Only the computer holding the token can send data, and a computer can keep the token for only a specific amount of time. If the computer with the token has no data to send, it passes the token to the next computer.

Because only the computer with the token can transmit data, this method prevents collisions. Computers no longer spend time waiting for collisions to be resolved, as they do in a CSMA/CD network. All computers have equal access to the medium, which makes token-passing networks best suited for time-sensitive environments, such as banking transactions and databases requiring precise timestamps. Also, because traffic moves in a specific "direction" around a ring topology, faster access methods (such as 100 Mbps token ring) can circulate two tokens at the same time without fear of collision. (By keeping the two sets of messages from overlapping, both tokens can circulate in order.)

However, token passing has two disadvantages. First, even if only one computer on the network has data to send, it must wait to receive the token. If its data is large enough to warrant two or more "turns" at the token, the computer must wait until the token makes a complete circuit before starting its second transmission. Second, the complicated process of

creating and passing tokens requires more expensive equipment than what's used on CSMA/CD networks. This additional expense and complication is in part what led to token ring quickly becoming second best in LAN technologies, compared with 100 Mbps and switched Ethernet. Because token ring is no longer a widely used LAN technology, additional operating details about it have been moved to Appendix B.

Fiber Distributed Data Interface Technology

Fiber Distributed Data Interface (FDDI) uses the token-passing media access method and dual rings for redundancy. The rings in an FDDI network are usually a physical ring of fiber-optic cable. FDDI transmits at 100 Mbps and can include up to 500 nodes over a distance of 100 kilometers (60 miles). FDDI full-duplex technology, an extension to standard FDDI, can support up to 200 Mbps. Like token ring, FDDI uses token passing; however, FDDI's token-passing scheme is based on IEEE 802.4 rather than IEEE 802.5. An FDDI network has no hubs; devices generally connect directly to each other. However, devices called "concentrators" can serve as a central connection point for buildings or sites in a campus setting.

Much like token ring, FDDI technology lost out to faster versions of Ethernet and is now obsolete for new network designs. It had its heyday in the early to mid-1990s when Ethernet was operating at only 10 Mbps and switched Ethernet was just being developed. If you want to buy FDDI products now, you'll be purchasing used products, as new products aren't available. Therefore, details on FDDI have been relegated to Appendix B.

Internet Access Technologies

This section describes the Internet access technologies most commonly used by small office/home office (SOHO) networks to share Internet access for all workstations on the network: cable modem, DSL, and satellite. These technologies are in contrast to the WAN technologies often used by large businesses to access the Internet and communicate with remote offices. WAN technologies are discussed in Chapter 12.

Cable Modem Cable modem networking is a broadband technology used to deliver Internet access to homes and businesses over standard cable television (CATV) coaxial cable. Because it's a broadband technology, data delivered to a cable modem shares the same cable as the channels delivered to your TV. In fact, Internet data simply travels on a TV channel that's not used by the cable company. The official standard governing cable modem operation is **Data Over Cable Service Interface Specification (DOCSIS)**. Although cable modems are considerably more complicated than dial-up modems, they are true modems in the sense that they modulate and demodulate signals.

Cable modem networks share some properties of traditional 10Base2 Ethernet. They are shared-media, bus topology networks at the point where data is delivered to a home. Other parts of a cable modem network use high-speed WAN technologies, as shown in Figure 3-23.

Cable modems have exploded in popularity because of the high speeds at which Internet data can be delivered to homes and businesses. With the newest DOCIS 3.0 standard (called Wideband Internet), Internet data can be delivered to customers at speeds up to 60 Mbps, although most ISPs don't support this standard in all areas as of this writing. Most cable Internet companies still deliver between 1 Mbps and 10 Mbps bandwidth to their customers.



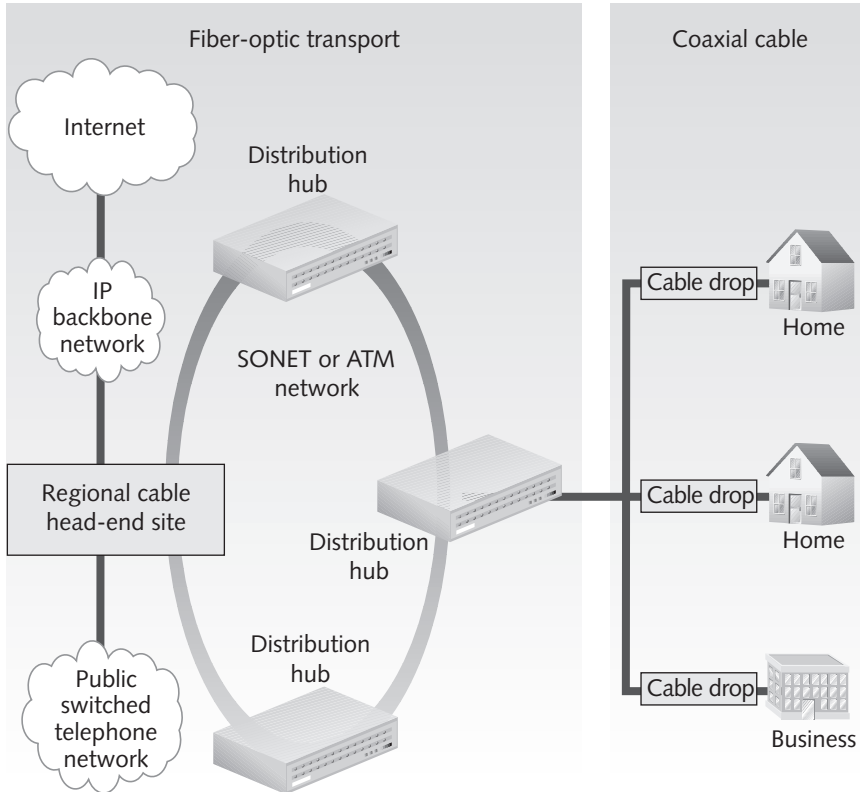


Figure 3-23 A typical cable modem network

Courtesy of Course Technology/Cengage Learning

Cable modems use an asymmetrical communication scheme—data rates going to the home (downstream rates) are higher than data rates coming from the home to the cable provider (upstream rates). Upstream data rates can be as much as 10 Mbps but are usually limited to between 256 Kbps and 1 Mbps.

Cable modems provide bandwidth to users as a form of shared media access. That is, all users on a CATV cable segment (usually part of a neighborhood or large building, for example) share available bandwidth. Therefore, more users (or more traffic per user) mean less bandwidth per user because access is shared. One powerful advantage of cable modem access, however, is that distance limitations don't govern functionality. As long as users have access to cable TV and the cable company offers Internet access on the local cable segment, users can install a cable modem and access the Internet.



Shared access in this context means the medium from your cable connection to the cable company (the ISP) is shared by other users in your location. After data reaches the ISP and goes out to the Internet, the medium is shared by all other Internet users.

Aside from performance issues caused by shared access media in cable modem networks, security was a concern in early cable networks because users who shared the same cable segment could eavesdrop on others' communication sessions. However, networks complying with the DOCSIS standard use a strong encryption key for each user connection, which ensures privacy and security on the shared portion of the cable network.

Cable Modem Operation From a user's standpoint, a cable modem is straightforward: Plug the cable from the wall outlet into the cable modem, and then plug a network cable from the cable modem to a PC or router (Figure 3-24).




Figure 3-24 A typical cable modem connection

Courtesy of Course Technology/Cengage Learning

One reason for the success of cable Internet access is that most people already have cable TV in their homes, so no additional cabling is required, and people are comfortable with cable TV technology. The cable modem, like your TV, has a tuner that “tunes in” the frequencies for upstream and downstream channels for Internet access. From there, these signals are converted (modulated) into the Ethernet signals your computer or router requires. When you send data, the signals are modulated into a form required by the cable medium. A cable modem even has a MAC address, which serves two purposes. The cable company uses it to determine whether your device is allowed on the network, and much like a NIC, a cable modem compares the destination address of incoming data to determine whether the modem should process arriving data.



TIP

For a wealth of information on cable modem technologies, see www.cablelabs.com/cablemodem/.



Digital Subscriber Line Digital Subscriber Line (DSL) warrants special mention in this section, as it competes with cable modem technologies for Internet access. This broadband technology uses existing phone lines to carry voice and data simultaneously. Many variations of DSL are available; the most prominent for home Internet access is **Asymmetric DSL (ADSL)**, named because the download and upload speeds differ substantially, so the data rates aren't symmetrical. ADSL splits the phone line into two frequency ranges: Frequencies below 4 KHz are used for voice transmission, and frequencies above 4 KHz are used to transmit data. Typical connection speeds for downloading data range from 256 Kbps to 8 Mbps; upload speeds are typically much slower, in the range of 16 Kbps to 640 Kbps.

To deliver digital services, DSL uses the same twisted-pair phone lines that deliver voice services. Unlike cable modem connections, DSL media aren't shared in the connection between the user and ISP, so they offer subscribers guaranteed bandwidth. Because bandwidth is guaranteed, however, upstream and downstream data rates are metered. Users must pay more for higher bandwidth connections. Nevertheless, DSL is a great SOHO technology, particularly in areas where cable modem is unavailable. Most DSL connections top out at around 1.5 Mbps bandwidth, but it's possible to achieve rates of about 6 Mbps downstream.

DSL's main disadvantage is its distance limitation (measured as the wire runs) between the user's location and the nearest central office (CO), where a copper-to-fiber interface device links to the telecommunication carrier's digital backbone. Depending on which vendor's equipment is used, this distance limitation varies between 17,500 feet (3.31 miles or 5.33 kilometers) and 23,000 feet (4.36 miles or 7.01 kilometers). Therefore, it's important to measure a connection point's distance from the local CO to determine whether DSL is a viable network option.

Another DSL variation is **Symmetric DSL (SDSL)**, and its upload and download speeds are the same, so if the download speed is 4 Mbps, the upload speed is also 4 Mbps. SDSL is often chosen for businesses operating a Web site because the amount of traffic uploaded and downloaded is likely to be similar.

DSL and cable modems share one important advantage over asynchronous modems—they're always on. Both technologies maintain constant connections to a remote server on the other side of the connection, so there's little or no delay to establish a connection, as with a conventional modem. Given the higher bandwidth, faster access, and lower cost of digital connections, it's no wonder that droves of users are switching from dial-up modems to digital alternatives for SOHO connections.

**TIP**

For more information on DSL, see www.dslreports.com.

Satellite Technologies If neither DSL nor cable modem are available and dial-up speeds are too frustrating, satellite Internet is another option. Speeds are comparable with where cable modem was several years ago and many DSL providers still are: download speeds of 1.5 Mbps and uploads speeds of about 256 Kbps. For quite a bit more money,

satellite download speeds can be increased to up to 5 Mbps with 300 Kbps upload speeds. Unfortunately, most satellite offerings have limitations on the daily amount of data that can be downloaded, making these options somewhat of a last resort for high-speed Internet.

Early implementations of satellite Internet required a regular dial-up modem for the upstream connection, but that's no longer the case. Two well-known satellite Internet providers are HughesNet and WildBlue. WildBlue focuses on residential customers only; HughesNet provides service for residential customers and offers some higher bandwidth options for businesses.

WiMAX: Wireless Internet Access Worldwide Interoperability for Microwave Access (WiMAX) comes in two flavors: 802.16-2004 (previously named 802.16d), or fixed WiMAX, and 802.16e, or mobile WiMAX. These standards provide wireless broadband to outlying and rural areas, where last-mile wired connections (connections between a service provider and homes or businesses) are too expensive or impractical because of rough terrain, and to mobile users so that they can maintain a high-speed connection while on the road. Fixed WiMAX delivers up to 70 Mbps of bandwidth at distances up to 30 miles, and mobile WiMAX has a coverage area of 3 to 10 miles. Developing standards promise up to 100 Mbps for mobile WiMAX and up to 1 Gbps for fixed WiMAX.

Besides providing wireless network service to outlying areas, fixed WiMAX is being used to deliver wireless Internet access to entire metropolitan areas rather than the limited-area hotspots available with 802.11. It can blanket an area up to a mile in radius, compared with just a few hundred feet for 802.11. Los Angeles has implemented WiMAX in an area of downtown that encompasses a 10-mile radius.

Mobile WiMAX brings broadband Internet roaming to the public. Although fixed WiMAX can create a wider hot spot than 802.11 wireless networks, network users are still confined to the coverage area. After a user leaves the coverage area of a transmitting station, his or her connection is dropped. Mobile WiMAX enables users to roam from area to area without losing the connection, which offers mobility much like cell phone users have. It's considered a fourth-generation (4G) wireless technology, and companies such as Clearwire (in its CLEAR product) are using it to provide Internet access to people who want to stay connected while on the move. CLEAR offers typical bandwidth between 3 and 6 Mbps. Several computer manufacturers, such as Dell, Toshiba, and others, offer WiMAX-enabled laptops, or you can buy a USB WiMAX modem to connect a laptop or desktop computer to a WiMAX network. In addition, Clearwire offers CLEAR Spot devices, which create mobile hotspots to connect up to eight Wi-Fi-enabled computers to a CLEAR network.



To read more about this evolving technology, go to www.wimax.com.

TIP



Chapter Summary

- Networks can be described by a physical and logical topology. The physical topology describes the arrangement of cabling that connects one device to another. The logical topology describes the path data travels between devices. The logical and physical topology can be, and often are, different.
- The main physical topologies are the bus, star, ring, and point-to-point. A physical bus topology is simple but is no longer in common use because of a number of weaknesses. A star topology, along with the extended star, is the most common for implementing LANs. A physical ring topology isn't in widespread use now but was used mainly in network backbones. Point-to-point topologies are used primarily in WANs and with wireless bridges. Several point-to-point connections can create a mesh topology for the purpose of redundancy.
- The main logical topologies are bus, ring, and switched. A logical bus can be implemented as a physical star or a physical bus and is used with hub-based Ethernet and Wi-Fi networks. A logical ring can be implemented as a physical ring or a physical star and is most commonly seen in token ring and FDDI networks. The switched topology uses a physical star and is used with Ethernet networks using a switch in the center of a star physical topology.
- A network technology defines the structure of frames and how a network interface accesses the medium to send frames. It often defines the media types that must be used to operate correctly.
- The most common network technology for LANs is Ethernet. It's described in IEEE 802.3 and has many subcategories, including 10BaseT, 100BaseT, and 1000BaseT, that use twisted-pair copper cabling. Ethernet uses the CSMA/CD access method, which is turned off when a full-duplex connection is established. Other Ethernet standards include fiber-optic implementations, such as 100BaseFX and 1000BaseLX, among others.
- Wi-Fi is a wireless technology based on Ethernet, but it uses the CSMA/CA media access method. The most common Wi-Fi standards are 802.11b, 802.11g, 802.11a, and 802.11n with speeds of 11 Mbps, 54 Mbps, 54 Mbps, and up to 600 Mbps, respectively.
- Token ring and FDDI are obsolete technologies that used a token-passing access method. Token ring operated at speeds of 4 Mbps and 16 Mbps and ran over twisted-pair cabling, whereas FDDI ran over fiber-optic cabling at 100 Mbps.
- Internet access technologies include cable modem, DSL, satellite, and WiMAX. Cable modem is among the most common of these technologies, as it's usually available wherever there's cable TV access. DSL is widely available, but you must live within a few miles of your phone provider. Satellite is becoming somewhat more common, replacing dial-up in areas that don't have other high-speed options. WiMAX is fairly new and provides wireless Internet access to outlying areas and can give mobile users in a metropolitan area fast Internet access.

Key Terms

1000BaseT Ethernet A technology defined by the IEEE 802.3ab standard, supports 1000 Mbps Ethernet (usually called Gigabit Ethernet) over Category 5 or higher UTP cable, using baseband signaling.

100BaseFX 100 Mbps Ethernet using baseband signaling over two strands of fiber-optic cabling.

100BaseTX A technology defined by IEEE 802.3u, it's the most commonly used Ethernet variety today. It runs over Category 5 or higher UTP cable and uses two of the four wire pairs: one to transmit data and the other to receive data. It runs at 100 Mbps, using baseband signaling.

10BaseT A technology defined by IEEE 802.3i, it's Ethernet running at 10 Mbps, using baseband signaling over Category 3 or higher twisted-pair cabling. Although still seen in older networks, newer networks use 100BaseT or faster technology.

10GBaseT A technology defined by IEEE 802.3an, it's 10 Gigabit Ethernet running over four pairs of Category 6A UTP cabling, using baseband signaling. Unlike the other BaseT Ethernet standards, 10GBaseT operates only in full-duplex mode.

ad hoc mode Sometimes called peer-to-peer mode, it's a wireless mode of operation typically used only in small or temporary installations. There's no central device, and data travels from one device to another to reach the destination device.

Asymmetric DSL (ADSL) A DSL variation in which the download and upload speeds differ substantially, so the data rates aren't symmetrical. Typical connection speeds for downloading data range from 256 Kbps to 8 Mbps; upload speeds are typically much slower, in the range of 16 Kbps to 640 Kbps. *See also* Digital Subscriber Line (DSL).

baseband A type of signaling used in networks, in which each bit of data is represented by a pulse of electricity (on copper media) or light (on fiber-optic media). These signals are sent at a single fixed frequency, using the medium's entire bandwidth. LAN technologies use baseband signaling.

broadband A type of signaling that uses analog techniques to encode binary 1s and 0s across a continuous range of values. Broadband signals move across the medium in the form of continuous electromagnetic or optical waves rather than discrete pulses. Signals flow at a particular frequency, and each frequency represents a channel of data, allowing multiple streams of data on a single wire. TV and cable Internet use broadband signaling.

Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) An access control method used by Wi-Fi networks, in which an acknowledgement is required for every packet sent, thereby avoiding most possibilities of a collision (collision avoidance).

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) A media access method in which a device must first listen (carrier sense) to the medium to be sure no other device is transmitting. If two devices transmit at the same time (multiple access), a collision occurs and is detected (collision detection). In this case, all devices involved in the collision wait for a random period of time before transmitting again.

collision domain The extent to which signals in an Ethernet bus topology network are propagated. All devices connected to a logical bus topology network are in the same collision domain. Switch and router ports delimit collision domains.



collision The result of two or more devices on the same medium transmitting simultaneously when CSMA/CD is the media access method in use. *See also* Carrier Sense Multiple Access with Collision Detection (CSMA/CD).

Cyclic Redundancy Check (CRC) The error-checking code in an Ethernet frame's trailer; it's the result of a mathematical algorithm computed on the frame data. When the destination device receives the frame, the calculation is repeated. If the results of this calculation don't match the CRC in the frame, it indicates the data was altered in some way.

Data Over Cable Service Interface Specification (DOCSIS) The official standard governing cable modem operation.

Digital Subscriber Line (DSL) A broadband technology that uses existing phone lines to carry voice and data simultaneously.

extended star topology An extension of the physical star topology, in which a central switch or hub is the central connecting point for other switches or hubs that have computers and other network devices attached, forming a star of stars. *See also* physical star topology.

Fiber Distributed Data Interface (FDDI) A technology that uses the token-passing media access method and dual rings for redundancy. The rings in an FDDI network are usually a physical ring of fiber-optic cable. FDDI transmits at 100 Mbps and can include up to 500 nodes over a distance of 100 kilometers.

frame types The frame formats that describe the content and length of a frame header.

Gigabit Ethernet *See* 1000BaseT Ethernet.

infrastructure mode An operational mode for Wi-Fi networks, in which wireless stations connect through a wireless access point before they can begin communicating with other devices.

logical topology The path data travels between computers on a network. The most common logical topologies are switched, bus, and ring.

media access control *See* media access method.

media access method A set of rules governing how and when the network medium can be accessed for transmission. The rules ensure that data is transmitted and received in an orderly fashion, and all stations have an opportunity to communicate. Also called media access control.

mesh topology A topology in which each device in the network is connected to every other device, providing multiple pathways in the event of a device or cable failure.

network backbone The cabling used to communicate between LANs or between hubs or switches. The backbone cabling often runs at a faster speed than the cabling used to connect computers because the backbone must carry data from many computers to other parts of the network.

physical bus topology A network topology in which a continuous length of cable connects one computer to another in daisy-chain fashion. There's no central interconnecting device.

physical ring topology A cabling arrangement in which each device is connected to another device in daisy-chain fashion, and the last device connects back to the first device forming a ring. Used by token ring and FDDI, the physical ring is rarely used now.

physical star topology A network topology that uses a central device, such as a hub or switch, to interconnect computers in a LAN. Each computer has a single length of cable going from its NIC to the central device. It's the most common physical topology in LANs.

physical topology The arrangement of cabling and how cables connect one device to another in a network. The most common physical topology is a star, but bus, ring, point-to-point, and mesh topologies are also used.

point-to-point topology A topology in which cabling creates a direct link between two devices; used most often in WANs or in wireless networks to create a wireless bridge.

reflection *See* signal bounce.

signal bounce The result of electricity bouncing off the end of a cable and back in the other direction. It causes corruption of data as the bouncing signal collides with signals behind it. A terminator at each cable end is needed to prevent signal bounce. Also called reflection.

signal propagation Signals traveling across a medium and through any connectors and connecting devices until the signal weakens enough to be undetectable or is absorbed by a termination device.

Symmetric DSL (SDSL) A DSL variation in which the download and upload speeds are equivalent, or symmetrical. *See also* Digital Subscriber Line (DSL).

terminator An electrical component called a resistor, placed at the ends of a physical bus network to absorb the signal instead of allowing it to bounce back up the wire.

token ring A technology based on the IEEE 802.5 standard; its cabling is in a physical star topology, but it functions as a logical ring. It uses the token-passing media access method, and only the computer holding the token can send data.

wireless bridge An operational mode of wireless networking usually used to connect two wired LANs that are separated from each other in such a way that using physical media is impractical. Can also be used to extend the reach of a wireless network.

Wireless Fidelity (Wi-Fi) The name given to the 802.11 series of IEEE standards that define four common varieties of wireless LANs: 802.11a, 802.11b, 802.11g, and 802.11n.

Worldwide Interoperability for Microwave Access (WiMAX) A wireless broadband technology defined in 802.16-2004 for fixed WiMAX and 802.16e for mobile WiMAX. WiMAX is considered a fourth-generation (4G) technology for bringing wireless Internet access to remote areas, large areas up to a mile radius, and mobile users.



Review Questions

1. Which of the following describes the arrangement of network cabling between devices?
 - a. Logical topology
 - b. Networking technology
 - c. Physical topology
 - d. Media access method
2. Which of the following is an advantage of a star topology? (Choose all that apply.)
 - a. Allows faster technologies than a bus does
 - b. Requires less cabling than a bus
 - c. Centralized monitoring of network traffic
 - d. No single point of failure

3. Which of the following is an example of a technology using a physical ring topology?
 - a. Token ring
 - b. FDDI
 - c. ADSL
 - d. IEEE 802.5
4. Which technology is likely to be implemented as a point-to-point physical topology?
 - a. Wi-Fi infrastructure mode
 - b. FDDI
 - c. Ethernet
 - d. Wireless bridge
5. Which of the following describes a hub-based Ethernet network?
 - a. Physical bus
 - b. Logical bus
 - c. Physical switching
 - d. Logical star
6. Which of the following is a characteristic of a logical ring topology? (Choose all that apply.)
 - a. It's used by Ethernet.
 - b. One technology uses an MAU.
 - c. It's used by FDDI.
 - d. Some technologies use a token.
 - e. It's the most popular logical topology.
7. Which best describes a typical wireless LAN?
 - a. Logical ring topology
 - b. Logical switching topology
 - c. Logical bus topology
 - d. Logical star topology
8. Which of the following is a characteristic of a switched logical topology? (Choose all that apply.)
 - a. Uses a physical bus topology
 - b. Creates dynamic connections
 - c. Sometimes called a shared-media topology
 - d. Uses a physical star topology
9. Which of the following is a characteristic of unshielded twisted-pair cabling? (Choose all that apply.)
 - a. Consists of four wires
 - b. Commonly used in physical bus topologies

- c. Has a distance limitation of 100 meters
 - d. Susceptible to electrical interference
10. Which of the following is a characteristic of fiber-optic cabling? (Choose all that apply.)
- a. Can be used in electrically noisy environments
 - b. Requires only a single strand of fiber for network connections
 - c. Carries data over longer distances than UTP does
 - d. Lower bandwidth capability
11. Which topology most likely uses coaxial cabling?
- a. Physical star
 - b. Logical ring
 - c. Physical bus
 - d. Logical switching
12. Which of the following is true of a MAC address?
- a. All binary 1s in the source address indicates a broadcast frame.
 - b. It's sometimes called a logical address.
 - c. A destination address of 12 hexadecimal Fs is a broadcast.
 - d. It's composed of 12 bits.
13. Which of the following is the most commonly used Ethernet frame type?
- a. Ethernet II
 - b. Ethernet SNAP
 - c. Ethernet 802.3
 - d. Ethernet 802.2
14. Which of the following is a field of the most common Ethernet frame type? (Choose all that apply.)
- a. ARP trailer
 - b. FCS
 - c. Destination MAC Address
 - d. Data
 - e. MAC type
15. Which access method uses a “listen before sending” strategy?
- a. Token passing
 - b. CSMA/CD
 - c. Token bus
 - d. Polling
16. Which of the following is true about full-duplex Ethernet? (Choose all that apply.)



- a. Stations can transmit and receive but not at the same time.
 - b. Collision detection is turned off.
 - c. It's possible only with switches.
 - d. It allows a physical bus to operate much faster.
17. Which of the following is defined by the extent to which signals in an Ethernet bus topology network are propagated?
- a. Physical domain
 - b. Collision domain
 - c. Broadcast domain
 - d. Logical domain
18. Which of the following is considered a property of Ethernet? (Choose all that apply.)
- a. Scalable
 - b. Best-effort delivery system
 - c. Guaranteed delivery system
 - d. Obsolete technology
19. Which of the following is true of IEEE 802.3an?
- a. Requires two pairs of wires
 - b. Uses Category 5 or higher cabling
 - c. Currently best for desktop computers
 - d. Operates only in full-duplex mode
20. Which of the following is a feature of 100BaseFX? (Choose all that apply.)
- a. Often used as backbone cabling
 - b. Best when only short cable runs are needed
 - c. The fastest of the Ethernet standards
 - d. Uses two strands of fiber
21. Which Wi-Fi standard can provide the highest bandwidth?
- a. 802.11a
 - b. 802.11b
 - c. 802.11n
 - d. 802.11g
22. Which of the following is true about infrastructure mode in wireless networks? (Choose all that apply.)
- a. Best used for temporary networks
 - b. Uses a central device
 - c. Resembles a physical bus and logical ring
 - d. Most like a logical bus and physical star

23. How many channels can be used on an 802.11b network in North America?
- 7
 - 9
 - 11
 - 13
24. Which media access method does Wi-Fi use?
- CSMA/CD
 - Token bus
 - Demand priority
 - CSMA/CA
25. Which of the following is true about the token ring technology? (Choose all that apply.)
- It uses a physical ring topology.
 - All computers have equal access to the media.
 - It uses RTS/CTS signaling before transmission can occur.
 - Only the computer with the token can transmit data.



Challenge Labs



Challenge Lab 3-1: Building an Extended Star Topology Network

Time Required: 30 minutes

Objective: Use hubs and switches to build an extended star topology network.

Required Tools/Equipment: Determine which type of devices and how many you need to build the network.

Description: In this lab, you build an extended star network, in which the computers are connected in a physical star and a logical bus topology, and the computers form the outer arms of the extended star. The center of the extended star should be a device that creates one collision domain per port. Build the network with as much equipment as you have available, distributing computers evenly around the outer edges of the extended star. Draw the final topology and label the devices. If you lack equipment, you can simply draw the topology without building the physical network. Then answer the following questions:

- What type of device are the computers attached to?
-

- What type of device is at the center of the extended star?

- How many collision domains are in this network?



Challenge Lab 3-2: Adding Wireless Access to the Extended Star Network

Time Required: 30 minutes

Objective: Add wireless networking to the extended star network you built in Challenge Lab 3-1.

Required Tools/Equipment: An access point or wireless router and some wireless NICs

Description: Add wireless networking to the extended star network you built in Challenge Lab 3-1. Expand the drawing you created to include the AP or wireless router. If you don't have the necessary equipment, just expand the drawing. Answer the following questions:

- Which device in your extended star did you connect the AP to and why?

- Which wireless mode are you using: ad hoc or infrastructure?

- What logical and physical topology does adding wireless bring to this network?



Challenge Lab 3-3: Installing inSSIDer

Time Required: 20 minutes

Objective: Install a wireless scanning tool and scan your network.

Required Tools/Equipment: A computer with a wireless NIC and access to the Internet or an already downloaded copy of inSSIDer

Description: In this lab, you download inSSIDer from <http://metageek.net> and install it on a computer with a wireless NIC. Your instructor might need to install it for you if you don't have the necessary permissions. After it's installed, start a scan of your network to look for access points. Answer the following questions:

- Approximately how many wireless networks did inSSIDer find?

- Which wireless channels are the most heavily used?

- If you were to set up a new wireless LAN based on what inSSIDer found, what channel would you use for the network?



Case Projects



Case Project 3-1

Old-Tech Corporation has 10 computers in its main office area, which is networked in a star topology using 10 Mbps Ethernet hubs, and wants to add five computers in the manufacturing area. One problem with the existing network is data throughput. Large files are transferred across the network regularly, and the transfers take quite a while. In addition, when two or more computers are transferring large files, the network becomes unbearably slow for users. Adding the manufacturing computers will only make this problem worse and result in another problem. Because the ceiling is more than 30 feet high, there's no easy way to run cables to computers, and providing a secure pathway for cables is next to impossible. Devise a solution to this company's networking problems. As part of your solution, answer the following questions:

- What changes in equipment are required to bring this company's network up to date to solve the shared-bandwidth problem?

- What topology and which type of device can be used in the manufacturing area to solve the cabling difficulties?

Case Project 3-2

EBiz.com has 250 networked computers and five servers and uses a star topology wired network to reach employees' offices, with a bus interconnecting three floors in its office building. Because of a staggering influx of Internet business, the network administrator's task is to boost network performance and availability as much as possible. The company also wants a network design that's easy to reconfigure and change because workgroups form and disband frequently, and their membership changes regularly. All computers must share sensitive data and control access to customer files and databases. Aside from the customer information and billing databases, which run on all servers, employees' desktop computers must run standard word-processing and spreadsheet programs.

Use the following write-on lines to evaluate the requirements for this network. After you finish, determine the best network topology or topology combination for the company. On a blank piece of paper, sketch the network design you think best suits EBiz.com's needs. Remember: High performance and easy reconfiguration are your primary design goals!

- What type of topology should be used in this network?

- Will the network be peer to peer or server based?

- How many computers will be attached to the network?

- What kind of networking device is easiest to reconfigure? What kind offers the best access to the network medium's bandwidth between pairs of devices?

Case Project 3-3

ENorm, Inc. has two sites in Pittsburgh that are four miles apart. Each site consists of a large factory with office space for 25 users at the front of the factory and up to 20 workstations in two work cells on each factory floor. All office users need access to an inventory database that runs on a server at the Allegheny Street location; they also need access to a billing application with data residing on a server at the Monongahela site. All factory floor users also need access to the inventory database at the Allegheny Street location.

Office space is permanently configured, but the manufacturing space must be reconfigured before each new manufacturing run begins. Wiring closets are available in the office space. Nothing but a concrete floor and overhead girders stay the same in the work cell areas. The computers must share sensitive data and control access to files. Aside from the two databases, which run on the two servers, office computers must run standard word-processing and spreadsheet programs. Work cell machines are used strictly for updating inventory and quality control information for the Allegheny Street inventory database. Workstations in the manufacturing cells are switched on only when they're in use, which might occur during different phases of a manufacturing run. Seldom is a machine in use constantly on the factory floor.

Use the following write-on lines to evaluate the requirements for this network. After you finish, determine the best network topology or topology combination for the company. On a blank piece of paper, sketch the network design you think best suits ENorm, Inc.'s needs.

- Will the network be peer to peer or server based?

- How many computers will be attached to the network?

- What topology works best for the offices, given the availability of wiring closets? What topology works best for the factory floor, given its need for constant reconfiguration?



This page intentionally left blank



Network Media

After reading this chapter and completing the exercises, you will be able to:

- Define the primary cables used in wired networking
- Describe the characteristics of the major types of fiber-optic media
- Explain the technologies used for wireless networking

Network media are the materials through which network signals travel between devices. They can be a physical material, such as copper wire or glass fiber, or simply the air. When a physical material is used as the medium, it's usually referred to as "wired networking," and when signals are transmitted through the air, the medium is aptly called "wireless networking."

In this chapter, you learn about common options for wired and wireless networking and where these options make sense. You learn about the characteristics of wired media and how to choose a media type to suit a situation and environment. You also learn how to install and terminate the most common types of LAN media. In addition, you learn about transmission technologies for making wireless network links for both short-range Wi-Fi networks and long-range wireless networks.

Wired Networking

Wired networking uses tangible physical media called cables. Cables used in networking come in two broad categories: copper wire and fiber optic. Regardless of the material used, all networking cables must support the basic tasks of sending and receiving bit signals. The composition of these signals (electricity or light), the speed at which these signals can be sent (bandwidth), and the distance they can effectively travel, make up the main differences between cabling types.

The following sections discuss cable characteristics, the criteria for choosing a particular type of cabling, and a variety of cable types, both copper and fiber optic.

Criteria for Choosing Network Media

All cables share certain fundamental characteristics you should know to understand their function and correct use. Even though copper cables differ radically from fiber-optic cables in composition and types of signals they carry, the characteristics described in the following sections apply equally to both types of cabling.

Bandwidth Rating Bandwidth, the number of bits per second that can be transmitted across a medium, is as much a function of the technology used to transmit bit signals as it is of the medium. For example, Category 5 UTP cabling was originally intended to support only up to 100 Mbps but was later upgraded to support up to 1000 Mbps when the 1000BaseT standard was developed.

What really determines the bandwidth of a particular cabling type is how fast a transmitting device, such as a NIC, can generate bit signals on the medium and whether these signals can be received accurately at the other end of the cable. Bit signals lose strength as they travel along the medium, so when judging whether a cabling type is suitable for a particular transmission speed, the maximum cable length must also be considered.

Another factor determining bandwidth is how bit signals are represented on the medium, a process called **encoding**. Different networking standards use different patterns of electrical or light pulses to represent a series of bits on the medium.

Although different media types and cable grades can support higher bandwidths than others, what's most important is choosing the media type and cable grade specified by the

networking standard you want to run. Keep in mind that today's 100BaseT network might be tomorrow's 1000BaseT network. So when possible, choose a cabling category that's compatible with the standard you want to implement now but will support the next level of speed your network is likely to need in the future.

Maximum Segment Length A **cable segment** is a length of cable between two network devices, such as a NIC and a switch. Any intermediate passive (unpowered) devices, such as wall jacks, are considered part of the total segment length.

Each cable type can transport data at a particular speed only so far before its signals begin to weaken past the point that a receiving station can read them accurately; this phenomenon is called **attenuation**. In addition, electrical signals are affected by electromagnetic interference, or "noise." The longer a signal travels down a cable segment, the more likely it is that electrical noise impairs the signal to the point that data can be misinterpreted. (For example, a 0 bit is read as a 1 bit.) An internetwork can be constructed of several cable segments, as long as the hardware connecting them (such as switches and routers) can capture and regenerate the incoming signal at full strength.

Interference and Eavesdropping Susceptibility How well a media type resists signal interference from outside sources depends on the medium's construction and type of signals it's designed to carry. Interference to electrical signals on copper media comes in the form of **electromagnetic interference (EMI)** and **radio frequency interference (RFI)**. Motors, transformers, fluorescent lights, and other sources of intense electrical activity can emit both EMI and RFI, but RFI problems are also associated with the proximity of strong broadcast sources in an environment (such as a nearby radio or TV station). RFI can also affect wireless networks if the frequencies are in the same range in which the wireless network is operating.

Another type of interference found in copper wires is a form of EMI called **crosstalk**, which is interference one wire generates on another wire when both wires are in a bundle (as all cabling in LANs is). When electrical signals travel across the medium, they create their own electromagnetic field. Although this field is weak, it can leak onto other wires, especially when their insulation is in contact with the other wire. Although it's not as common now, you might have experienced crosstalk while talking on a landline phone and hearing another conversation faintly. With phone wires, crosstalk is merely an annoyance because people can filter out this noise easily, but in networking, excessive crosstalk can render the network connection unusable.

Because electrical signals traveling down a copper wire create an electromagnetic field that can be detected outside the wires, copper wire is susceptible to electronic eavesdropping. It might sound like the stuff of spy movies, but with the right type of equipment, an eavesdropper simply needs to get close to a copper cable to extract data from it. In the absence of sensitive electronic equipment, if eavesdroppers have physical access to the connecting equipment and the copper wire is slightly exposed, they would have no problem installing a listening device directly on these wires.

Fiber-optic media carries light signals and is impervious to interference. In addition, because no magnetic field is present, eavesdropping is a difficult proposition with fiber-optic cable. To eavesdrop, someone needs access to the glass strands carrying the optical signals to



install a device that captures data and prevents the connection from being broken. Not impossible, but extremely difficult.

When choosing a cable type, the environment in which the medium operates is one of the most crucial factors in the decision. The decision is usually between copper cabling and fiber-optic cabling for high-performance applications and between copper cabling and wireless for less bandwidth-heavy applications.

Cable Grade Building and fire codes include specific cabling requirements, usually aimed at the combustibility and toxicity of the jacket and insulation covering most cables. Polyvinyl chloride (PVC) covers the cheapest and most common cables (for example, the 120-volt cord in lamps and other household appliances). Unfortunately, when this material burns, it gives off toxic fumes, which makes it unsuitable for cables strung in ceilings or inside walls.

The space between a false ceiling and the true ceiling in most office buildings, called the “plenum,” is commonly used to aid air circulation for heating and cooling. Any cables in this space must be plenum-rated, which typically means they’re coated with Teflon because of its low combustibility and the nontoxic fumes it produces when burned. These cables can be used in the plenum or inside walls without being enclosed in conduit. Although plenum-rated cable is nearly twice as expensive as non-plenum-rated cable, eliminating the need for conduit makes installing plenum-rated network cabling much cheaper. UTP cabling is usually marked as communication cable riser (CMR) or communication cable plenum (CMP). CMR is suitable only for building risers, such as elevator shafts or in cable trays, and can’t be used in spaces that carry environmental air. CMP is suitable for use in plenum spaces. Before installing any type of cable, check all local fire and building codes because requirements vary widely.

Connection Hardware Every type of cable has connectors that influence the kinds of hardware the cable can connect to and affect the costs of the resulting network. This section also explains whether these connectors are easy to attach or need specialized equipment and whether building cables should be left to professionals. You learn how to install the connectors used in UTP cabling, which are the least expensive connectors. Fiber-optic connectors tend to be expensive, as are the tools used to attach them.

Other Media Considerations Additional media considerations include ease of installation, testability, and of course cost:

- *Ease of installation*—The difficulty of installing a cable plant has a bearing on your choice of media. **Cable plant** is the term for all the cables and connectors tying your network together. Sometimes you have to make a tradeoff between the highest quality available and the cost and time factors involved in installing the medium correctly. Some factors that must be considered are a medium’s minimum bend radius, which limits the angle at which a cable can be bent to run around corners; the cost and time needed to terminate the medium, which involves installing connectors and attaching media to patch panels and jacks; and the physical environment. (Cinderblock or plaster walls, concrete floors, and high ceilings can make installing a cable plant cost prohibitive, for example.) You might decide to make parts of your network wireless because of some of these factors.

- *Testability*—How difficult and expensive is it to test the medium after it’s installed? Declaring a cable installation successful just because computers can communicate doesn’t really constitute a test. A network that “works” might be crippled by excessive transmission errors caused by poor cable termination. A true test of cabling, whether it’s copper or fiber optic, is to install it, add the connectors and other termination points, and then test it with a device that can certify whether the cable meets the requirements for its category. Simple testers that check for basic electrical or optical connectivity are inexpensive (a few hundred dollars or less) but don’t give you a true picture of your cable plant. Copper wire certifiers that do a full battery of Category 5 and above tests start at about \$1000, and those capable of fiber-optic testing range up to more than \$10,000.
- *Total cost*—When figuring the total cost for media, you must include the cabling, connectors, termination panels, wall jacks, termination tools, testing equipment, and, of course, time. The complexity of a large media installation (for a new building, for example) can be daunting, which is why there are companies specializing in media installation. In almost all cases, fiber-optic cabling costs considerably more than copper cabling for all components. When you need fiber-optic cabling, however, there’s really no substitute. Some people opt for a wireless network because of the cost of wired components, but wireless networks are often not the solution when there are many users requiring high bandwidth. As a network administrator, you need to factor in all costs as well as users’ needs before deciding which media types to use and in which situations. A combination of types tends to be the norm in today’s networks.

Now that you know the general characteristics of cabling as well as which characteristics influence selecting a cable type or collection of cable types, you can understand the importance of the strengths and weaknesses of cabling types discussed in the following sections.

Coaxial Cable

For many years, coaxial cable—often called “coax” for short—was the predominant form of network cabling. Inexpensive and easy to install, coaxial cable was the networker’s choice for many years, until the early 1990s. Now the main use for coaxial cable in networking is in connecting a cable modem to the wall outlet your cable TV/Internet provider installs. For this reason, details on coax cable have been relegated to Appendix B.

Twisted-Pair Cable

Twisted-pair (TP) cable comes in two types: unshielded and shielded. It consists of one or more pairs of insulated strands of copper wire twisted around one another and housed in an outer jacket or sheath (shown in Figure 4-1). These twists are important because they cause the electromagnetic fields that form around a wire carrying bit signals to wrap around one another and improve resistance to crosstalk and EMI from outside sources. In general, the more twists per unit length, the better the resistance to EMI and crosstalk. More expensive TP wire is usually more twisted than less expensive kinds and, therefore, provides a better pathway for higher bandwidth networks.

Unshielded Twisted-Pair Cable Most networks use UTP cabling, with STP used only where electrical noise is an excessive problem. The UTP cable used in LANs consists of four pairs of insulated wires; other UTP types contain fewer pairs. UTP is also used as






Figure 4-1 Twisted-pair cable

Courtesy of Course Technology/Cengage Learning

phone wire, but because voice applications are much less demanding than networking in bandwidth and signal quality, the type of cable used for phone connections is usually unsuitable as network cabling.

UTP cabling is rated according to categories devised by the Telecommunications Industry Association (TIA) and the Electronic Industries Alliance (EIA); the American National Standards Institute (ANSI) has also endorsed these standards. The ANSI/TIA/EIA 568 Commercial Building Wiring Standard defines standards for the kinds of wiring used in commercial environments and helps ensure consistent performance from wiring products. Currently, the ANSI/TIA/EIA 568 standard includes eight categories for UTP wiring; these categories also govern the number of twists per foot or meter:

- *Category 1*—Applies to traditional UTP phone cabling, which is designed to carry voice but not data. This cabling is, therefore, labeled as **voicegrade**. Most UTP installed before 1982 falls into this category. This standard is no longer recognized by TIA/EIA.
- *Category 2*—Certifies UTP cabling for bandwidth up to 4 Mbps and consists of four pairs of wire. Because 4 Mbps is slower than most current networking technologies (except for older token ring installations), Category 2 is unlikely to be seen in networking environments and is no longer recognized by TIA/EIA.
- *Category 3*—Certifies UTP cabling for bandwidth up to 10 Mbps with signaling rates up to 16 MHz. This category supports 10BaseT Ethernet and 4 Mbps token ring networks with maximum segment lengths of 100 meters. Category 3 consists of four pairs, with each pair having a minimum of three twists per foot (10 twists per meter). Cat 3 remains in use in some older networks but should be replaced when networks

are upgraded. Most networks have already migrated to 100 Mbps and 1000 Mbps speeds, and Cat 3 isn't suitable for these speeds.

- *Category 4*—Certifies UTP cabling for bandwidth up to 16 Mbps with signaling rates up to 20 MHz. This category supports primarily 10BaseT Ethernet and 16 Mbps token ring and is the first ANSI/TIA/EIA designation that labels cables as **datagrade** (meaning they're capable of carrying data) rather than voicegrade. Cat 4 consists of four twisted pairs.
- *Category 5*—Certifies UTP cabling for bandwidth up to 100 Mbps with signaling rates up to 100 MHz. This category supports 100BaseTX, Asynchronous Transfer Mode (ATM) technologies at 25 and 155 Mbps, and Copper Distributed Data Interface (CDDI) at 100 Mbps. Category 5 also consists of four twisted pairs with an average of three to four twists per inch. Cat 5 cabling has been superseded by Category 5e and can be used in Gigabit Ethernet (1000BaseT), but Cat 5e is the minimum recommendation because of the additional tests required for it. Cat 5 cable is no longer widely available.
- *Category 5e*—Category 5 enhanced UTP cabling, as the name suggests, is an enhancement to Category 5 UTP. It differs mainly in the tests it must undergo and was designed to correct some shortcomings in Cat 5 cabling, particularly in Gigabit Ethernet and full-duplex operation. Cat 5e is an acceptable cable type for 1000BaseT Ethernet, but Category 6 should be considered for new installations. Cat 5e consists of four pairs and is rated for 100 MHz signaling rates; it comes in both shielded and unshielded versions.
- *Category 6*—This standard, published in June 2002 by the TIA/EIA, is the recommended UTP cabling standard for Ethernet applications over copper media at speeds up to 1 Gbps. Category 6 cabling uses the same type of modular jack as lower categories and is backward-compatible with Category 5 and Category 5e cable plants. It's specified to operate at signaling rates of 250 MHz. Some Cat 6 cabling includes a spline, or separator, in the jacket for additional separation between pairs of wires. However, this separator isn't a requirement. Cat 6 is the preferred cabling for 1000BaseT (Gigabit Ethernet) networks, but it can also support 10GBaseT for distances under 55 meters. Cat 6 is a four-pair cable and comes in both shielded and unshielded versions.
- *Category 6a*—Published in February 2008, Category 6a (Category 6 augmented) is suitable for signaling rates up to 500 MHz and is the category specified for 10GBaseT networks with segments up to 100 meters. It comes in both shielded and unshielded versions.

Two additional categories aren't yet TIA/EIA standards and might never be in the United States. However, Europe has accepted the Category 7 and Category 7a standards, which specify a fully shielded twisted-pair cable (each wire pair is shielded, as is the outer sheath) with performance characteristics well above earlier cabling standards. Signaling rates are specified at up to 600 MHz for Cat 7 and 1000 MHz for Cat 7a. Because of a different connecting hardware design, these cables and connectors aren't likely to be backward-compatible. Cat 7 and 7a are ISO/IEC 11801 standards, and their use in the upcoming 40 and 100 Gigabit Ethernet standards is uncertain. These two categories of cable might have a short life if they aren't specified in a widely adopted networking standard.





There's some mention of Category 8 cabling, operating at 1200 MHz, that would come in only shielded versions, but there's little information about it as of this writing. A company called Hyperline is selling a product called Category 8 cable, but no standard has been published as of this writing.

Categories 5, 5e, and 6 are by far the most popular types of UTP cabling. Their huge installed base guarantees that developers of new high-speed networking technologies will strive to make their technologies compatible with these categories; for example, Category 5 cable, originally designed for 10 Mbps Ethernet, is capable (although not recommended) of running at speeds up to 1 Gbps. Table 4-1 summarizes the characteristics of the two most common UTP cabling types.

Table 4-1 Category 5e and 6 UTP cabling characteristics

Characteristic	Value
Maximum cable length	100 m (328 ft.)
Bandwidth	Up to 1000 Mbps
Bend radius	Minimum four times the cable diameter or 1 inch
Installation and maintenance	Easy to install, no need to reroute; the most flexible
Cost	Least expensive of all cabling options
Connector type	RJ-45 plug, RJ-45 jack, and patch panels
Security	Moderately susceptible to eavesdropping
Signaling rates	100 MHz for Cat 5e; 250 MHz for Cat 6
Interference rating	Susceptible to EMI and crosstalk

Shielded Twisted-Pair Cable As its name indicates, STP includes shielding to reduce crosstalk and limit the effects of external interference. For most STP cables, this means the wiring includes a wire braid inside the cladding or sheath material as well as a foil wrap around each wire pair. This shielding improves the cable's transmission speed and resistance to interference, which allows using STP in electrically noisy environments or very high-bandwidth applications. Unfortunately, no standards for STP correspond to the ANSI/TIA/EIA 568 Standard for UTP, but you can readily find STP versions of Cat 5e (shown in Figure 4-2), Cat 6, and Cat 6a. These STP versions are sometimes referred to as foiled twisted pair (FTP), and the shielding surrounds all four wire pairs rather than each wire pair.



Figure 4-2 Cat 5e shielded twisted pair

Courtesy of Course Technology/Cengage Learning

Twisted-Pair Cable Plant Components A twisted-pair cable plant requires more than just the cabling, which is usually sold in spools of 1000 feet. In addition, you find most of the following components:

- *RJ-45 connectors*—Whether STP or UTP, most twisted-pair cabling uses registered jack 45 (RJ-45) connectors to plug into network interfaces or other networked devices. This connector looks much like the RJ-11 connector on modular phone jacks, but it's larger and contains eight wire traces rather than the four or six in an RJ-11. An RJ-45 connector (see Figure 4-3), often called an **RJ-45 plug**, is most commonly used in patch cables, which are used to connect computers to hubs and switches and computers to RJ-45 wall jacks.



Figure 4-3 An RJ-45 plug

Courtesy of Hyperline Systems

- *Patch cable*—A **patch cable** (see Figure 4-4) is a short cable for connecting a computer to an RJ-45 jack or connecting a patch-panel port to a switch or hub. Patch cables can be made with inexpensive tools, two RJ-45 plugs, and a length of TP cable, which you do later in Hands-On Project 4-1. Although making a patch cable is easy, most network administrators prefer buying ready-made cables to save time.



Figure 4-4 A patch cable

© spilman/Shutterstock.com



- *RJ-45 jacks*—An **RJ-45 jack** (shown in Figure 4-5) is what you plug an RJ-45 connector into when the computer is in a work area away from hubs and switches. It has a receptacle for an RJ-45 plug on one side and a place to terminate, or “punch down,” the TP cabling on the other side. RJ-45 jacks are usually placed behind wall plates when cables are run inside walls but can also be recessed into the floor or placed in surface-mounted boxes if the cabling runs on the outside of walls.



Figure 4-5 An RJ-45 jack

Courtesy of Hyperline Systems

- *Patch panels*—Patch panels are used to terminate long runs of cable from the work area (where the computers are) to the wiring closet (where the switches and hubs are). Patch panels are like RJ-45 jacks, in that they have a receptacle on one end and punchdown terminals on the other, but a patch panel can usually accommodate 12, 24, or 48 cables. Figure 4-6 shows the front side of a patch panel, where a patch cable plugs in, and the back side, where long runs of cable are terminated.




Figure 4-6 Patch panel front and back

Courtesy of Hyperline Systems

- *Distribution racks*—Distribution racks (also called 19-inch racks because the upright rails are 19 inches apart) hold network equipment, such as routers and switches, plus patch panels and rack-mounted servers. They’re usually found in wiring closets and equipment rooms. Figure 4-7 shows a typical distribution rack.

The following sections explain how to use these components to construct a cable plant, using a method called structured cabling.

Structured Cabling: Managing and Installing a UTP Cable Plant

Entire books are written on cable installation and management, and the details are beyond the scope of this book. However, understanding some basic methods and terminology of cable installation and management will give you a good foundation. As mentioned, the TIA/EIA




Figure 4-7 A distribution rack

Courtesy of Course Technology/Cengage Learning

developed the document “568 Commercial Building Wiring Standard,” which specifies how network media should be installed to maximize performance and efficiency. This standard defines what’s often referred to as “structured cabling.”



The 568 Commercial Building Wiring Standard covers all media types, but the discussion in this section focuses on UTP cabling, the most common media for LANs and internetworks.

Structured cabling specifies how cabling should be organized, regardless of the media type or network architecture. Although a variety of logical topologies can be used, structured cabling relies on an extended star physical topology. TIA/EIA 568 can be applied to any size network and divides the details of a cable plant into six components. A small LAN in a 10-computer business might need only two or three of these components, but large networks typically use most or all of these components:

- Work area
- Horizontal wiring
- Telecommunications closets
- Equipment rooms
- Backbone or vertical wiring
- Entrance facilities

Network cabling standards are designed to ensure that standards for equipment rooms and wiring closets, including limitations on media, are adhered to, which helps limit the possible reasons for network failure or poor performance. If the network cable plant is in good working order and meets standards, a network administrator's job is easier. Structured cabling facilitates troubleshooting as well as network upgrades and expansion.

Work Area The **work area**, as the name suggests, is where workstations and other user devices are located—in short, the place where people work. Faceplates and wall jacks are installed in the work area, and patch cables connect computers and printers to wall jacks, which are connected to a nearby telecommunications closet. Patch cables in the work area should be limited to less than 6 meters long (about 20 feet). The TIA/EIA 568 standard calls for at least one voice and one data outlet on each faceplate in each work area. The connection between wall jack and telecommunications closet is made with horizontal wiring. Figure 4-8 shows the components of the work area.




Figure 4-8 Work area components

Courtesy of Course Technology/Cengage Learning

Horizontal Wiring **Horizontal wiring** runs from the work area's wall jack to the telecommunications closet and is usually terminated at a patch panel. Acceptable horizontal wiring types include four-pair Cat 5e or Cat 6 or two fiber-optic cables. The total maximum distance for horizontal wiring is up to 100 meters, which includes the cable running from the wall jack to the patch panel plus all patch cables. However, horizontal wiring from the wall jack to the patch panel should be no longer than 90 meters to allow up to 10 meters for patch cables.

Telecommunications Closet The **telecommunications closet (TC)** provides connectivity to computer equipment in the nearby work area. In small installations, it can also serve as the entrance facility (explained later in "Entrance Facilities"). Typical equipment includes patch panels to terminate horizontal wiring runs, hubs and switches to provide network connectivity, and patch cables to connect patch panels to hubs and switches. In smaller installations, network servers can be housed in the TC. Larger installations usually have connections from the TC to an equipment room (discussed next). Figure 4-9 shows the relationship and connections between the work area, horizontal wiring, and TC.




Figure 4-9 Work area, horizontal wiring, and telecommunications closet

Courtesy of Course Technology/Cengage Learning

Equipment Rooms The **equipment room** houses servers, routers, switches, and other major network equipment and serves as a connection point for backbone cabling running between TCs. An equipment room can be the main cross-connect of backbone cabling for the entire network, or it might serve as the connecting point for backbone cabling between buildings. In multi-building installations, each building often has its own equipment room.

Backbone Cabling **Backbone cabling** (or vertical cabling) interconnects TCs and equipment rooms. This cabling runs between floors or wings of a building and between buildings to carry network traffic destined for devices outside the work area. It's often fiber-optic cable but can also be UTP if the distance between rooms is less than 90 meters. When it connects buildings, backbone cabling is almost always fiber optic because of UTP's distance limitations and because fiber doesn't propagate lightning strikes or electrical imbalances between buildings. Multimode fiber-optic cable can extend up to 2000 meters, whereas single-mode fiber can reach distances up to 3000 meters when used as backbone cabling between the main cross-connect and TCs. Between equipment rooms and TCs, the distance is limited to 500 meters for both fiber-optic cable types; from the main cross-connect to equipment rooms, fiber-optic cable can run up to 1500 meters. Figure 4-10 shows how backbone cabling can connect TCs to an equipment room.

Entrance Facilities An **entrance facility** is the location of the cabling and equipment that connects a corporate network to a third-party telecommunications provider. It can also serve as an equipment room and the main cross-connect for all backbone cabling. This




Figure 4-10 Backbone cabling connects TCs and equipment rooms

Courtesy of Course Technology/Cengage Learning

location is also where a connection to a WAN is made and the point where corporate LAN equipment ends and a third-party provider’s equipment and cabling begins—also known as the “demarcation point.”

Installing UTP Cabling One skill required of a network technician is terminating UTP cables. Cable **termination** means putting RJ-45 plugs on the cable to make a patch cable or punching down cable wires into terminal blocks on a jack or patch panel. To create a patch panel, a technician needs the following tools:

- Bulk UTP cabling
- Wire cutters or electrician’s scissors
- Cable stripper
- Crimping tool
- Cable tester
- RJ-45 plugs

To terminate cable at an RJ-45 jack or a patch panel, you need the following tools:

- Bulk UTP cabling
- Wire cutters or electrician's scissors
- Cable stripper
- Type 110 punchdown tool
- Cable tester
- RJ-45 jack and patch panel

Some tools you need to perform these tasks are shown in Figure 4-11.





Figure 4-11 Cable installation and termination tools

Courtesy of Course Technology/Cengage Learning

The quality of the tools needed for cable installation varies considerably, usually according to cost. If you expect to be doing a lot of cable termination, it pays to invest in high-quality tools, particularly a cable tester. If you're installing only a few dozen to a few hundred cables, you might get away with less expensive tools and a basic cable tester. However, if you have a cable-installation business, you want high-quality tools, including a cable tester that certifies the cable plant for the category of cable installed.


Hands-On Project 4-1 walks you through making a patch cable. One of the most important aspects of making a cable or terminating a cable at a jack or patch panel is to get the colored wires arranged in the correct order. There are two competing standards for the arrangement of wires: TIA/EIA 568A and TIA/EIA 568B. Either standard is okay to follow, as long as you stick to one throughout your network. The arrangement of wires for both standards is shown in Figures 4-12 and 4-13.



Pin #	Color	Pair #	Function
1	White with green stripe	3	Transmit +
2	Green	3	Transmit -
3	White with orange stripe	2	Receive +
4	Blue	1	Unused
5	White with blue stripe	1	Unused
6	Orange	2	Receive -
7	White with brown stripe	4	Unused
8	Brown	4	Unused

Figure 4-12 TIA/EIA 568A cable pinouts

Courtesy of Course Technology/Cengage Learning



Pin #	Color	Pair #	Function
1	White with orange stripe	2	Transmit +
2	Orange	2	Transmit -
3	White with green stripe	3	Receive +
4	Blue	1	Unused
5	White with blue stripe	1	Unused
6	Green	3	Receive -
7	White with brown stripe	4	Unused
8	Brown	4	Unused

Figure 4-13 TIA/EIA 568B cable pinouts

Courtesy of Course Technology/Cengage Learning

Straight-Through Versus Crossover Cable When you make a standard patch cable, you use the same wiring standards on both ends of the cable so that each wire is in the same corresponding location on both ends of the cable (pin 1 goes to pin 1, pin 2 to pin 2, and so forth). This type of cable is also called a **straight-through cable**. Another type of cable, called a **crossover cable**, uses the 56B standard on one end and the 56A standard on the other end. This arrangement crosses the transmit and receive wires so that transmit on one end connects to receive on the other end. This type of cable is often needed when you connect two devices of the same type to one another—for example, connecting a hub to a hub, a switch to a switch, a hub to a switch, or a PC to a PC.

Medium Dependent Interface Network devices connecting with RJ-45 plugs over twisted-pair cabling are classified as **medium dependent interface (MDI) devices** or **MDI crossed (MDI-X) devices**. You might even see these abbreviations on some hubs and switches. For communication to take place between two devices, the wires one device transmits on must be connected to the wires the other device receives on, and vice versa. For example, the 568 standards have pins 1 and 2 labeled as transmit and pins 3 and 6 labeled as receive. Clearly, not all devices can transmit on pins 1 and 2 and receive on pins 3 and 6; otherwise a standard patch cable wouldn't work between these devices because one device's transmit signals would be going to the transmitter of the other device—like having a phone's earpiece at your mouth and the mouthpiece at your ear.

MDI devices transmit on pins 1 and 2 and receive on pins 3 and 6. Examples include PC NICs and routers. MDI-X devices, usually hubs and switches, receive on pins 1 and 2 and transmit on pins 3 and 6. Therefore, a straight-through patch cable works for the most common connection of a PC NIC to a switch or hub. When a switch needs to be connected to a switch (or a PC to a PC), you use a crossover cable so that the transmit and receive wires get crossed, and you end up with transmit going to receive and vice versa. Thankfully, developers of NICs, switches, routers, and even some hubs have started doing this job for you by making ports on some devices auto-sensing. Auto-sensing means they can detect whether you're trying to connect transmit wires to transmit wires, and the port reconfigures its transmit and receive wires, thus making a crossover cable unnecessary. Not all devices support auto-sensing, so it's best to have crossover cables handy in case you need them.



Hands-On Project 4-1: Making a Patch Cable

Time Required: 20 minutes

Objective: Create a 568B straight-through patch cable.

Required Tools/Equipment: Wire cutter and cable stripper, RJ-45 crimping tool, 2 to 4 feet of Cat 5/5e or Cat 6 cable, two RJ-45 plugs, cable tester (optional)

Description: In this project, you make a patch cable according to the instructions. The instructor will inspect the cable for the correct wire order and strain relief. If possible, a cable tester is used to test for conductivity and wiremap, at a minimum.

1. Strip approximately 2 inches of the outer jacket off one end of the cable with the cable stripper. Be careful not to nick the inner wires' insulation. Most UTP cable strippers are calibrated to score the cable's outer jacket so that you can simply break it off. Cable



strippers differ in the techniques you use with them, so refer to the instructions that came with yours or ask your instructor.

2. Untwist the four pairs of wires.
3. Here comes the tricky part: Arrange the wires from left to right (as you're looking down on them) so that they're in the following order: white with orange stripes, orange, white with green stripes, blue, white with blue stripes, green, white with brown stripes, and brown. This order adheres to the 568B wiring standard (see Figure 4-14).




Figure 4-14 The correct arrangement of wires

Courtesy of Course Technology/Cengage Learning

4. Clip the eight wires so that a little more than a half-inch of wire extends beyond the outer jacket.
5. While holding the RJ-45 plug in one hand with the clip facing away from you, insert the eight wires into the connector, making sure the tops of wires extend to the front of the connector and the cable jacket goes far enough into the connector so that the jacket will be caught by the crimp bar (see Figure 4-15).




Figure 4-15 Correct RJ-45 plug installation

Courtesy of Course Technology/Cengage Learning

6. Now insert the RJ-45 connector into the crimping tool, and make sure the wires don't slip. Squeeze the handle on the crimping tool firmly. It might take a little hand strength or using two hands, depending on the crimping tool's quality. This tool does two things. First, it forces the eight small contacts at the top of the plug down onto the wires; the contacts are pushed just far enough in that they slice through the insulation on each wire, thereby making an electrical contact with the wire. Second, the strain-relief bar is pushed in to grab the cable's outer jacket, making it more difficult to pull the wires out of the plug.
7. Repeat the process for the other end of the cable, and test with a cable tester, if available. Congratulations! You have made a patch cable.
8. If time and materials allow, make a crossover cable, using the 568B wiring scheme on one end and the 568A wiring scheme on the other. Keep your tools handy for the next project.



Hands-On Project 4-2: Terminating UTP Cable

Time Required: 20 minutes

Objective: Terminate UTP cable at a patch panel and an RJ-45 jack.

Required Tools/Equipment: Wire cutter and cable stripper, 2 to 4 feet of Cat 5/5e or Cat 6 cable, 110 punchdown tool, Cat 5/5e or Cat 6 patch panel (a 568A or 568B patch-panel can be used; 568B panels are more common), RJ-45 jack, cable tester (optional)

Description: In this project, you punch down one end of a cable to the back of a patch panel.

1. Strip approximately 2 inches of the outer jacket off one end of the cable with the cable stripper. Be careful not to nick the inner wires.
2. Leave the wire pairs twisted. Arrange the wires according to the color coding on your patch panel. The color-coding will vary, depending on whether it's a 568A or 568B patch panel, and the wires might be arranged in a straight line or split between the two rows of terminals.
3. Center the cable so that each wire is equally distant from the terminal in which it will be placed. On each wire pair, separate the wires about one-half inch or less from the end of the jacket so that the two wires form an oval, and slip the pair over the middle terminal



for that wire pair (see Figure 4-16). Pull each wire pair down firmly so that the wires stay in place.




Figure 4-16 Placing wires on the patch panel terminals

Courtesy of Course Technology/Cengage Learning

4. Next, use the 110 punchdown tool. For each wire, place the tool over the wire so that the slot in the tool lines up with the wire. The tool's blade should be facing the end of the wires, not the cable jacket (see Figure 4-17).




Figure 4-17 Positioning the punchdown tool

Courtesy of Course Technology/Cengage Learning

5. Push the punchdown tool down firmly until you hear it snap. Don't be afraid to give it a good hard push. The blade should cut the wire or at least score it so that you can gently twist the end off. Do this for all eight wires.
6. That's it! A correct termination should have no more than one-half inch of untwisted wire outside the jacket. Repeat this process for the other end of the cable, but this time, terminate the cable onto an RJ-45 jack. Keep your cables and tools ready for the next project.



Hands-On Project 4-3: Conducting End-to-End Testing

Time Required: 10 minutes

Objective: Test your terminations and patch cable with a live connection.

Required Tools/Equipment: The patch cable you made, an additional patch cable, the patch panel and RJ-45 jack to which you terminated the cable, a PC, and a hub or switch

Description: Working in groups of at least two, connect a PC, using the patch cable you made, to the RJ-45 jack you punched down. Using an additional patch cable, connect the patch panel to a hub or switch. Then you use the `ping` command to verify connectivity between computers.

1. Using the patch cable you made in Hands-On Project 4-1, connect your PC's NIC to the RJ-45 jack you punched down in Hands-On Project 4-2.
2. Using the additional patch cable, connect the port on the patch panel you punched down to a hub or switch.
3. Turn on the PC and the hub or switch, if necessary.
4. Verify that you have a link light at the hub or switch and at your PC's NIC. Log on to your PC and give your PC's IP address to another student who's connected to the hub or switch.
5. Ping another student's computer after getting his or her IP address. If the ping is successful, your cable termination was a success.
6. If you're sharing computers, allow the next group of students to test their cabling.
7. Shut down your computer if no one else is using it for testing.

Why Two Transmit and Two Receive Wires? As you can see from the cable pinout diagrams in Figures 4-12 and 4-13, one wire pair is used for transmit (labeled transmit+/transmit-) and one wire pair is used for receive (labeled receive+/receive-). The plus and minus symbols indicate that the wires carry a positive or negative signal. This **differential signal** mitigates the effects of crosstalk and noise on the cable. It does so because when a bit signal is transmitted, it's transmitted as a positive voltage and a negative voltage (v). For example, if a 1 bit is defined as $+2v$, the bit is transmitted as $+2v$ and $-2v$. The receiver reads the difference between the two values, which is $4v$. EMI and crosstalk manifest as positive voltages, so what happens if the signal is hit by a burst of EMI that adds $1v$ to the signal? You have the following:



Original signal with no EMI:

Transmit+	Transmit-	Differential result
+2v	-2v	+4v

Signal with EMI adding 1v to both transmit+ and transmit- wires:

Transmit+	Transmit-	Differential result
+2v + 1v = 3v	-2v + 1v = -1v	+4v

As you can see, in both cases, the result stays at +4v because the differential signal effectively cancels out the EMI. However, this canceling effect works only if the same amount of EMI is imposed on both wires. The closer the wires are, the more likely that EMI will affect both wires equally. This phenomenon is one reason for using twisted wires: The wires are so tightly coupled that both external EMI and crosstalk are likely to affect both wires equally and be canceled out.

Although UTP is the most common media type for LANs, it has its limitations in bandwidth, noise susceptibility, and length. In addition, UTP wiring shouldn't be used outside to connect between buildings. Copper wire is susceptible to the elements, and its electrical conducting properties change slightly depending on the temperature. A more important reason to not use any type of copper wire between buildings is that it can carry a harmful electrical charge based on the ground potential between buildings, if the buildings are fed from different transformers. When any of these limitations eliminate UTP as an option, fiber-optic cable, discussed in the next section, is the likely solution.

Fiber-Optic Cable

Fiber-optic cable trades electrical pulses for pulses of light to represent bits. Because no electrical signals ever pass through the cable, fiber-optic media is as immune to electrical interference as any medium can get. Therefore, light pulses are unaffected by EMI and RFI. This characteristic also makes fiber-optic cables highly secure. They emit no external signals that might be detected, unlike electrical or broadcast media, thereby eliminating the possibility of electronic eavesdropping. In particular, fiber-optic cable is a good medium for high-bandwidth, high-speed, long-distance data transmission because of its lower attenuation characteristics and vastly higher bandwidth potential. Today, commercial implementations at 10, 40, and 100 Gbps are in use.

Figure 4-18 shows a typical fiber-optic cable. A slender cylinder of glass fiber called the core is surrounded by a concentric layer of glass called the cladding. The fiber is then jacketed in a thin transparent plastic material called the buffer. These three components make up what's labeled as the optical fiber in Figure 4-18. The fiber is optionally surrounded by an inner sheath made of colored plastic. A strengthening material, usually made of Kevlar, comes next, followed by an outer sheath. Sometimes the core consists of plastic rather than glass fibers; plastic is more flexible and less sensitive to damage than glass, but attenuation is more of a problem than with glass.




Figure 4-18 Fiber-optic cable

Courtesy of Course Technology/Cengage Learning

Each fiber-optic strand carries data in only one direction, meaning fiber-optic network connections consist of two or more strands, each in a separate inner sheath. However, these cables can be enclosed in a single sheath or can be two separate cables, each with its own sheath. Just as you have UTP patch cables, you also find fiber-optic patch cables, usually to connect from a fiber-optic patch panel to a switch or router. Fiber-optic cable used as backbone cabling often comes in bundles of 12 or more fiber strands. Even if you're using only two strands at first, it's a good idea to run cable containing more fiber than you need, in case a strand breaks during installation or you need additional strands for future growth.

Some testing has shown that glass fibers can carry several terabits (1000 gigabits) per second (Tbps). There's really no end in sight for the bandwidth capacity of optical fiber. As network bandwidth needs increase and the limits of copper wire are reached, fiber-optic cable will probably replace copper for all types of network connections. Table 4-2 summarizes fiber-optic cable characteristics.

Table 4-2 Fiber-optic cable characteristics

Characteristic	Value
Maximum cable length	2 km (6562 ft.) to 100 km (62.14 miles)
Bandwidth	10, 40, and 100 Gbps and higher
Bend radius	30 degrees per foot
Installation and maintenance	Difficult to install and reroute; sensitive to strain and bending
Cost	Most expensive of all cabling options
Connector type	Several types (see bulleted list in the next section)
Security	Not susceptible to eavesdropping
Interference rating	None; least susceptible of all cable types

Fiber-Optic Connectors

A wide variety of connectors can be used with fiber-optic media, depending on the light-emitting sources used to generate light pulses and the corresponding light-detecting sensors used to detect them. Figure 4-19 shows some connectors described in the following list:




Figure 4-19 Fiber-optic connectors

Courtesy of Course Technology/Cengage Learning

- *Straight tip*—Straight tip (ST) connectors join fibers at cross-connections or to optical devices. They're used most often in Ethernet networks with fiber-optic cable as backbone cabling. An ST connector locks onto the jack when twisted.
- *Straight connection*—Straight connection (SC) connectors push on, which makes them easy to install and requires less space for an attachment. They make a strong connection and can be used when splicing fiber-optic cables. An SC connector is a one-piece component, with two receptacles for sending and receiving fibers. A notch in its jacket ensures the correct orientation when inserted.
- *Locking connection*—Locking connection (LC) connectors push on and pull off with an RJ-45-style latching mechanism. They're about half the size of SC connectors, which makes them good for high-density applications, in which many fibers are concentrated in one location.
- *Medium interface connector*—A medium interface connector (MIC) is used for Fiber Distributed Data Interface (FDDI). Like SC connectors, MIC connectors are one-piece constructions.
- *Subminiature type A*—The company Amphenol originally designed subminiature type A (SMA) connectors for microwave use and later modified them for fiber-optic use. Two SMA versions are widely available: The 905 uses a straight ferrule, which is a metal sleeve for strengthening the connector, and the 906 uses a stepped ferrule with a plastic sleeve to ensure precise alignment of fibers. Like ST connectors, SMAs use two connectors for each fiber strand.

- *Mechanical transfer registered jack*—A mechanical transfer registered jack (MT-RJ) connector looks a little like an RJ-45 connector. It provides a high-density fiber-optic connection by using two fiber-optic cables. Compared with other connector types, MT-RJ connectors take only half the space for the same number of cable terminations. They're also easy to install and require only one connector for a two-fiber termination.

Fiber-Optic Installation

Installing fiber-optic networks is somewhat more difficult and time consuming than copper media installation. However, advances in connector technology have made field termination of fiber-optic cables almost as fast and easy as copper terminations. The connectors and test equipment required for termination are still considerably more expensive than their copper counterparts, but the trend toward easier, more affordable fiber-optic networks continues. Fiber-optic cable to the desktop, although not common, is becoming a feasible option for more companies.

There are several methods for terminating fiber-optic cables because of the many connectors and cable types available, so installation details are beyond the scope of this book. Before embarking on a fiber-optic termination task, you need to purchase a fiber-optic termination kit, which can range from several hundred to several thousand dollars. Some tools in a typical fiber-optic termination kit include the following:

- *Buffer tube stripper*—A tightly calibrated tool designed for stripping buffer tubes off the glass fiber strand without breaking the fiber
- *Cable stripper*—Used to remove the fiber cable's outer sheath; much like the cable stripper used with UTP
- *Crimper*—Used with connectors that use crimping as the method to fix the connector to the cable
- *Diamond cleaver*—Used to cut glass fiber cleanly without shattering the end
- *Inspection scope*—Used for examining the end of a fiber strand to make sure it's clean and polished
- *Polishing tool*—Used to polish the end of a cleaved (cut) strand of fiber

Fiber-Optic Cable Types

Fiber-optic cables come in two main types: single-mode fiber (SMF) cables, which include a single extremely small-diameter fiber (typically 8 microns) at the core, and multimode fiber (MMF) cables, which use a considerably larger diameter fiber (50 and 62.5 microns are standard sizes) at the core. SMF cable costs more and generally works with laser-based emitters but spans the longest distances and is used in higher-bandwidth applications. MMF cables cost less and work with lower-power light emitting diodes (LEDs), which span shorter distances.

In the past, fiber-optic cable's high cost and difficult installation meant it was used only when a network required extremely high bandwidth or needed to span long distances between wired network segments. However, because of the falling costs of fiber and its advantages in immunity to interference, high-bandwidth capability, and increased security, fiber-optic cable is now used almost exclusively for all network backbone connections. It's also the medium of



choice for long-haul telecommunications, in which large amounts of voice and data traffic are aggregated, such as between telecommunication providers and ISPs.

Wireless Networking

Wireless technologies are playing a bigger role in all kinds of networks. Since 1990, wireless options have increased, and the cost of these technologies continues to decrease. As wireless networking has become more affordable, demand has increased, and as demand increases, so does production of wireless equipment, which brings prices down even more. For this reason, wireless networks can now be found in most towns and cities in the form of hotspots, and many home users have turned to wireless networks so that their computers are no longer tethered to a network cable.

The adjective “wireless” might lead you to believe that wireless networks have no cabling of any kind. However, wireless networks are often used with wired networks to interconnect geographically dispersed LANs or groups of mobile users with wired servers and resources on a wired LAN. Networks including both wired and wireless components are called “hybrid networks.” Indeed, even in home or small business networks with workstations connecting to a wireless AP or router, the AP or router usually connects to the Internet via a wired connection to a cable modem or similar device. Probably the only truly wireless networks are ad hoc networks or small infrastructure networks put together for the purpose of sharing files among a small group of people.

Wireless Benefits

Wireless networking has a lot of appeal in many circumstances and can offer the following capabilities:

- Create temporary connections to existing wired networks.
- Establish backup or contingency connectivity for existing wired networks.
- Extend a network’s span beyond the reach of wire-based or fiber-optic cabling, especially in older buildings where rewiring might be too expensive.
- Allow businesses to provide customers with wireless networking easily, thereby offering a service that gets customers in and keeps them there.
- Enable users to roam around a corporate or college campus with their machines.

Each capability supports uses that extend the benefits of networking beyond conventional limits. Although wireless networking is invariably more expensive than cable-based alternatives, sometimes these benefits can more than offset the extra costs. Today, common applications for wireless networking technologies include the following:

- Ready access to data for mobile professionals, such as doctors or nurses in hospitals or delivery personnel. For instance, United Parcel Service (UPS) drivers maintain connections to a server at the home office; their handheld computers send and receive delivery updates and status information to a network server over a wireless phone connection. Doctors can carry lightweight mobile devices on their patient visits and have wireless access to patient information at all times.

- Delivering network access to isolated facilities or disaster-stricken areas. For example, the Federal Emergency Management Agency (FEMA) uses battery-powered wireless technologies to install field networks in areas where power and connections might be unavailable.
- Access in environments where layout and settings change constantly. For instance, film studios often include wireless network components on the set so that information is always available, no matter how the stage configuration changes.
- Improved customer services in busy areas, such as check-in or reception centers. For example, Hertz employees use handheld units to check in returned rental vehicles right in the parking lot.
- Network connectivity in structures, such as historical buildings, where in-wall wiring is impossible to install or prohibitively expensive.
- Home networks where installing cables is inconvenient. More people who own multiple computers are installing inexpensive wireless networks so that family members can share Internet connections and files. Figure 4-20 shows an example of using wireless in a home network.




Figure 4-20 A typical home wireless network

Courtesy of Course Technology/Cengage Learning

Types of Wireless Networks

Depending on the role wireless components play in a network, wireless networks can be subdivided into the following categories:

- *Local area networks (LANs)*—In LANs, wireless components act as part of an ordinary LAN, usually to provide connectivity for mobile users or in changing environments or perhaps across areas that couldn't otherwise be networked. Examples include older buildings where installing wiring is impractical or areas that encompass public or common property where cabling might not be permitted.



- *Extended LANs*—In **extended LANs**, an organization might use wireless components to increase a LAN's span beyond normal distance limitations for wire-based or fiber-optic cables.
- *Internet service*—A company that wants to be a high-speed ISP but doesn't have a media infrastructure available, as cable and phone companies do, can use wireless technologies to bring Internet access to homes and businesses.
- *Mobile computing*—With mobile computing, users communicate by using a wireless networking medium, such as radio or cell phone frequencies, that enable them to move while remaining connected to a network.

Wireless LAN Components

The wireless components of most LANs behave like their wired counterparts, except for the media and related hardware. The operational principles are much the same: Attaching a network interface of some kind to a computer is still necessary, but the interface attaches to an antenna and an emitter rather than to a cable. Users can still access the network just as though cable connects them to it.

Another component is required to link wireless users with wired users or resources. At some point on a cabled network, a transmitter/receiver device, called a **transceiver** or an access point (AP), must be installed to translate between wired and wireless networks. This device broadcasts messages in wireless format that must be directed to wireless users and relays messages sent by wireless users to resources or users on the wired side of its connection. An AP includes an antenna and a transmitter to send and receive wireless traffic but also connects to the wired side of the network. This connection enables the device to shuttle traffic back and forth between a network's wired and wireless sides.

Wireless LAN Transmission

All wireless communication depends on sending and receiving signals broadcast through the air to carry information between network devices. These signals take the form of waves in the electromagnetic (EM) spectrum. The frequency of the wave forms used for communication is measured in cycles per second, usually expressed as **hertz (Hz)**. The entire EM spectrum starts with low-frequency waves, such as those used for electrical power (60 Hz in the United States) and telephone (0 to 3 kilohertz [KHz] for traditional voice systems) and goes all the way through the visible light frequencies to the highest frequencies in existence, at which gamma rays and other high-energy particles operate.

In wireless communication, frequency affects the amount and speed of data transmission. The transmission's strength or power determines the distance that broadcast data can travel and still remain intelligible. In general, however, the principles governing wireless transmissions dictate that lower-frequency transmissions can carry less data more slowly over longer distances, and higher-frequency transmissions can carry more data faster over shorter distances.

The middle part of the EM spectrum is commonly divided into several named frequency ranges (bands). The following are the most common frequencies for wireless data communication:

- *Radio*—10 KHz to 1 GHz (gigahertz)
- *Microwave*—1 GHz to 500 GHz
- *Infrared*—500 GHz to 1 THz (terahertz)

The important principles to remember about a broadcast medium are the inverse relationship between frequency and distance and the direct relationship between frequency and data transfer rate and bandwidth. It's also important to understand that higher-frequency technologies often use tight-beam broadcasts and require a clear line of sight between sender and receiver to ensure correct delivery.

Wireless LANs make use of four main technologies for transmitting and receiving data, discussed in the following sections:

- Infrared
- Laser
- Narrowband (single-frequency) radio
- Spread-spectrum radio

Infrared LAN Technologies Infrared (IR) wireless networks use infrared light beams to send signals between pairs of devices. These devices typically generate signals strong enough to prevent interference from light sources in most office environments. Infrared works well for LAN applications because of its high bandwidth, which makes 10 to 100 Mbps transmission rates easy to deliver. The four main kinds of infrared LANs include the following:

- Line-of-sight networks require an unobstructed view, or a clear line of sight, between the transmitter and receiver.
- Reflective wireless networks broadcast signals from optical transceivers near devices to a central hub, which then forwards signals to their intended recipients.
- Scatter infrared networks bounce transmissions off walls and ceilings to deliver signals from sender to receiver. TV remotes work in this fashion. This approach limits maximum reception distances to approximately 30 meters (100 feet). Because bounce technologies introduce signal delays, scatter infrared results in lower bandwidth than line of sight.
- Broadband optical telepoint networks provide broadband services. This technology offers high speed and wide bandwidth, can handle high-end multimedia traffic, and matches the capabilities of most wired networks.

IR transmissions are sometimes used for virtual docking connections that enable portable computing devices to communicate with wired computers or peripheral devices, such as printers. Even though infrared offers reasonable networking speeds and convenience, infrared LANs are hampered by the typical 100-foot distance limitation. Because infrared light is close in frequency to visible light (and most visible light sources emit strongly in infrared frequencies), infrared is prone to interference problems from fluorescent and other light sources in most work environments. These devices are often called **IrDA devices**, named after the Infrared Device Association, a trade association for designers and manufacturers of infrared equipment.

Laser-Based LAN Technologies Laser-based transmissions also require a clear line of sight between sender and receiver. Any solid object or person blocking a beam interrupts data transmissions. To protect people from injury and excess radiation, laser-based LAN devices are subject to many of the same limitations as infrared but aren't as susceptible to interference from visible light sources.



Narrowband Radio LAN Technologies Narrowband radio (also called “single-frequency radio”) LANs use low-powered, two-way radio communication, much like what’s used in taxis, police radios, and other private radio systems. Receiver and transmitter must be tuned to the same frequency to handle incoming and outgoing data. Unlike light-based communications, such as infrared or laser, narrowband radio requires no line of sight between sender and receiver, as long as both parties stay within the broadcast range of these devices—typically, a maximum range of approximately 70 meters (230 feet).

In the United States, government agencies, such as the Federal Communications Commission (FCC), regulate nearly all radio frequencies. Organizations that want frequencies for their exclusive use in specific locales must complete a time-consuming, expensive application process before being granted the right to use them. Because of the difficulty in securing exclusive use, the FCC sets aside certain frequencies for unregulated use, such as the ones at which cell phones and remote-control toys operate. As wireless networking and other forms of wireless communication become more popular, crowding of these frequencies could become a problem.

Depending on the frequency, walls or other solid barriers can block signals and prevent transmission and reception. Interference from other radio sources is also possible, particularly if the devices broadcast in the unregulated frequency ranges, as most wireless LAN technologies do. As with any broadcast technology, anyone within range of the network devices could eavesdrop on network communication. For narrowband radio technologies, this range is quite short. Table 4-3 summarizes the characteristics of narrowband wireless LAN technologies.

Table 4-3 Narrowband wireless LAN characteristics

Characteristic	Value
Frequency ranges	Unregulated: 902–928 MHz, 2.4 GHz, 5.72–5.85 GHz
Maximum distance	50–70 m (164–230 ft.)
Bandwidth	1–10 Mbps
Installation and maintenance	Easy to install and maintain
Interference	Highly susceptible
Cost	Moderate
Security	Highly susceptible to eavesdropping within range

Other single-frequency LAN technologies operate at higher power ratings. Networks of this type can usually transmit as far as the horizon and even farther by using repeater towers or signal-bouncing techniques. This kind of technology is well suited for communicating with mobile users but much more expensive than lower-powered alternatives. In addition, transmission equipment is more expensive and usually requires FCC licensing. Most users of this technology, even in the largest organizations, purchase this service from a communications carrier instead of operating their own facilities.

Lack of security can be a serious concern with this kind of networking technology. Anyone with the correct receiver can eavesdrop on communications, which explains why encryption of traffic is common for networks operating at these frequencies. Table 4-4 summarizes the characteristics of high-powered single-frequency radio networks.

Table 4-4 High-powered single-frequency LAN characteristics

Characteristic	Value
Frequency ranges	Unregulated: 902–928 MHz, 2.4 GHz, 5.72–5.85 GHz
Maximum distance	Line of sight, unless extension technologies are used
Bandwidth	1–10 Mbps
Installation and maintenance	Difficult, highly technical, requires licensing
Interference	Highly susceptible
Cost	Expensive to very expensive
Security	Highly susceptible to eavesdropping

Spread-Spectrum LAN Technologies Spread-spectrum radio addresses several weaknesses of single-frequency communications, whether high or low power. Instead of using a single frequency, spread-spectrum uses multiple frequencies simultaneously, thereby improving reliability and reducing susceptibility to interference. Also, using multiple frequencies makes eavesdropping more difficult.

The two main kinds of spread-spectrum communications are frequency hopping and direct-sequence modulation. Frequency hopping switches data between multiple frequencies at regular intervals. Transmitter and receiver must be tightly synchronized to maintain communication. The hardware handles the timing of hops and chooses the next frequency without sending any information about this activity, so eavesdropping is nearly impossible. Because frequency-hopping technologies use only one frequency at a time, however, their effective bandwidth is usually 1 Mbps or lower and seldom exceeds 2 Mbps.

Direct-sequence modulation breaks data into fixed-size segments called “chips” and transmits the data on several different frequencies at the same time. The receiving equipment knows what frequencies to monitor and how to reassemble the arriving chips into the correct sequences of data. It’s even possible to transmit dummy data on one or more channels, along with real data on other channels, to make it more difficult for eavesdroppers to re-create the original data. Typically, these networks operate in unregulated frequencies and provide bandwidths from 2 to 6 Mbps, depending on the number of dummy channels used. 802.11b/g/n networks using the 2.4 GHz frequency use direct sequence spread spectrum (DSSS).



Orthogonal frequency divisional multiplexing (OFDM) is a spread-spectrum technology used by the 802.11a 5 GHz standard and the 802.16 WiMAX standard.

Table 4-5 summarizes the characteristics of spread-spectrum LAN technologies.

Table 4-5 Spread-spectrum LAN characteristics

Characteristic	Value
Frequency ranges	Unregulated: 902–928 MHz or 2.4 GHz, 5 GHz
Maximum distance	Limited to cell boundaries but often extends over several miles
Bandwidth	1–2 Mbps for frequency hopping, 2–6 Mbps for direct-sequence modulation
Installation and maintenance	Depends on equipment; ranges from easy to difficult
Interference	Moderately resistant
Cost	Inexpensive to moderate
Security	Not very susceptible to eavesdropping

Wireless Extended LAN Technologies Certain kinds of wireless networking equipment extend LANs beyond their normal cable-based distance limitations or provide connectivity across areas where cables are not allowed (or able) to traverse. For instance, **wireless bridges** can connect networks up to 3 miles (4.4 km) apart. These LAN bridges permit linking locations by using line-of-sight or broadcast transmissions. They can also make it unnecessary to route dedicated digital communication lines from one site to another through a communications carrier. Normally, upfront expenses for this technology are as much as 10 times higher, but it eliminates recurring monthly service charges from a carrier. This savings can quickly make up for (and exceed) the initial expense. Spread-spectrum radio, infrared, and laser-based equipment are readily available commercially.

Longer-range wireless bridges are also available, including spread-spectrum solutions that work with Ethernet or token ring over distances up to 25 miles. As with shorter-range wireless bridges, the communication cost savings over time can justify the cost of a long-range wireless bridge. When it's connected correctly, this equipment (in long-range and short-range varieties) can transport both voice and data traffic. Table 4-6 summarizes the characteristics of wireless extended LAN technologies.

Table 4-6 Wireless extended LAN characteristics

Characteristic	Value
Frequency ranges	Spread-spectrum, infrared, laser
Maximum distance	1–3 miles for short-range, up to 25 miles for long-range
Bandwidth	1–6 Mbps for spread-spectrum, 2–100 Mbps for infrared and laser
Installation and maintenance	Depends on equipment; ranges from easy to difficult
Interference	Highly resistant
Cost	Inexpensive to moderate
Security	Not very susceptible to eavesdropping

Microwave Networking Technologies Microwave systems deliver higher transmission rates than radio-based systems do, but because the frequencies are so high, transmitters and receivers must share a common clear line of sight. Microwave communication usually requires FCC approval and licensing and is more expensive than radio systems. Experts distinguish between two types of microwave systems: terrestrial and satellite.

Terrestrial refers to line-of-sight transmissions between special microwave towers or between transmitters and receivers mounted on tall buildings, mountaintops, or other locations with long, clear lines of sight. **Terrestrial microwave** systems use tight-beam, high-frequency signals to link sender and receiver. By using relay towers, microwave systems can extend a signal across continental-scale distances.

In fact, many communications carriers use microwave towers to send traffic across sparsely populated areas where traffic is moderate and distances make laying cable expensive. The tight-beam nature of these systems means transmitters and receivers must align precisely for the best results. Some low-powered microwave systems are available for short-range LAN use, but they, too, require a clear line of sight between transmitters and receivers. Table 4-7 summarizes the characteristics of terrestrial microwave networks.

Table 4-7 Terrestrial microwave LAN/WAN characteristics

Characteristic	Value
Frequency ranges	4–6 GHz or 21–23 GHz
Maximum distance	Typically 1–50 miles
Bandwidth	1–10 Mbps
Installation and maintenance	Difficult
Interference	Varies depending on power and distance; longer distances are more prone to weather disturbances
Cost	Expensive
Security	Highly susceptible, but signals are usually encrypted

The other main alternative for microwave transmission is satellite. Instead of aiming at transmitters or receivers within a clear line of sight on the ground, **satellite microwave** systems send and receive data from geosynchronous satellites that maintain fixed positions in the sky. This is how TV signals and some long-distance phone signals travel from one side of the world to another: The sender beams the signal to a satellite visible on the horizon, the satellite relays the signal to one or more satellites until it comes onto the receiver's horizon, and then the satellite redirects the signal to the receiver.

Most organizations can't fund launching satellites, so most satellite microwave systems must lease frequencies on satellites operated by global communications carriers. Because this approach is prohibitively expensive, even multinational companies with legitimate needs to send data around the globe typically choose to pay for their communications time rather than exclusive use of their own frequency.

Even more than terrestrial microwave, satellite communication covers a broad area and can be received by anyone with the right reception equipment. That's why microwave transmissions are routinely encrypted—to make sure only their intended recipients can access their contents. Table 4-8 summarizes the characteristics of satellite microwave communications.



Table 4-8 Satellite microwave WAN characteristics

Characteristic	Value
Frequency ranges	11–14 GHz
Maximum distance	Global reach
Bandwidth	1–10 Mbps
Installation and maintenance	Prohibitively difficult
Interference	Prone to EM interference, jamming, atmospheric disturbances
Cost	Prohibitive
Security	Not very susceptible to eavesdropping

For extending the reach of a network to its ultimate dimensions, microwave technologies currently offer the broadest reach. That’s why they are labeled LAN/WAN (terrestrial) or WAN (satellite) technologies.

LAN Media Selection Criteria

In LANs and internetworks, there are three main media choices: UTP, fiber optic, and wireless. For UTP, the choices are usually Cat 5e or Cat 6 for most applications, although you might opt for a shielded version. Fiber-optic cabling is often the top choice for connecting wiring closets and buildings and possibly in electrically noisy environments and for ultra-high speed connections to servers. Wireless networks typically supplement a wired network to accommodate mobile users or are used for SOHO networks that don’t need the higher bandwidth wired networks can provide. Following is a summary of criteria to explore when you’re having difficulty choosing between media types:

- *Bandwidth*—How fast must the network be? Higher bandwidth means more expensive cable and higher installation costs, usually fiber-optic cable. However, if you need a 40 or 100 Gigabit Ethernet network, fiber-optic is really your only choice.
- *Budget*—How much money can you spend on cabling? Sometimes budget alone dictates a choice. A typical UTP cable installation costs \$100 to \$200 per cable run, whereas fiber-optic might cost twice this much. Wireless media have no physical installation costs, but you need to install access points and verify connectivity from all locations.
- *Environmental considerations*—How electrically noisy is the deployment environment? How important is data security? Sometimes high-EMI environments or security requirements can dictate cable choices, regardless of other factors. The more weight either factor has, the more likely a choice fiber-optic cable is (or in lower bandwidth applications, a secure wireless network).
- *Span*—What kind of distance must the network span? Longer spans might require fiber-optic cabling or wireless technologies used between buildings. Strategic placement of small switches or hubs for use with UTP wiring gives UTP surprising reach in many office environments where workers tend to cluster in groups, even if these groups are widely scattered.

- *Existing cable plant*—For a new installation, only the previously listed criteria need to be considered, but for an upgrade, the existing cable plant must be considered. For example, if some existing cable is to remain, is it compatible with the speeds and new equipment that are planned?

Networks combining fiber-optic, UTP, and wireless media have almost become the norm, with fiber-optic cables providing a backbone that ties together clusters of devices networked with UTP cable through switches and wiring centers. With wireless networks, users can stay connected with their Wi-Fi-enabled phones, laptops, and iPads. Table 4-9 condenses the most important information for the cable types covered in this chapter.



Table 4-9 Comparison of LAN media characteristics

Type	Maximum cable length	Bandwidth	Installation	Interference	Cost
UTP	100 m	10–10000 Mbps	Easy	High	Cheapest
STP	100 m	16–10000 Mbps	Moderate	Moderate	Moderate
Fiber optic	2–100 km	100 Mbps–10 Gbps	Moderate	None	Most expensive
Wireless	100-300 feet	11 to 300 Mbps	Easy	Moderate	None for physical media

Chapter Summary

- Wired networking media come in two main categories: copper and fiber optic. Cable characteristics include bandwidth rating, maximum segment length, susceptibility to interference and eavesdropping, and cable grade.
- Twisted-pair cabling come in shielded or unshielded varieties. Most networks use UTP, but STP can be used in electrically noisy environments. Cat 5e and Cat 6 are the most common cable types in network today.
- Twisted-pair cabling components consist of connectors, patch cable, jacks, patch panels, and distribution racks. A structured cabling plant consists of work areas, horizontal wiring, telecommunications closets, equipment rooms, backbone cabling, and entrance facilities.
- Fiber-optic cable uses pulses of light to represent bits and is immune to EMI, RFI, and electronic eavesdropping. Commercial implementations of up to 100 Gbps are in use. Each network connection requires two strands of fiber-optic cable: one for transmit and one for receive. Fiber-optic cable comes in single-mode or multimode; single-mode uses lasers and can carry data longer distances, and multimode uses LEDs.
- Wireless networks can be subdivided into LANs, extended LANs, and mobile computing. The components of a wireless LAN are a NIC, an antenna, and a transceiver or access point. Wireless networks send signals in the form of electromagnetic waves. Different network types use different frequencies for signal transmission.
- Different technologies are used to transmit and receive data, including infrared, laser, narrowband radio, and spread-spectrum radio. Infrared can deliver speeds up to 100 Mbps and is used in some LAN applications. Laser-based technologies require line of sight between sender and receiver, as does infrared, but laser isn't as susceptible to

interference from other light sources. Narrowband radio uses low-power two-way radio communication and is highly susceptible to interference. Spread-spectrum LANs are the most common and are used for 802.11 b/g/n Wi-Fi networks.

- Criteria for choosing LAN media include needed bandwidth, budget, environmental factors, the distance the network must span, and the existing cable plant, if any. Networks combining fiber-optic, UTP, and wireless have become the norm.

Key Terms

attenuation Weakening of a signal as it travels the length of the medium.

backbone cabling Network cabling that interconnects telecommunications closets and equipment rooms. This cabling runs between floors or wings of a building and between buildings to carry network traffic destined for devices outside the work area. It's often fiber-optic cable but can also be UTP. Also called "vertical cabling."

cable plant The collection of all cables and connectors tying a network together.

cable segment A length of cable between two network devices, such as a NIC and a switch. Any intermediate passive (unpowered) devices, such as wall jacks, are considered part of the total segment length.

crossover cable A type of patch cable that uses the 586B standard on one end and the 586A standard on the other end. This arrangement crosses the transmit and receive wires so that transmit on one end connects to receive on the other end. Often used to connect two devices of the same type to one another—for example, connecting a switch to a switch.

crosstalk Interference one wire generates on another wire when both wires are in a bundle.

datagrade A grade of cable suitable for data networking.

differential signal A method for transmitting data in which two wires of opposite polarity are used. One wire transmits using positive voltage and the other uses negative voltage. Differential signals enhance signal reliability by providing a canceling affect on EMI and crosstalk.

electromagnetic interference (EMI) A disturbance to the operation of an electronic circuit or its data, caused by devices that emit an electromagnetic field.

encoding term The method used to represent bits on a medium.

entrance facility The location of cabling and equipment that connects a corporate network to a third-party telecommunications provider. It can also serve as an equipment room and the main cross-connect for all backbone cabling.

equipment room A room that houses servers, routers, switches, and other major network equipment and serves as a connection point for backbone cabling running between telecommunications closets.

extended LANs A LAN that's expanded beyond its normal distance limitations with wireless communication.

fiber-optic cable A cable type that carries data over thin strands of glass by using optical (light) pulses to represent bits.

hertz (Hz) A unit expressing how many times per second a signal or electromagnetic wave occurs.

horizontal wiring The network cabling running from the work area's wall jack to the telecommunications closet, usually terminated at a patch panel. The total maximum distance for horizontal wiring is 100 meters.

infrared (IR) A very long wavelength light source in the invisible spectrum that can be used to transmit data wirelessly.

IrDA devices Devices that use infrared signals to communicate. IrDA stands for Infrared Device Association.

MDI crossed (MDI-X) devices Network devices that connect by using RJ-45 plugs over twisted-pair cabling; they transmit over pins 3 and 6 and receive over pins 1 and 2 of an RJ-45 connector.

medium dependent interface (MDI) devices Network devices that connect by using RJ-45 plugs over twisted-pair cabling; they transmit on pins 1 and 2 and receive on pins 3 and 6 of an RJ-45 connector.

narrowband radio Low-powered, two-way radio communication systems, such as those used in taxis, police radios, and other private radio systems. Also called “single-frequency radio.”

patch cable A short cable for connecting a computer to an RJ-45 jack or connecting a patch-panel port to a switch or hub. *See also* straight-through cable.

radio frequency interference (RFI) Similar to EMI, but RFI is usually interference caused by strong broadcast sources. *See also* electromagnetic interference (EMI).

RJ-45 jack A device used in the work area in wall plates and surface-mounted boxes to plug a patch cable that connects a computer to the horizontal wiring.

RJ-45 plug A connector used to terminate twisted-pair cable for making patch cables. It has eight wire traces to accommodate a standard twisted-pair cable with four wire pairs.

satellite microwave Microwave communication systems that send and receive data from satellites that maintain fixed positions in the sky.

spread-spectrum radio A radio communication system that uses multiple frequencies simultaneously, thereby improving reliability and reducing susceptibility to interference over narrowband radio.

straight-through cable A standard patch cable that uses the same wiring standards on both ends so that each wire is in the same location on both ends of the cable (pin 1 goes to pin 1, pin 2 to pin 2, and so forth). *See also* patch cable.

structured cabling A specification for organizing cabling in data and voice networks, regardless of the media type or network architecture.

telecommunications closet (TC) Usually an enclosed space or room that provides connectivity to computer equipment in the nearby work area; can also serve as the entrance facility in small installations. Typical equipment includes patch panels to terminate horizontal wiring runs, hubs, and switches.

termination The attachment of RJ-45 plugs on a cable to make a patch cable or punching down the cable wires into terminal blocks on a jack or patch panel.

terrestrial microwave Line-of-sight transmissions between microwave towers or between transmitters and receivers mounted on tall buildings, mountaintops, or other locations with long, clear lines of sight.

transceiver A device that transmits and receives. In wireless networking, an access point is a transceiver.

twisted-pair (TP) cable A cable containing one or more pairs of insulated strands of copper wire twisted around one another and housed in an outer sheath.

voicegrade A grade of cable that’s not suitable for data networking but is suitable for voice communication.



wireless bridge A wireless network arrangement that connects networks up to 3 miles (4.4 km) apart, allowing locations to be linked with line-of-sight or broadcast transmissions.

work area The location of workstations and other user devices—in short, the place where people work with computers and other network devices.

Review Questions

1. Which of the following is a common characteristic of a networking medium? (Choose all that apply.)
 - a. Bandwidth rating
 - b. Interference susceptibility
 - c. Broadband rating
 - d. Maximum segment length
2. Which of the following types of fiber-optic connectors provides high density and requires only one connector for two cables?
 - a. SC
 - b. ST
 - c. MT-RJ
 - d. RJ-45
3. Which of the following conditions requires cables not to exceed a recommended maximum length?
 - a. Diminution
 - b. Capacitance
 - c. Bandwidth
 - d. Attenuation
4. Which of the following is the process for representing bit signals on the medium?
 - a. Encryption
 - b. Encoding
 - c. Decryption
 - d. Decoding
5. What happens to signals as they travel the length of the medium?
 - a. They decode.
 - b. They amplify.
 - c. They attenuate.
 - d. They implode.

6. Which of the following is UTP susceptible to? (Choose all that apply.)
 - a. EMI
 - b. Crosstalk
 - c. Signal enhancement
 - d. LEDs
7. The space between a false ceiling and the true ceiling where heating and cooling air circulates is called the _____ .
 - a. Duct-equivalent airspace
 - b. Conduit
 - c. Return air
 - d. Plenum
8. What type of connector is used most commonly with TP network wiring?
 - a. RJ-11
 - b. RJ-45
 - c. BNC
 - d. MT-RJ
9. You have been hired to install a network at a large government agency that wants to reduce the likelihood of electronic eavesdropping on its network. What type of cable is most resistant to eavesdropping?
 - a. UTP
 - b. STP
 - c. Coaxial
 - d. Fiber optic
10. Which of the following is a characteristic of unshielded twisted-pair cable? (Choose all that apply.)
 - a. Consists of four wires
 - b. Commonly used in physical bus topologies
 - c. Has a distance limitation of 100 meters
 - d. Is susceptible to electrical interference
11. Which of the following is a characteristic of fiber-optic cabling? (Choose all that apply.)
 - a. Can be used in electrically noisy environments
 - b. Requires only a single strand of fiber for network connections
 - c. Carries data over longer distances than UTP
 - d. Has low bandwidth



12. You're preparing to install a conventional Ethernet network in your new office building, but your boss tells you to be ready to handle a switchover to 1 Gbps Ethernet next year. What types of cable could you install? (Choose all that apply.)
 - a. Cat 5
 - b. Fiber optic
 - c. Cat 4
 - d. Cat 6
 - e. Coax
13. When two cables run side by side, signals traveling down one wire might interfere with signals traveling on the other wire. What is this phenomenon called?
 - a. RFI
 - b. Attenuation
 - c. Impedance
 - d. Crosstalk
14. What characteristic of twisted-pair cabling helps mitigate the effects of crosstalk?
 - a. Differential signals
 - b. Copper conductors
 - c. Four pairs of wires
 - d. 100-ohm impedance
15. What is the wireless device used to link buildings without cable?
 - a. Hub
 - b. Router
 - c. Gateway
 - d. Bridge
16. Which of the following is a wiring standard for twisted-pair cable connections? (Choose all that apply.)
 - a. IEEE 802.3a
 - b. TIA/EIA 568A
 - c. IEEE 802.3b
 - d. TIA/EIA 568B
17. Which of the following is a component of a structured cabling system? (Choose all that apply.)
 - a. Patch cables
 - b. RJ-11 plugs
 - c. Coax cable
 - d. Horizontal wiring

18. Where are you most likely to find vertical cabling? (Choose all that apply.)
 - a. Equipment rooms
 - b. In the work area
 - c. Connecting TCs
 - d. Connecting a work area to a TC
19. Which of the following is a tool needed to make a patch cable? (Choose all that apply.)
 - a. 110 punchdown tool
 - b. Cable stripper
 - c. Crimping tool
 - d. RJ-45 jack
20. Which type of connection is most likely to require a crossover cable?
 - a. PC to hub
 - b. Hub to router
 - c. Router to switch
 - d. PC to router
21. Which UTP limitations can be solved by fiber-optic cable? (Choose all that apply.)
 - a. Bandwidth
 - b. EMI susceptibility
 - c. Installation cost
 - d. Segment length
22. How many strands of fiber-optic cable are needed for a network connection?
 - a. 1
 - b. 2
 - c. 4
 - d. 8
23. Which statement is true about fiber-optic cables?
 - a. MMF uses lasers and has a thicker core.
 - b. SMF uses lasers and has a thinner core.
 - c. MMF uses LEDs and has a thinner core.
 - d. SMF uses LEDs and has a thicker core.
24. Which type of wireless network requires a clear line of sight between transmitter and receiver? (Choose all that apply.)
 - a. Infrared
 - b. Narrowband radio
 - c. Spread-spectrum LAN
 - d. Terrestrial microwave



25. Which of the following wireless technologies does a 802.11 wireless network using the 2.4 GHz frequency range use?
- Infrared
 - Narrowband radio
 - Frequency hopping
 - Direct-sequence spread spectrum

Case Projects



Case Project 4-1

During the design of most real-world networks, you'll discover that using more than one type of networking medium is common. The usual reasons for needing more than one type of medium include the following:

- Two or more areas must be interconnected, and the distance separating them is greater than the maximum segment length for the type of medium used in (or best suited for) each area.
- A connection must pass through a high-interference environment (across some large transformers, near heavy-duty electrical motors, and so on). Failure to use a different type of medium increases the risk of impeding data flow. This reason is especially common for choosing fiber-optic cable or wireless in many networks, particularly when connecting floors in an office building and the only available pathway is the elevator shaft.
- Certain parts of an internetwork might have to carry more traffic than other parts. Typically, the segment where traffic aggregates is the backbone, a common cable segment that interconnects subsidiary networks. (Think of a tree trunk as the backbone and its major branches as cable segments.) Often, a higher-capacity cable is used for a backbone (for example, fiber-optic cable or Cat 6 cable rated for Gigabit Ethernet), along with a higher-speed networking technology for attachments to the backbone. This arrangement means outlying segments might use conventional 10 or 100 Mbps Ethernet, and the backbone uses 1 Gbps or 10 Gbps Ethernet.

Using this information, suggest solutions that involve at least two types, if possible, of networking media to address the following problems:

- A—XYZ Corp. is planning a new network. Engineers in the design shop must have connections to accountants and salespeople in the front office, but all routes between the two areas must traverse the shop floor, where arc welders and metal-stamping equipment create potent amounts of EMI and RFI. Given that both the design shop and front office use 10BaseT (twisted-pair Ethernet), how might you interconnect these two areas? What medium guarantees immunity from interference?

- B—After the front-office network at XYZ Corp. is set up, an accountant realizes that if the loading dock connected to the network, dock workers could log incoming and outgoing shipments and keep the inventory more current. Even though the loading dock is nowhere near the shop floor, the dock is 1100 feet from the front office. What kinds of cable will work to make this connection? What kind would you choose and why?
- C—ABC Company occupies three floors in a 10-story building, where the elevator shaft provides the only path to all three floors. In addition, users on the ninth and tenth floors must access a collection of servers on the eighth floor. Explain what kind of connections would work in the elevator shaft. If more than one choice is possible, pick the best option and explain the reasons for your choice. Assuming that interfloor connections might someday need to run at much higher speeds, reevaluate your choice. What's the best type of medium for open-ended bandwidth needs? Explain your answer.
- D—Very Big ISP (VBISP) Corporation wants to increase the bandwidth it can access at its downtown location in New York City. The distance between locations is about 20 miles, and the bandwidth needed between locations is at least 50 Mbps. What media types could work to provide this connection?
- E—Following a year of major sales increases in the Pacific Rim, MarTexCo decides to open a second plant in Malaysia. The company wants the new plant to be able to access the headquarters database in Des Moines, Iowa, in real time, but long-haul telephone connections aren't possible, owing to the lack of communications infrastructure at the Malaysia location. What kind of wireless networking alternative makes the most sense when considering network links that span an appreciable portion of the globe? Explain why laying cable might not be feasible.



Case Project 4-2

XYZ Corp.'s Nashua, NH, facilities are two office buildings 400 feet apart, each with its own LAN. To connect the two networks, you plan to dig a trench and lay cable in conduit between the two buildings. You want to use fiber-optic cable, but your budget-conscious facilities manager wants to use 100 Mbps Ethernet over twisted-pair cable. Which of the following reasons can you use to justify fiber-optic cable in this case, and why?

- a: Twisted pair won't span a 400-foot distance.
- b: Fiber-optic cable is cheaper and easier to work with than twisted pair.
- c: Twisted pair is a conductive cable and can, therefore, carry current based on the difference in ground potential between the two buildings.
- d: Fiber-optic cable leaves more room for growth and future needs for increased bandwidth than twisted pair does.

Case Project 4-3

TVBCA has just occupied a historic building in downtown Pittsburgh in which 15 employees will work. Because of codes for historic buildings, TVBCA isn't permitted to run cables inside walls or ceilings.

Required result: Employees must be able to share files and printers, as in a typical LAN environment, without using cables.

Optional desired results: Employees must be able to use their laptops and move freely throughout the office while maintaining a network connection. Because of the size of some computer-aided design (CAD) files employees use frequently, data transfer speeds should be more than 20 Mbps.

Proposed solution: Install an 802.11n wireless access point and configure each computer and laptop with a wireless network card. Which of the following results does the proposed solution deliver? Explain your answer.

- a: The proposed solution delivers the required result and both optional desired results.
- b: The proposed solution delivers the required result and only one of the two optional desired results.
- c: The proposed solution delivers the required result but neither optional desired result.
- d: The proposed solution does not deliver the required result.

Network Protocols

After reading this chapter and completing the exercises, you will be able to:

- Describe the purpose of a network protocol, the layers in the TCP/IP architecture, and the protocols in each TCP/IP layer
- Explain IP address configuration and subnetting

For effective communication across a network, computers must be capable of transmitting data reliably and efficiently. Network protocols are designed to accomplish this goal, with some protocols emphasizing reliability and others efficiency. Network protocols often work together at different layers of the network communication process to provide both reliability and efficiency. Network administrators must understand the role and function of protocols, as much of their time is spent configuring and troubleshooting the protocols used by the network's clients and servers. This chapter discusses network protocols in general but focuses on the most common suite of protocols used in today's networks: TCP/IP.

TCP/IP's Layered Architecture

The term “protocol” isn't specific to the field of networking. In general, **protocols** are rules and procedures for communication and behavior or etiquette. Just as two people must share a common set of rules for verbal communication—a language—computers must also “speak” the same language and agree on the rules of communication. You use protocols in other ways. Texting, e-mail, and Facebook communication, for example, have their own rules of etiquette and, especially for texting, their own language.

Until fairly recently, you had a choice of network protocols you could install on your computer, depending on the computing environment. A small network in the 1990s running Windows 3.1 or Windows 95 probably ran the Windows-specific NetBEUI protocol. A network with Novell NetWare 4.x servers typically ran IPX/SPX. Both these protocols are obsolete now and are found only in networks that haven't been upgraded in more than a decade. Today, you can focus on the TCP/IP protocol suite, the protocol of the Internet and the default protocol all contemporary OSs run.

When a set of protocols works cooperatively, it's called a **protocol stack** or **protocol suite**. The most common protocol stack is **Transmission Control Protocol/Internet Protocol (TCP/IP)**, the Internet protocol suite. Although you can see by its name that TCP/IP consists of at least two protocols—TCP and IP—this protocol suite is actually composed of more than a dozen protocols operating at different layers of the communication process.

Recall the communication process explained in Chapter 1 and animated in Simulation 1. This discussion was an introduction to the idea of communication taking place in layers. The protocols in TCP/IP can also be divided into four layers, with similar names and functions. Figure 5-1 shows the layers of the TCP/IP protocol suite and which protocols operate at each layer. This layered architecture is usually referred to as the “TCP/IP model.”



Many books and Web sites about TCP/IP call the Internetwork layer the “Internet layer,” but the term “internetwork” describes the layer's function more accurately, especially because many people use the term Internet interchangeably with the term World Wide Web. Also, the Network access layer is often referred to as the “Network interface layer.” Although both terms describe this layer's function, “Network access” has been used in this book.

The TCP/IP protocol suite includes more protocols than are shown in Figure 5-1, but the ones listed are some of the most common protocols used in networks. Before you examine each layer and protocol more closely, take a look at an example of how the layers work together.

Layer name	TCP/IP protocols			
Application	HTTP	FTP	DHCP	TFTP
	SMTP	POP3	DNS	SNMP
Transport	TCP		UDP	
Internetwork	ICMP	ARP	IPSec	
	IPv4 and IPv6			
Network access	Ethernet, token ring, FDDI, WAN technologies			

Figure 5-1 The TCP/IP layered architecture

Courtesy of Course Technology/Cengage Learning

Suppose you start your Web browser and have configured your home page as *http://www.course.com*. The Web browser formats a request for a page on the *www.course.com* Web server by using the Application-layer protocol HTTP. The request looks something like Figure 5-2.

get the course.com home page

Figure 5-2 The Application layer creates data

Courtesy of Course Technology/Cengage Learning

You've learned about packets and frames, but the unit of information the Application layer works with is simply called "data." The Application-layer protocol HTTP passes the request down to the Transport-layer protocol: in this case, TCP. Notice that the four Application-layer protocols on the left of Figure 5-1 use TCP as the Transport-layer protocol, and the Application-layer protocols on the right use UDP. (The difference between TCP and UDP is explained later in "Role of the Transport Layer.") TCP has its own job to do, so it adds a header to the request that looks like Figure 5-3.

TCP header get the course.com home page

Figure 5-3 The Transport layer adds its header to make a segment

Courtesy of Course Technology/Cengage Learning

The unit of information the Transport layer works with is called a **segment**. TCP passes the segment to the Internetwork layer. The Internetwork layer has a number of subprotocols, but most operate by following the basic rules and format of IP. IP then places its header on the segment, making it a packet (see Figure 5-4).



Figure 5-4 The Internetwork layer creates a packet

Courtesy of Course Technology/Cengage Learning

The packet is almost ready for delivery to the network medium, with one more stop at the Network access layer, where the NIC operates. As you know, NICs work with frames, so a frame header and trailer are added (see Figure 5-5).




Figure 5-5 The frame is created and ready for delivery on the medium

Courtesy of Course Technology/Cengage Learning

The frame is then delivered to the network medium as bits on its way to the *www.course.com* server, where the Web server software processes it and returns a Web page to the computer that originated the request. Now that you have an idea of how these protocols work together, examine the roles of these four layers more closely, starting from the bottom: the Network access layer.



Hands-On Project 5-1: Viewing TCP/IP Layers in Windows and Configuring Your IP Address

Time Required: 10 minutes

Objective: View the properties of your computer's network connection, identify the TCP/IP layers, and configure your IP address.

Required Tools/Equipment: Your classroom computer

Description: In this project, you view the properties of your computer's local area connection and identify the TCP/IP layers. This project is similar to Hands-On Project 1-3, but you're viewing the TCP/IP protocol suite layers instead of the more general layers of the networking process. Next, you configure your IPv4 address.

1. Start your computer and log on as **NetAdmin**.
2. Open the Network and Sharing Center by clicking **Start, Control Panel**. Under Network and Internet, click **View network status and tasks**.
3. In the left pane of the Network and Sharing Center, click **Change adapter settings**. Right-click **Local Area Connection** and click **Properties** to open the Local Area Connection Properties dialog box.
4. The Connect using text box displays the network interface card. In the list box under it, you see several items. Client for Microsoft Networks, File and Printer Sharing for Microsoft Networks, Internet Protocol Version 4, and Internet Protocol Version 6 are the items you're interested in right now, as they're the most necessary software components for making network communication work.
5. For each component, write which TCP/IP layer or layers you think it operates in:
 - NIC displayed in the Connect using text box: _____

- Client for Microsoft Networks: _____
 - File and Printer Sharing for Microsoft Networks: _____
 - Internet Protocol Version 4: _____
 - Internet Protocol Version 6: _____
6. Next, you configure your IP address settings. Click **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.
 7. If your IP settings have the “Obtain an IP address automatically” option enabled, click **Use the following IP address**. You use this option to set a static IP address. If your address is already static, make a note of it, and skip entering the information in Step 8. Click **OK**.
 8. For the following IP address settings fields, enter the information shown unless your instructor tells you to use different values, and then click **OK** when you're finished:
 - IP address: 192.168.100.XX (replacing XX with your student number)
 - Subnet mask: 255.255.255.0
 - Default gateway: provided by your instructor
 - Preferred DNS server: provided by your instructor
 9. Click **Close**. If you're prompted to set a network location, click **Work network**, and then click **Close**.
 10. To test your configuration, open a command prompt window and try to ping your default gateway address and your preferred DNS server address. If either ping is unsuccessful, inform your instructor and troubleshoot your settings.
 11. Close all open windows, but leave your computer running for the next project.

The following list recaps your IP address configuration and explains each item's purpose:

- The IP address provides your computer with a unique internetwork identity on a logical IP network.
- The subnet mask defines which part of the IP address is the network ID and which is the host ID.
- The default gateway is a router in your network that your computer sends packets to when the destination is a remote network.
- The preferred DNS server is the address of a DNS server that resolves computer names to IP addresses.

These items are explained in more detail in this chapter and throughout this book.



Hands-On Project 5-2: Identifying the TCP/IP Layers in a Frame

Time Required: 10 minutes

Objective: Capture packets and view the TCP/IP layers in the frame.



Required Tools/Equipment: Your classroom computer with Wireshark installed

Description: In this project, you capture some frames generated by your Web browser and examine the captured frames to identify the TCP/IP layers.

1. If necessary, log on to your computer as **NetAdmin**.
2. Start Wireshark and click **Capture Options**. In the Capture Filter text box, type **tcp port http**, and then click **Start**.
3. Start a Web browser, and after the home page loads, exit the browser.
4. In Wireshark, click the **Stop the running live capture** toolbar icon to stop the capture. Scroll up to the first packet summary line, if necessary.
5. Click a packet summary in the top pane with HTTP in the protocol field and an Info line beginning with GET. In the middle pane are summaries of each protocol header (see Figure 5-6). You can ignore the first line starting with Frame X (with X representing the frame number), as it gives information about the frame, such as the time it arrived, its length, protocols in the frame, and so forth.

```

Frame 4 (629 bytes on wire, 629 bytes captured)
Ethernet II, Src: Supermic_67:7e:6c (00:30:48:67:7e:6c), Dst: Cisco_42:22:c0 (00:0c:85:42:22:c0)
Internet Protocol, Src: 172.31.210.1 (172.31.210.1), Dst: 174.129.210.177 (174.129.210.177)
Transmission Control Protocol, Src Port: 52091 (52091), Dst Port: http (80), Seq: 1, Ack: 1, Len: 575
Hypertext Transfer Protocol
  
```

Figure 5-6 Summary of protocol headers in Wireshark

Courtesy of Course Technology/Cengage Learning

6. Click to expand the line beginning with **Ethernet II**. Examine the information in this header (discussed in more detail in the following sections). Write which layer of the TCP/IP model the Ethernet II header represents, and then click again to collapse this header:

7. Click to expand the line beginning with **Internet Protocol**. Examine the information in this header (discussed in more detail in the following sections). Write which layer of the TCP/IP model the Internet Protocol header represents, and then click again to collapse this header:

8. Click to expand the line beginning with **Transmission Control Protocol**. Examine the information in this header (discussed in more detail in the following sections). Write which layer of the TCP/IP model the Transmission Control Protocol header represents, and then click again to collapse this header:

9. Click to expand the line beginning with **Hypertext Transfer Protocol**, and examine the information. This data portion of the frame is what a Web server actually sees and responds to. In this case, the HTTP command is **GET**, which means HTTP is requesting a page (or part of a page) from the Web server. Write which layer of the TCP/IP model the HTTP protocol represents, and then click again to collapse this header:

10. Exit Wireshark and click **Quit without Saving** when prompted.
11. Close all open windows, but leave your computer running for the next project.

Role of the Network Access Layer

Strictly speaking, the Network access layer isn't composed of TCP/IP protocols. As you can see, network technologies, such as Ethernet and token ring, operate at this layer. So this layer is part of the TCP/IP architecture only to the extent that the layer above—the Internetwork layer—has the capability to communicate with any network technologies following the rules of the Network access layer. Some tasks the Network access layer performs have already been discussed but are worth repeating here:

- Provides a physical (MAC) address for the network interface
- Verifies that incoming frames have the correct destination MAC address
- Defines and follows media access rules
- Receives packets from the Internetwork layer and encapsulates them to create frames
- Deencapsulates received frames and sends the resulting packets to the Internetwork layer
- Often provides frame error detection in the form of a CRC code
- Transmits and receives bit signals
- Defines the signaling needed to transmit bits, whether electrical, light pulses, or radio waves
- Defines the media and connectors needed to make a physical network connection

As you learn in Chapter 6, the last three items in this list are tasks the Physical layer performs in the more detailed OSI networking model, which splits the Network access layer into two separate layers.

Role of the Internetwork Layer

The Internetwork layer is where administrators usually do the most network configuration. It's where the IP protocol operates, and it can be looked at as the heart of the TCP/IP protocol suite. IP addresses, of course, are defined here, and routing takes place in this layer, too. Without routing, the Internet and World Wide Web wouldn't exist. With all the complexity of configuring routing and managing IP addresses, this layer is also where most errors in network configuration occur. In a large internetwork, a lot of your time is typically spent unraveling the intricacies of the Internetwork layer.

The Internetwork layer is responsible for four main tasks, discussed in the following sections:

- Defines and verifies IP addresses
- Routes packets through an internetwork
- Resolves MAC addresses from IP addresses
- Delivers packets efficiently

Defines and Verifies IP Addresses An IP address is assigned to every computer and network device using TCP/IP for communication. IP addresses are used for two main purposes: to identify a network device at the Internetwork layer and to identify the network on which a device resides. When an IP address is assigned to a computer or network device (referred to as an “IP host” or just “host”), the host's Internetwork-layer identity is defined. When a host receives an IP packet, it compares the packet's destination IP address with its own address to verify that the packet was delivered correctly. If the destination address matches or is a broadcast or recognized multicast address, the packet is processed;



otherwise, it's discarded. When a host sends a packet, the IP protocol places its own IP address in the packet header's source field before sending the packet to the network interface.

The IP address is also used to identify the network on which a host resides. Every IP address contains two parts: a network ID and a host ID. This format is similar to a 10-digit phone number, with a three-digit area code identifying the region of the country where the number was assigned and a seven-digit number identifying the particular phone. IP addresses aren't as straightforward in their format, as you discover later in "IP Addressing," but there's always a portion of an IP address that identifies the network on which the host resides.

Routes Packets Through an Internetwork The next task of the Internetwork layer is determining the best way to get a packet from network to network until it reaches its destination. If there were only one way for a packet to get from here to there, this aspect of the Internetwork layer's job would be pretty ho-hum. However, much like the nation's road system, most large networks, such as the Internet, have multiple paths for getting from location A to location B. Which path to take isn't always a clear-cut decision. Some paths are heavily traveled, and some are lightly traveled; some paths have construction or accidents, and others are clear sailing.

As mentioned, routers work at the Internetwork layer, and their job is to select the best path to the destination. If a path becomes unavailable or congested, they select an alternative, if available. Routers use the network ID portion of IP addresses along with their routing tables to determine on which network a destination device can be found and the best way to get packets to their destination. Chapter 7 discusses routers in more detail.

Resolves MAC Addresses from IP Addresses As you've learned, every frame sent to the network medium contains both physical (MAC) and logical (IP) source and destination addresses. When a packet is ready to be sent to the Network access layer, the destination device's MAC address must be retrieved before the frame header can be constructed. TCP/IP uses Address Resolution Protocol (ARP) for this task. ARP is discussed in more detail later in "Address Resolution Protocol," but in a nutshell, it returns a computer's MAC address by asking the network which computer is assigned a particular IP address.

Delivers Packets Efficiently Internetwork-layer protocols focus mainly on efficient delivery of packets. The secret to achieving this efficiency is low processing overhead. Internetwork-layer protocols don't include features such as flow control, delivery confirmation, or message reassembly; these features require considerable overhead to ensure reliable delivery, at the cost of efficiency.

Internetwork protocols rely on protocols in the Transport and Application layers to provide advanced reliability features. Protocols at this layer are concerned with one packet at a time, with no concern for packets that came before or after it and with no confirmation that delivery was successful. This communication strategy is called connectionless communication, and protocols using it are called connectionless protocols.

When using a **connectionless protocol**, no lasting connection is made from source to destination. A connectionless protocol relies on an upper-layer protocol to ensure the packet's safe journey. This process is much like delivering a first-class letter via the U.S. mail. You drop the letter in the mailbox and hope it makes it to its destination. Usually it does, but when you want to be certain it

was received (or notified if it wasn't), you must add a layer of complexity by sending the letter certified mail, which requires an acknowledgement that the letter was received. In networking, protocols providing this acknowledgement are called connection-oriented protocols. In the TCP/IP model, these protocols are in the Transport layer, discussed later in "Role of the Transport Layer."

Protocols at the Internetwork Layer

IP is the underlying basis for most Internetwork-layer protocols, which means they just send specialized versions of IP packets. The protocols operating at this layer are too numerous to describe in this book, so the following sections focus on the most commonly used:

- IPv4
- IPv6
- ICMP
- ARP
- IPSec

Internet Protocol Version 4 Internet Protocol Version 4 (IPv4), or just IP, is an Internetwork-layer protocol that provides source and destination addressing and routing for the TCP/IP protocol suite. IP is a connectionless protocol, so it's efficient but unreliable. Note that "unreliable" doesn't mean it fails often. In this context, it simply means IP has no method for ensuring that data is delivered to the destination. IP assumes the Transport or Application layer provides reliable data delivery in applications that require it.

IPv4 is the most common version in networks and the first version that was in widespread use. Earlier versions never really made it out of the lab. One of IP's most important functions is the definition of logical addresses, naturally called **IP addresses**. IPv4 defines a 32-bit dotted decimal address: 172.31.149.10, for example. Each grouping of numbers separated by a dot (period) is an 8-bit value that can range from 0 to 255. Because an IP address has 32 bits, a total of 2^{32} addresses are possible, which is approximately 4 billion. That might seem like a lot of addresses, but as you learn later, many are wasted, and available addresses to assign to devices on the Internet are running out.

As mentioned, in the IP address format, part of the address specifies the network where the computer assigned the address is located, and the rest of the address specifies a unique host ID in the network. For example, using the address 172.31.149.10, 172.31 might be the network ID and 149.10 might be the host ID. This topic can be complex and is covered in more detail later in "IP Addressing."

IP works with packets, and when it receives a message from the layers above, it adds an IP header. So far in this chapter, only the destination (the intended recipient) and source (the sending machine) addresses of the IP header have been discussed. There's quite a bit more to an IP header, and the following list describes the more important fields in the order they appear in a packet:

- *Version*—This field simply indicates which version of IP is in use. Today, the possibilities are 4 and 6. Because computers can run both versions at the same time, the Version field tells the computer whether the packet should be processed by IPv4 or IPv6.
- *Time to live*—The TTL field is a safeguard that prevents a packet from wandering aimlessly through an internetwork, which can be caused by a network misconfiguration.



Before a packet is sent to the network, the TTL field is given a value, usually 64, 128, or 255. As the packet travels through the internetwork, the TTL value is decremented at each router. If the TTL value reaches 0, the packet is deemed to have expired, and the router that decremented the packet to 0 discards it. Most routers also send a message back to the source address as notification that the packet expired.

- *Protocol*—This field is a numeric code specifying the type of IP packet or the next layer protocol contained in the packet. For example, the Protocol value in an ICMP packet (used by the Ping program) is 1. If the packet contains a Transport-layer protocol, such as TCP, the value is 6. There are more than 140 different types of IP packets.
- *Checksum*—This field is a value produced by a mathematical calculation on data in the header that protects the IP header’s contents. When a network device receives a packet, this value is recalculated and compared with the value in the Checksum field. If they match, the header hasn’t been altered.
- *Source address*—This field is self-explanatory.
- *Destination address*—This field is self-explanatory, too, but it can be one of three types, as in a MAC address: unicast (intended for a single computer), broadcast (sent to all computers in the network), or multicast (sent to a group of computers).

Internet Protocol Version 6 Internet Protocol version 6 (IPv6) has many of the features of IPv4 but addresses IPv4’s shortcomings, not the least of which is its inadequate 32-bit address space. The advantage of a layered approach to networking is that IPv6 can run on computers alongside IPv4 without needing to change the Transport layer or Network access layer. Most Application-layer protocols require no changes either, except those dealing directly with IP addresses, such as DHCP and DNS. IPv6 is discussed in more detail later in “Introduction to Internet Protocol Version 6.”

Address Resolution Protocol Address Resolution Protocol (ARP) is used to resolve a logical (IP) address to a physical (MAC) address. When a system begins a conversation with a host and doesn’t have its MAC address to create the frame header, it sends an ARP broadcast frame requesting the MAC address corresponding to the host’s IP address. A network device configured with the specified IP address responds with an ARP reply message containing its MAC address. Then the packet is sent to the Network access layer, and the frame can be constructed.

This process requires additional explanation, as you might be wondering what happens when the two computers are on separate networks on a corporate internetwork or even miles apart on the Internet. As discussed in Chapter 2, routers are responsible for getting packets from one network to another, and they don’t forward broadcast packets, which makes routers the delimiting device for broadcast domains. If routers did forward broadcasts, when any computer on the Internet sent a broadcast, the message would be forwarded to every LAN on the Internet, and the Internet would be overrun with broadcasts.

When a computer using TCP/IP wants to communicate with another computer, it must know the destination computer’s IP address. Usually, the application sending the message knows the address, or it’s resolved by using a name lookup. If the destination’s IP address is on the same network as the source, the source computer sends an ARP request message in the form of a broadcast. All computers in the broadcast domain process the ARP request, and the computer with this IP address sends back an ARP reply containing its MAC

address. After the MAC address is received, the frame header can be constructed, and the frame is delivered to the destination.

If a computer has to send an ARP broadcast every time it wants to send an IP packet to a particular destination, the network would have one broadcast ARP frame for every frame carrying actual data, which is a big waste of bandwidth. To avoid sending an ARP request every time an IP packet is sent, PCs and other devices store learned IP address/MAC address pairs in an **ARP cache**, a temporary location in RAM. (You viewed your ARP cache in Hands-On Project 1-4 with the `arp -a` command.) So a computer or router has to send an ARP broadcast only once for each destination host it communicates with on its network.



ARP cache entries aren't kept indefinitely. Most computers keep each entry for only a few minutes after it's last used to avoid storing inaccurate information, which could result from a changed NIC or IP address.

If the destination computer is on another network, the computer uses ARP to retrieve the MAC address of the router configured as its default gateway. The packet is delivered to the router, and the router determines where the packet should go next to get to its destination. When the packet gets to the destination network, the router on the destination network uses ARP to get the destination computer's MAC address. Figure 5-7 illustrates this process, and it's animated in Simulation 9. Notice that the destination MAC address in the original

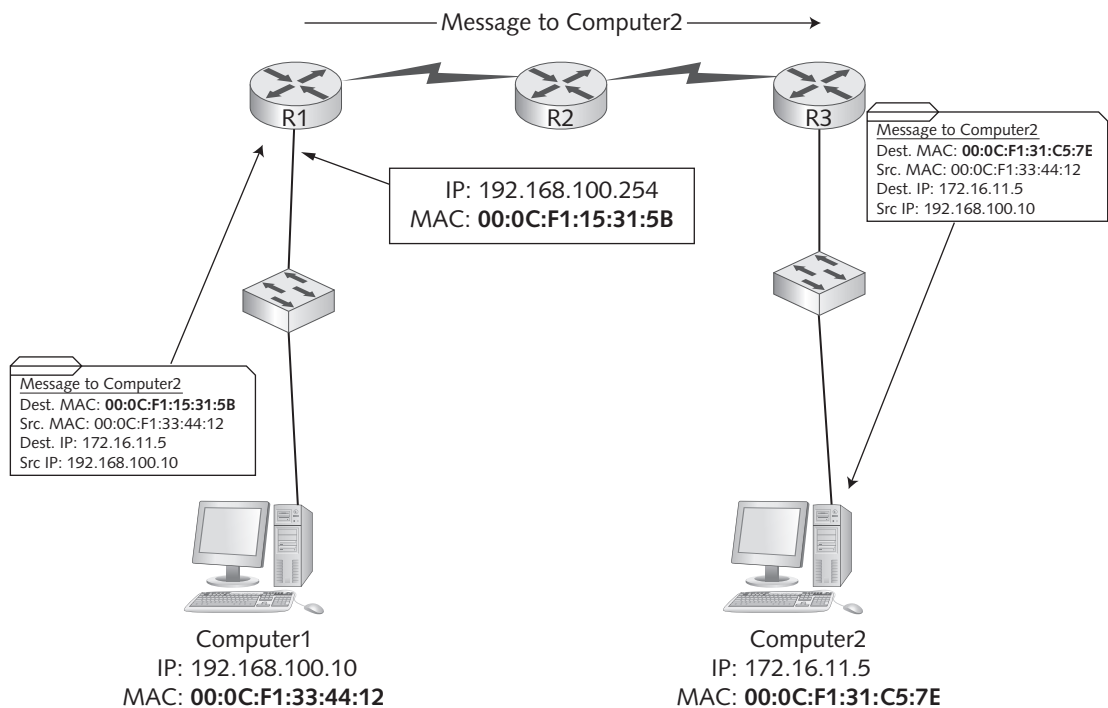


Figure 5-7 How MAC addresses are used in an internetwork

Courtesy of Course Technology/Cengage Learning

message is the MAC address of router R1, but the destination IP address remains the same throughout the journey. Only when the message gets to the destination network does the MAC address become Computer2's address. Notice also that the source MAC address in the frame going from router R3 to Computer2 has changed, showing that the frame is coming from router R3.



Simulation 9: The changing frame header

Internet Control Message Protocol Internet Control Message Protocol (ICMP) is used to send error and control messages between systems or devices. It's an encapsulated IP protocol, meaning it's wrapped in an IP header. In essence, ICMP is just a specialized IP packet with its own header.

ICMP has many message types, but the two most people know are ICMP Echo (sent by the Ping program) and ICMP Reply (sent by the target of the ping). Ping uses ICMP Echo packets to request a response from another computer to verify whether it's available for communication. The response, if received, is an ICMP Reply packet indicating not only that the remote host is reachable, but also how long the message's round trip from sender to receiver took.

The Tracert program uses ICMP Echo packets to determine the route a packet takes through an internetwork. It sends an ICMP Echo packet with the TTL value in the IP header set to 1. When the packet reaches the first router on its way to the destination, the router decrements the TTL value to 0, discards the packet, and sends a TTL-Expired ICMP packet to the sending machine to notify it that the packet expired. Tracert receives the TTL-Expired message containing the router's IP address and then has the address of the first router in the path to the destination. Next, Tracert sends an ICMP Echo packet with a TTL of 2. When the packet gets to the second router, it again expires, and the second router sends a TTL-Expired message. In this way, Tracert discovers the IP address of every router between the source and destination computers.



These uses of ICMP are the most common, but ICMP has more than 20 message types, many of which are now obsolete. To learn more about ICMP message types, see www.iana.org/assignments/icmp-parameters. The Internet Assigned Numbers Authority (IANA) is in charge of everything related to numbers used in Internet protocols.

Internet Protocol Security Internet Protocol Security (IPSec) works with IPv4 to ensure secure delivery of packets. Most OSs now support IPSec as a feature that can be enabled for certain types of communication between specific computers. In other words, IPSec can be used to secure sensitive network transmissions between computers needing the extra security.

IPSec provides security by using authentication and encryption. It authenticates the identity of computers transmitting data with a password or some other form of credentials, and it

encrypts data so that if packets are captured, the data will be unintelligible. IPSec requires additional network and computer resources, so it should be enabled only for highly sensitive communication and in environments where security risks are high.



Hands-On Project 5-3: Examining the Internetwork Layer

Time Required: 10 minutes

Objective: Capture packets and view the Internetwork layer.

Required Tools/Equipment: Your classroom computer with Wireshark installed

Description: In this project, you capture some ICMP packets and examine the IP header information.

1. If necessary, log on to your computer as **NetAdmin**.
 2. Start Wireshark and click **Capture Options**. In the Capture Filter text box, type **icmp**, and then click **Start**.
 3. Ping your default gateway or DNS server. If you don't remember these IP addresses, use the `ipconfig` command to display them.
 4. In Wireshark, click the **Stop the running live capture** toolbar icon to stop the capture. Scroll up to the first packet summary line, if necessary.
 5. Click a packet summary in the top pane with **ICMP** in the Protocol field and **Echo (ping)** request in the Info field.
 6. Click to expand the line beginning with **Internet Protocol**. In the header, find the fields discussed previously in "Protocols at the Internetwork Layer." Write the values in the Version, Time to live, Protocol, and Checksum fields:
-
-
7. Click to expand the line beginning with **Internet Control Message Protocol**, and examine the information in the ICMP header. The Type field specifies the type of ICMP message. The data portion of the ICMP field is simply a string of letters. It doesn't matter what's in the data part of the Echo Request message, as long as the reply contains the same data.
 8. Exit Wireshark and click **Quit without Saving** when prompted.
 9. Close all open windows, but leave your computer running for the next project.



Hands-On Project 5-4: Capturing ARP and ICMP Packets

Time Required: 10 minutes

Objective: Use Wireshark to capture packets created by the Tracert program.

Required Tools/Equipment: Classroom computers with Wireshark installed



Description: In this project, you use Wireshark to capture ARP and ICMP packets generated by the Tracert program.

1. If necessary, log on to your computer as **NetAdmin** and open a command prompt window.
2. Type **arp -d** and press **Enter** to clear your ARP cache.
3. Start Wireshark and click **Capture Options**. In the Capture Filter text box, type **arp or icmp**, and then click **Start**.
4. At the command prompt, type **tracert books.tomsho.com** and press **Enter**. When Tracert is finished, click the **Stop the running live capture** toolbar icon in Wireshark to stop the capture. Scroll to the first packet summary line, if necessary.
5. Find the ARP packets your computer has generated by looking in the Info column for “Who has *A.B.C.D*, Tell 192.168.100.*XX*” (replacing *A.B.C.D* with the address of your default gateway and *XX* with your student number). Click this packet summary line.
6. Notice that the Dst (for destination) address is **ff:ff:ff:ff:ff:ff**, indicating a broadcast. In the middle pane, click to expand the **Ethernet II** line. Notice that the Type field is **ARP (0x806)**, which tells the Network access layer which Internetwork-layer protocol should receive the packet. Click again to collapse this line.
7. Click to expand the **Address Resolution Protocol (request)** line. Examine the information in the ARP header. The ARP message has fields to indicate what technology is used in the Network access layer (Ethernet) and the protocol type that needs the MAC address (IP, in this case). Click again to collapse this line.
8. Next, in the top pane, click the ARP reply message immediately following the ARP request. The Info column should be similar to “*A.B.C.D* is at 0A:1B:2C:3D:4E:5F.” The MAC address in the ARP reply is the MAC address of your default gateway. Explore the Network access and Internetwork headers for this frame. (*Note:* You might also find an ARP request and ARP reply for your DNS server if it’s in the same network as your computer.)
9. In the top pane, click the first **ICMP Echo (ping) request** message from your computer to the destination computer at *books.tomsho.com*. The IP address should be 67.210.126.125, but IP addresses can change, so it might be different.
10. In the middle pane, click to expand the **Internet Protocol** line. Notice that the value in the Time to live line is 1.
11. In the top pane, click the **ICMP Time-to-live exceeded** message that follows the Ping request. This message was generated by the first router en route to *books.tomsho.com*. Notice that the source address is the address of your default gateway.
12. Find the next ICMP Echo (Ping) Request message and view the TTL value. Tracert sends three Echo Request messages for each TTL value, so the first three Echo Request messages have a TTL value of 1. Find the fourth ICMP Echo (Ping) Request message and view the TTL value, which should be 2. The Time-to-live exceeded message following it is from the next router down the line. Tracert follows this pattern until reaching the destination device (*books.tomsho.com*).
13. Exit Wireshark, but leave the command prompt window open if you’re continuing to the next project.



Hands-On Project 5-5: Using the `arp` Command

Time Required: 10 minutes

Objective: Use the `arp` command to view and change the ARP cache.

Required Tools/Equipment: Your classroom computer.

Description: In this project, you use the `arp` command to view and then delete the ARP cache, and you use the `ping` command to generate ARP cache entries.

1. If necessary, log on to your computer as **NetAdmin**.
2. Some tasks require opening the command prompt window as an administrator. To do this, click **Start**, point to **All Programs, Accessories**, right-click **Command Prompt**, and click **Run as administrator**. In the User Account Control (UAC) message box, click **Yes**.
3. To display your current ARP cache, type `arp -a` and press **Enter**. A list of IP address/MAC address pairs is displayed. The Type field (third column) indicates whether the entry is static or dynamic. Windows 7 generates static entries automatically, but dynamic entries are generated by network communication. If you don't have any entries, the message "No ARP Entries Found" is displayed.
4. To delete the ARP cache, type `arp -d` and press **Enter**. To verify that the entries have been deleted, type `arp -a` and press **Enter** again.
5. Type `ping 192.168.100.XX` (replacing `XX` with the IP address of another computer in your network) and press **Enter**. Display your ARP cache again. You should see the IP address you pinged along with its MAC address.
6. Clear the ARP cache again. Type `ping www.course.com` and press **Enter**. Display the ARP cache again. You'll probably see two new entries in your ARP cache. On the following lines, list these two new entries and state why they were generated. Compare the entries in the ARP cache with the IP address of `www.course.com`. Do you see this IP address in the ARP cache? Write your answer along with an explanation.

7. Leave the command prompt window open for the next project.



Hands-On Project 5-6: Using the Netstat Program

Time Required: 10 minutes

Objective: Use the Netstat program to view network interface and IP protocol status and statistics.

Required Tools/Equipment: Your classroom computer

Description: In this project, you use Netstat to view statistics about your network interface and the IP protocol. Then you generate traffic with Ping and Tracert to see the statistics of different packet types change.



1. If necessary, log on to your computer as **NetAdmin** and open a command prompt window.
2. To display statistics about your Ethernet interface, type **netstat -e** and press **Enter**. These statistics include the number of bytes and packets received and sent through the Ethernet interface. If any errors are indicated in the display, you might have problems with your network connection that are slowing the network down. If the error packets approach 1% of the total number of packets, something is probably wrong with your NIC or physical interface.
3. To see statistics for all protocols, type **netstat -s** and press **Enter**. To limit the display to just IP statistics, type **netstat -ps IP** and press **Enter**.
4. To see your network statistics updated every 5 seconds, type **netstat -ps IP 5** and press **Enter**. Press **Ctrl+C** to stop the program.
5. To display ICMP information, type **netstat -ps ICMP** and press **Enter**. A variety of ICMP message types are displayed along with how many of each type of message were received and sent. Most, if not all, will be Echo and Echo Reply messages.
6. Type **ping 5.5.5.5** and press **Enter**. This command should generate ICMP Destination Unreachable messages. To see whether the number of Destination Unreachable messages has increased, type **netstat -ps ICMP** and press **Enter**.
7. The ICMP TTL-Expired messages used in Tracert are called Time Exceeded messages in Netstat. Type **tracert books.tomsho.com** and press **Enter**. To see whether the number of Time Exceeded messages has increased, type **netstat -ps ICMP** and press **Enter**.
8. To display your computer's routing table, type **netstat -r** and press **Enter**. Every computer has a routing table it uses to decide which interface to send packets to. The first entry lists the network destination as 0.0.0.0, which is the entry for your default gateway.
9. Leave the command prompt window open for the next project.

Role of the Transport Layer

Without the Transport layer in the TCP/IP protocol suite, large internetworks would be in big trouble. So many things can go wrong with complex, constantly changing networks that without some reliability measures, successful transfers of large amounts of data would be the exception rather than the norm. In environments such as the Internet, using only connectionless protocols simply wouldn't work. The more robust protocols in the Transport layer provide the reliability needed to handle the unpredictable nature of the Internet (or any large internetwork, for that matter).

The Transport layer has two protocols. **Transmission Control Protocol (TCP)** is connection oriented and designed for reliable transfer of information in complex internetworks. **User Datagram Protocol (UDP)** is connectionless and designed for efficient communication of generally small amounts of data. Both protocols perform the following tasks:

- Work with segments.
- Provide a means to identify the source and destination applications involved in a communication.
- Protect data in the segment with a checksum.

Working with Segments Both Transport-layer protocols work with units of data called segments. For outgoing data, in which the Application-layer protocol requires the services of a Transport-layer protocol, the Application layer passes data to TCP or UDP, depending on which protocol it was designed to use. Both TCP and UDP add a header to the data to make it a segment. The Transport-layer protocol then passes the segment to the Internetwork-layer protocol, usually IP. With incoming data, the Internetwork-layer protocol deencapsulates the packet and forwards the resulting segment to the Transport-layer protocol. The Transport-layer protocol processes the segment, deencapsulates it, and sends the resulting data up to the Application layer.

Identifying Source and Destination Applications Have you ever wondered how your computer keeps track of the myriad network applications you run? At any time, you might be running a Web browser, an e-mail application, and a chat program and have a file open on a file server. When one of these applications receives data from the network, a frame is received by the NIC, which sends a packet up to the IP protocol, which then sends a segment to TCP or UDP. Now what? Eventually, data that's received must go to an application or a network service.

The Transport-layer header provides the information needed to determine the application the received data is sent to. TCP and UDP use **port numbers** to specify the source and destination Application-layer protocols. Using an envelope analogy, if the IP address is the zip code and the street number is the MAC address, the port number specifies the person in the house who should read the letter. In other words, the MAC address and IP address get the packet to the computer, and the port number gets the data to the application or service.

The IANA assigns a dedicated port number to every well-known network service. For example, Web servers are assigned port 80, so when your computer formats a message to a Web server, the destination port number in the TCP header is 80. Likewise, when your e-mail application requests messages from your mail server, it sends the request to port 110, the Post Office Protocol (POP3) port number. Most client applications are assigned a random port number when they make a request to a server. So when you start a Web browser, for example, the Web browser window is assigned a port number. When the request for a Web page goes out, the source port number in the TCP header contains the number assigned to that Web browser window so that the Web server knows to which port the reply should be sent. If you open another Web browser window or tab, another port number is assigned, and so forth. The port number is a 16-bit value, so you can open as many as 65,000 windows!



TIP

You can see the list of well-known port numbers at www.iana.org/assignments/port-numbers.



NOTE

An IP application that doesn't use a Transport-layer protocol, such as the Ping program and routing protocols, can rely on the Internetwork layer to provide application information. As you've seen, the IP packet header includes the Protocol field for just this purpose.



Protecting Data with a Checksum To protect data integrity, TCP and UDP provide a checksum similar to the CRC in the Network access layer, but the CRC isn't always a perfect mechanism for ensuring that data wasn't corrupted on the way to its destination. Routers and switches have been known to corrupt data, recalculate the CRC code, and send the corrupted data on its way. In this situation, the receiver has no way of knowing the data was corrupted because the CRC was calculated after the corruption. Intermediate devices don't recalculate the checksum in the Transport layer, so if data corruption occurs along the way, the final receiving station detects the checksum error and discards the data. To ensure reliability, calculating a checksum is as far as UDP goes. All other reliability features at the Transport layer are the domain of TCP.

TCP: The Reliable Transport Layer

If an application requires reliable data transfer, it uses TCP as the Transport-layer protocol. TCP provides reliability with the following features that aren't available in UDP:

- Establishing a connection
- Segmenting large chunks of data
- Ensuring flow control with acknowledgements

Each feature relies on TCP being a connection-oriented protocol. TCP establishes a connection with the destination, data is transferred, and the connection is broken.

Establishing a Connection: The TCP Handshake Establishing a connection with TCP is similar to making a phone call. You dial the number and wait for your party to answer, usually with a greeting. The caller then states his or her name and says who he or she wants to talk to. If everything is agreeable, a conversation begins.

A TCP session begins when a client sends a TCP synchronization (SYN) segment to the destination device, usually a server. A destination port number (typically a well-known port, such as 80) is specified, and a source port number is assigned dynamically. When the server receives the SYN segment, it usually responds by sending one of two segments: an acknowledgement-synchronization (ACK-SYN) segment or a reset connection (RST) segment. If an RST segment is returned, the server refused the request to open a session, possibly because the destination port is unknown. If an ACK-SYN segment is returned, the client completes the **three-way handshake** by sending an ACK segment back to the server. The client is then ready to begin sending or requesting data. You capture and examine a three-way handshake in Challenge Lab 5-1.

Segmenting Data One safeguard TCP provides is segmenting data before sending it to the Internetwork layer and reassembling the data at the destination before sending it up to the Application layer. When TCP receives data from the Application layer, the size of the data might be too large to send to the Internetwork layer in one piece. Remember that Ethernet can send only frames that are a maximum of 1518 bytes. It's TCP's job to break the data into smaller segments before handing each segment to the Internetwork layer. Each segment is labeled with a sequence number so that if segments arrive at the destination out of order, they can be reassembled in the correct order by using the sequence number. Programs that work with large amounts of data, such as Web browsers and file transfer programs, use Application-layer protocols that work with TCP for this reason. Applications

that work with small amounts of data can use UDP, which doesn't disassemble or reassemble data.

Ensuring Flow Control with Acknowledgements Another role of TCP is to provide **flow control**, which prevents a destination from becoming overwhelmed by data, resulting in dropped packets. TCP does this by establishing a maximum number of bytes, called the window size, that can be sent before the destination must acknowledge receipt of the data. If a sending machine hasn't received an acknowledgement before sending the number of bytes established by the window size, it stops sending data. If no acknowledgement is received in a specified timeout period, the sender retransmits the data from the point at which an acknowledgement was last received.

Role of the Application Layer

The Application layer provides network services to user applications that access network resources. For example, when you run Microsoft Word and need to open a file on a network server, Word contacts Client for Microsoft Networks, an Application-layer service, which provides the details of accessing files on the server. Client for Microsoft Networks implements an Application-layer protocol called Server Message Block (SMB), which is also known as Common Internet File System (CIFS). Linux uses NFS and Samba file-sharing Application-layer protocols.

In some cases, the Application-layer protocol or service is built into the user application, as with a Web browser or e-mail client. For example, a Web browser contains the software that implements Hypertext Transfer Protocol (HTTP). Whether the Application-layer protocol is implemented by the user application or by a network service, the process is the same: When data is ready to be sent, it's transferred from the Application-layer protocol to the Transport layer and down the protocol stack until a frame is transmitted as bits to the network medium.

Application-layer protocols also provide authentication and data-formatting services as needed. For example, if a client attempts to access a server that's password protected, the Application layer is responsible for handling the exchange of packets that allow user logon. If data needs to be formatted or translated in some way for the user application, as with some types of data encryption, the Application layer provides that service for user applications. For example, when you connect to a secure Web site with HTTPS, the authentication and encryption that occur with HTTPS are Application-layer functions.



Some functions of the TCP/IP model's Application layer are separated into additional layers in the OSI model discussed in Chapter 6. For example, the Session layer handles network logon, and the Presentation layer handles data encryption and decryption.

With most Application-layer protocols, both a client and a server version exist. For HTTP, the client is a Web browser and the server is a Web server, such as Microsoft Internet Information Services (IIS) or the popular Apache Web Server that's often used on Linux servers. Client for Microsoft Networks has File and Printer Sharing for Microsoft Networks as its server counterpart.

Most Application-layer protocols facilitate a client's access to data, such as an e-mail message or a document. However, the Application layer contains some specialized protocols for



making a network easier to use and configure. Examples include protocols for name resolution and dynamic IP address assignment. Several Application-layer protocols are discussed in more detail in the next sections, but to sum up, the Application layer provides these functions:

- Access by applications to network services
- Client/server data access
- Name resolution
- Dynamic address assignment
- Authentication/user logon
- Data formatting and translation

HTTP: Protocol of the World Wide Web HTTP is the protocol Web browsers use to access data on the World Wide Web. Originally, its main purpose was simply to transfer static Web pages written in HTML. Now HTTP is also used for general file transfer and downloading and displaying multimedia files. Because it's often used to transfer large amounts of data over the Internet, it uses TCP as its Transport-layer protocol, and the default TCP port number is 80. Figure 5-8 shows a typical HTTP message as it might look at the Internetwork layer before going to the Network access layer to be framed. The TCP header contains important source and destination port numbers.




Figure 5-8 An HTTP message

Courtesy of Course Technology/Cengage Learning

POP3, IMAP, and SMTP: E-mail Protocols E-mail clients use the **Post Office Protocol version 3 (POP3)** protocol to download incoming messages from an e-mail server to their local desktops. POP3 clients download e-mail from the mail server running at the user's ISP, and these message are then deleted from the server. POP3 uses TCP port 110.

Internet Message Access Protocol (IMAP) has advanced message controls, including the capability to manage messages locally yet store them on a server, plus numerous fault-tolerance features. IMAP downloads only e-mail headers initially, and then downloads the message body and attachments when the message is selected. IMAP uses TCP port 143.

Simple Mail Transfer Protocol (SMTP) is the standard protocol for sending e-mail over the Internet. POP3 is used to retrieve e-mail, and SMTP is used to send it. SMTP uses TCP port 25. All three e-mail protocols use the TCP Transport-layer protocol to ensure reliable delivery of large messages.

Dynamic Host Configuration Protocol Some drawbacks of using TCP/IP in a large network include detailed configuration of devices and keeping track of assigned

addresses and to which machine they're assigned. To make these tasks easier, **Dynamic Host Configuration Protocol (DHCP)** was developed.

To use DHCP, a server must be configured with a block of available IP addresses and other IP address configuration information. To receive its IP address from the server, each computer must be configured to request its address configuration. A computer requests IP address information from the DHCP server in the form of a broadcast message. Each time a computer requests an address, the server assigns one until it has no more addresses to assign.

A computer just leases the address the server assigns to it. The network administrator defines the lease time when the DHCP server is configured. This time can be as little as a few minutes to an infinite period, in which case the lease never expires. A typical lease time is one day or a few days. When 50% of the lease time has elapsed, the computer attempts to renew the lease from the same DHCP server that responded to the initial DHCP request. If no response is received, the computer waits until 87.5% of the lease time has elapsed; at that point, a broadcast DHCP renewal request is sent. If no response has been received when the lease expires, the computer broadcasts a DHCP request for a new IP address. If no DHCP server responds, one of two things happens: TCP/IP stops functioning, or the computer assigns itself an address from a special range of addresses beginning with 169.254.

These special addresses are reserved for **Automatic Private IP Addressing (APIPA)**. An address in the APIPA range is assigned automatically to an APIPA-enabled computer when an IP address is requested via DHCP but no DHCP server responds to the request. Using APIPA rather than a DHCP server to assign addresses is recommended only for small networks that aren't attached to the Internet.

A major benefit of using DHCP is how easily computers can be moved. When a computer is moved to a new network segment and turned on, it requests its configuration from a DHCP server on that segment. This type of address assignment shouldn't be used for systems that require a static address, such as Web servers, DNS servers, and DHCP servers, because computers with these network services are usually expected to maintain the same IP address.

DHCP uses the UDP Transport-layer protocol because DHCP servers are usually located on the same network as the DHCP client, and DHCP messages are short in length. Recall that UDP is a connectionless protocol and provides few reliability features, so it works best when the amount of data in each transaction is small.



All major OSs include a DHCP client service, and most server OSs and routers include the DHCP server component.

Domain Name System Domain Name System (DNS) is a name-to-address resolution protocol that keeps a list of computer names and their IP addresses. Through a correctly configured workstation, a user can use a computer's name—for instance, Server1 or www.course.com—rather than a numerical address, such as 207.46.134.189, to communicate with the computer. For example, when you enter “www.course.com” in your Web browser's address box, the Web browser contacts the DNS client service on your computer. The



DNS client contacts the DNS server specified in your OS's IP configuration and requests that the name "www.course.com" be resolved to an IP address. The DNS server responds with the IP address assigned to the computer named www at the course.com domain. From there, using the IP address returned, your Web browser application can contact the Web server to request a Web page.

DNS uses the UDP Transport-layer protocol because DNS messages usually consist of a single packet of data, so there's no need for the reliability measures TCP offers. The DNS system used throughout the Internet is organized as a treelike hierarchy (see Figure 5-9). The tree consists of these domain levels: root, top, second, subdomain, and host. All levels below the root level have branches, each of which has a name. When you put all the names of a branch together, separated by a period, you have the **fully qualified domain name (FQDN)** of the network resource, such as *www.course.com*.

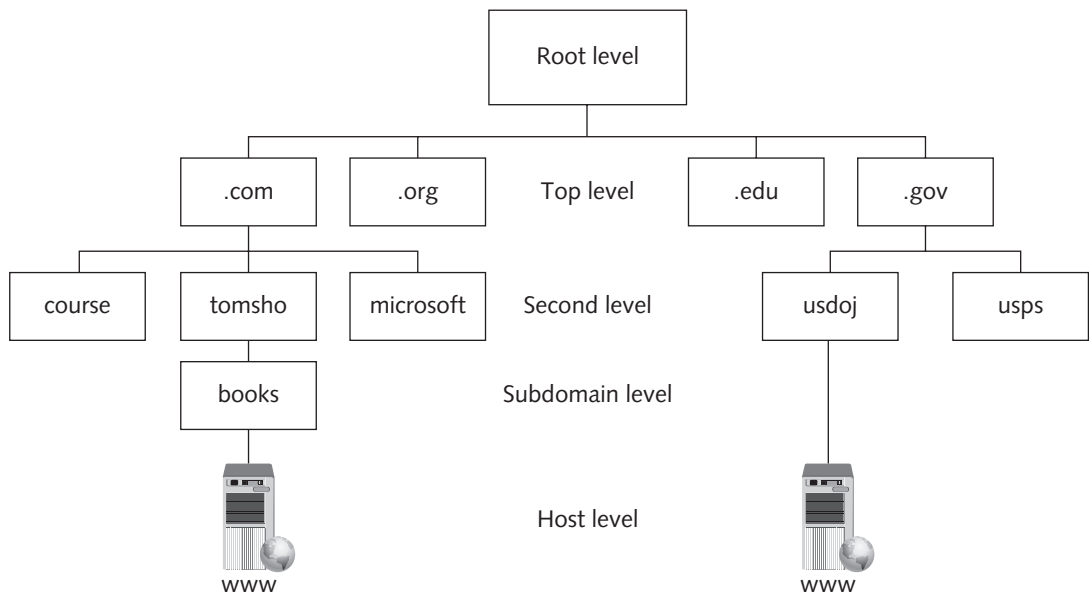


Figure 5-9 DNS hierarchical tree structure

Courtesy of Course Technology/Cengage Learning

The top-level domains are organized into categories—such as commercial (.com), nonprofit organizations (.org), government (.gov), and education (.edu)—or country of origin, indicated by a two-letter country code. The second-level domains are usually the name of a company or an institution. The subdomain level is optional and can consist of several names separated by a period. An example is a department or branch of an organization. Finally, the host level represents individual computers hosting network services. For example, in *www.books.tomsho.com*, com is the top-level domain name, tomsho is the second-level domain, books is the subdomain, and www is the hostname.

Because of the hierarchical nature of DNS, not every DNS server needs to maintain a database of all domain names, computer names, and IP addresses for the entire Internet. Most DNS servers maintain addresses for the domain in which they're installed. The domain might be a single secondary-level domain, such as xyzcorp.com, or if you own a business

hosting Web sites for other companies, you might maintain hundreds or thousands of domains, but that number is small compared with the entire Internet.

When a DNS server is installed, the administrator creates one or more domain names or zones. A zone is named by using the second-level and top-level domain names and the sub-domain, if necessary. Most of the information a DNS zone contains consists of hostname/IP address pairs, called host records. An administrator can create host records manually, and they can be created dynamically with Dynamic DNS (DDNS), which enables a computer to create its own DNS record.

In addition to host records, a DNS server database is loaded with a list of IP addresses that point to root servers around the world. These servers supply the addresses of top-level domain servers, which are used to provide addresses of second-level domain servers. This hierarchical organization allows any DNS client anywhere in the world to access the DNS servers for any domain.



TIP

You can view a map of the root servers around the world at <http://public-root.com/root-server-locations.htm>.

To speed up communication, DNS clients in most OSs maintain a DNS cache, much as ARP maintains an ARP cache. The DNS cache stores name and IP address pairs along with other pertinent DNS data for names that have been resolved recently. This cache prevents your DNS client from having to request that a DNS server do a name lookup for a name that was resolved recently. Additionally, the cache contains a text file called *Hosts*, which stores name and IP address pairs. This file usually contains only the name “localhost” mapped to 127.0.0.1, but you can add entries manually by editing the file. If there are computers you access frequently by name and you don’t expect their addresses to change, you can add an entry for these computers in the *Hosts* file, thereby preventing a network DNS lookup from occurring when you access them. You examine the *Hosts* file in Hands-On Project 5-7.

Many other Application-layer protocols work with TCP or UDP, but the protocols discussed in this chapter cover the most commonly used ones in most networks. For all these protocols to work, TCP/IP needs a correctly addressed network, so in the next section, you turn your attention to IP addressing.



Hands-On Project 5-7: Working with DNS Tools

Time Required: 10 minutes

Objective: Use Ipconfig and Nslookup to work with DNS.

Required Tools/Equipment: Your classroom computer

Description: In this project, you use Ipconfig to display and delete your DNS cache, and then view your *Hosts* file. You also use Nslookup to query your DNS server.

1. If necessary, log on to your computer as **NetAdmin** and open a command prompt window.



2. To see a list of recent DNS lookups, type **ipconfig /displaydns** and press **Enter**. To delete the entries, type **ipconfig /flushdns** and press **Enter**. Display the DNS cache again. Unless there are entries in your Hosts file, you should get the message “Could not display the DNS Resolver Cache.”



At the command prompt, you can press the up and down arrow keys to access recent commands you have entered.

TIP

3. To perform a DNS lookup, type **ping www.course.com** and press **Enter**. Display the DNS cache again.
4. You should see a DNS record for *www.course.com* that includes the IP address and other information. Another field in the DNS cache is a TTL value, which is different from the TTL in an IP packet. This DNS TTL value is sent by the DNS server maintaining the *www.course.com* record. It’s measured in seconds and tells your DNS client how long to cache the DNS record as a safeguard against clients holding on to DNS records whose IP addresses might have changed.
5. To open your computer’s Hosts file, click **Start**, point to **All Programs, Accessories**, right-click **Notepad**, and click **Run as administrator**. In the UAC message box, click **Yes**. Click **File, Open** from the Notepad menu. In the Open dialog box, navigate to **C:\Windows\System32\Drivers\Etc**. In the File type drop-down list, click **All Files**. Double-click the **hosts** file to open it.
6. After the last line in the file, type **67.210.126.125 books**, and then save the file and exit Notepad.
7. At the command prompt, type **ipconfig /displaydns** and press **Enter** to see that the entry is in your DNS cache. Type **ping books** and press **Enter**. Delete the DNS cache (see Step 2), and then display the DNS cache again. Notice that the books entry remains in the cache because the Hosts file data always stays in the cache.
8. Type **nslookup www.course.com** and press **Enter**. Your DNS server’s name and IP address are displayed, along with the name and IP address of *www.course.com*. You use Nslookup to look up a host’s IP address without actually communicating with it.
9. Type **nslookup** and press **Enter**. You enter Nslookup’s interactive mode. Type **www.google.com** and press **Enter**. Notice that more than one address is returned along with one or more aliases (other names that *www.google.com* goes by). Type **www.google.com** again (or press the up arrow to repeat the last line you typed) and press **Enter**. You should see the IP addresses returned in a different order. (If you don’t, keep trying, and the order will change.) The *www.google.com* page can be reached by a number of different IP addresses, and the addresses are returned in a different order so that a different server is used each time, which is called load balancing.
10. Type **198.60.123.100** and press **Enter**. Nslookup is also used to do reverse lookups, in which the IP address is given and the hostname is returned.
11. To set the DNS server that Nslookup uses to a public DNS server run by Google, type **server 8.8.8.8** and press **Enter**. Type **www.microsoft.com** and press **Enter**. If you’re ever concerned that your DNS server isn’t working correctly, you can test it with Nslookup

and compare the results of your DNS server with the results from another server, such as Google's.

12. Leave the command prompt window open for the next project.



Hands-On Project 5-8: Working with the DHCP Client

Time Required: 10 minutes

Objective: Use Ipconfig to work with your DHCP client.

Required Tools/Equipment: Your classroom computer

Description: In this project, you change your IP settings to use DHCP and then see how to work with DHCP by using Ipconfig.

1. If necessary, log on to your computer as **NetAdmin**.
 2. Open the Network and Sharing Center by clicking **Start, Control Panel**. Under Network and Internet, click **View network status and tasks**.
 3. In the left pane of the Network and Sharing Center, click **Change adapter settings**. Right-click **Local Area Connection** and click **Properties** to open the Local Area Connection Properties dialog box.
 4. Click **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**. If your IP address settings were set manually in Hands-On Project 5-1, write these settings on the following lines. At the end of this project, you set your IP address again by using these settings.
-
5. Click **Obtain an IP address automatically**. Click **OK** and then **Close**.
 6. Open a command prompt window, if necessary. Type **ipconfig /all** and press **Enter** to view detailed IP configuration information. Under Ethernet adapter Local Area Connection, you see information about DHCP, including its status (enabled or not), the DHCP server's IP address, and lease information.
 7. Occasionally, you might need to force your computer to renew its lease (for example, if changes are made on the DHCP server and you need to get the newest IP configuration). To renew a DHCP lease, type **ipconfig /renew** and press **Enter**. Display your detailed IP configuration again to see that the lease information has changed.
 8. To release your IP address configuration, type **ipconfig /release** and press **Enter**. This command's output shows that your IP configuration has been deleted. To request a new IP address configuration, type **ipconfig /renew** and press **Enter**. (Note that you might not get the same IP address you had before.) Using these commands can help you troubleshoot DHCP-related problems.
 9. Close the command prompt window, and set your IP configuration to the values you wrote down in Step 4. Stay logged on for the challenge labs at the end of the chapter.



IP Addressing

As you've learned, IP is responsible for addressing and routing in the TCP/IP environment. IP addresses are 32-bit (4-byte) logical addresses. The 32 bits are grouped into four 8-bit **octets**, and each octet is represented by a decimal number from 0 to 255. The four decimal numbers are separated by periods in a format called **dotted decimal notation**, as in 172.24.208.192.

As discussed, an IP address is divided into two distinct parts. One part designates which logical network the computer is a part of; the remainder of the address represents the host ID for that computer. For example, a computer with the address 172.24.208.192 resides on the 172.24 network, and its host ID is 208.192. In this case, the complete network address is expressed as 172.24.0.0, with the trailing zeros indicating a network address because a host ID can't be 0. The computer next to it might have the address 172.24.18.26, but even though their host IDs are quite different, both computers are on the same network because they share the same network address (172.24).

You can determine how many host addresses are in a network (the **address space**) by looking at the host ID's size. In the address 172.24.208.192, for example, the host ID occupies the third and fourth octets of the address, which allows 16 bits for the address space. With 16 bits of address space, you can use the formula 2^{16} , which yields 65,536. An address with a network number of 201.55.66 leaves only one octet (or 8 bits) for the host ID, or 2^8 , which yields 256 possible addresses.

IP Address Classes

IP addresses are categorized in ranges referred to as Classes A, B, C, D, or E. Only IP addresses in the A, B, and C classes are available for host assignment. Although the IP address class system has been somewhat superseded by a more flexible way to manage IP addresses, called Classless Interdomain Routing (CIDR, discussed later in this chapter in "Classless Interdomain Routing"), the class system is a basis for determining which part of an IP address is the network ID and which part is the host ID. The first octet of an address denotes its class. Note the following facts about IP address classes:

- The value of the first octet for Class A addresses is between 1 and 127. Class A addresses are intended for use by large corporations and governments. An IP address registry assigns the first octet, leaving the last three octets for network administrators to assign to hosts. This allows 24 bits of address space or 16,777,214 hosts per network address. In a Class A IP address such as 21.155.49.211, for example, the network ID is 21.0.0. So the first address in the 21.0.0.0 network is 21.0.0.1, and the last address is 21.255.255.254.
- Class B addresses begin with network IDs between 128 and 191 and are intended for use in medium to large networks. An IP address registry assigns the first two octets, leaving the third and fourth octets available for administrators to assign as host addresses. In the Class B address 172.17.11.4, for example, the network ID is 172.17.0. Having two octets in the host ID allows 65,534 hosts per network address.
- Class C addresses are intended for small networks. An IP address registry assigns the first three octets, ranging from 192 to 223. In the Class C address 211.255.49.254, for example, the network ID is 211.255.49. These networks are limited to 254 hosts per network.

- Class D addresses are reserved for multicasting, in which a packet is addressed so that more than one destination can receive it. Applications using this feature include video-conferencing and streaming media. In a Class D address, the first octet is in the range 224 to 239. Class D addresses can't be used to assign IP addresses to host computers.
- Class E addresses have a value from 240 to 255 in the first octet. This range of addresses is reserved for experimental use and can't be used for address assignment.

A couple of notes about this list: First, if you did your math, you would see that a Class C address provides 2^8 bits of address space, which yields 256 addresses, not 254. Note, too, that the number of addresses specified for Classes A, B, and C are two fewer than the address space suggests. This discrepancy happens because each network has two reserved addresses: the address in which all host ID bits are binary 0s and the address in which all host ID bits are binary 1s. For example, all the host bits in address 198.44.19.0 are binary 0s, and this address represents the network number and can't be assigned to a computer. The host bits in address 198.44.19.255 are binary 1s; this address is the broadcast address for the 198.44.19.0 network and can't be assigned to a computer.

The other note concerns the 127.0.0.0 network. Although technically a Class A address, it's reserved for the **loopback address**, which always refers to the local computer and is used to test the functioning of TCP/IP. A packet with a destination address starting with 127 is sent to the local device without reaching the network medium. Likewise, the reserved name **localhost** always corresponds to the IP address 127.0.0.1 so that a local machine can always be referenced by this name.



Even though localhost and the loopback address are usually associated with the address 127.0.0.1, any address in the 127.0.0.0 network (except 127.0.0.0 and 127.255.255.255 in most OSs) references the local machine.

Private IP Addresses

Because of the popularity of TCP/IP and the Internet, unique IP addresses to assign to Internet-accessible devices are almost exhausted. To help alleviate this problem, TCP/IP's technical governing body reserved a series of addresses for private networks—that is, networks whose hosts can't be accessed directly through the Internet. This nonprofit governing body, the Internet Engineering Task Force (IETF; www.ietf.org), is responsible for TCP/IP standards and characteristics. The reserved addresses are as follows:

- Class A addresses beginning with 10 (one Class A network address)
- Class B addresses from 172.16 to 172.31 (16 Class B network addresses)
- Class C addresses from 192.168.0 to 192.168.255 (256 Class C network addresses)

The addresses in these ranges can't be routed across the Internet, which is why any organization can use them to assign IP addresses to their internal hosts. If access to the Internet is necessary, a process called Network Address Translation (NAT) is used, explained next.

IPv6 eliminates the need for private addressing because it provides a 128-bit address space, compared with IPv4's mere 32 bits. You learn more about IPv6 later in this chapter in "Introduction to Internet Protocol Version 6."



Several IP address registries around the world cooperatively manage the total collection of valid IP addresses. Their activities are overseen by the Internet Assigned Numbers Authority (IANA), a nonprofit agency responsible for Internet addressing and address management.

Network Address Translation

Although subnetting can alleviate the IP address shortage problem, as you learn later in “Subnet Masks,” it simply makes more efficient use of existing addresses. **Network Address Translation (NAT)** helps more by allowing an organization to use private IP addresses while connected to the Internet. Recall that there are three ranges of private IP addresses (one range for each class), and these addresses can’t be used as the source or destination address in a packet on the Internet.

Anyone can use private IP addresses for address assignment to internal computers and devices, and because the addresses aren’t sent to the Internet, there’s no address conflict. What if you want your computers to have access to the Internet, however? That’s where NAT comes in. An organization can, for example, assign all its workstations’ addresses in the 10.x.x.x private network. Say an organization has 1000 workstations. Although these addresses can’t be used on the Internet, the NAT process translates a workstation address (as a packet leaves the corporate network) into a valid public Internet address. When data returns to the workstation, the address is translated back to the original 10.x.x.x address. NAT is usually handled by a network device that connects the organization to the Internet, such as a router. As shown in Figure 5-10, when station 10.0.0.1 sends a packet to the Internet, the NAT router intercepts the packet and replaces its source address with 198.60.123.101 (a public Internet address). When a reply comes back addressed to 198.60.123.101, the NAT router replaces the destination address with 10.0.0.1.

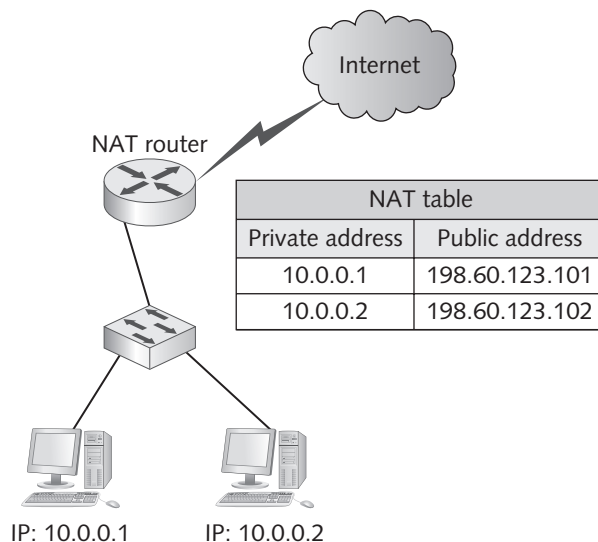


Figure 5-10 Private addresses are translated to public addresses with NAT

Courtesy of Course Technology/Cengage Learning

This process allows any number of companies to use private IP addresses in their own network but requires a public IP address only when a workstation attempts to access the Internet. NAT reduces the number of public IP addresses needed because a public address is required only if a computer accesses the Internet. NAT does have a drawback, in that one public address is required for every private address. However, it's usually used only for Web servers and other devices that must be accessed through the Internet.

An extension of NAT, called **Port Address Translation (PAT)**, allows several hundred workstations to access the Internet with a single public Internet address. This process relies on each packet containing not only source and destination IP addresses, but also source and destination TCP or UDP port numbers. With PAT, the address is translated into a single public IP address for all workstations, but a different source port number (which can be any value from 1024 to 65,535) is used for each communication session, allowing a NAT device to differentiate between workstations. The typical router used in home and small business networks is already configured to use PAT.

Figure 5-11 shows an example of how PAT is used. Notice that the public address is the same for both entries; only the port number differs. When an Internet server responds to 198.60.123.100 on port 3105, however, the router knows to translate the destination address in the packet to 10.0.0.2 port 12441. Notice also that the public address in the NAT/PAT table is the same as the router's Internet-connected interface. Although this configuration isn't necessary, it's common in home/small office routers. In Simulation 10, you see an animation of how PAT works.

NAT/PAT table	
Private address: Port	Public address: Port
10.0.0.1:2562	198.60.123.100:5311
10.0.0.2:12441	198.60.123.100:3105




Figure 5-11 PAT uses the port number to allow using a single public IP address

Courtesy of Course Technology/Cengage Learning



Simulation 10: Demonstrating NAT/PAT



For an excellent tutorial on NAT, see www.howstuffworks.com/nat.htm.

Classless Interdomain Routing

As mentioned, addressing by class has been superseded by a more flexible addressing method. To use all available addresses more efficiently, a different addressing scheme called **Classless Interdomain Routing (CIDR)** is now used on the Internet. With this scheme, when an address is assigned, the network and host IDs don't always begin and end on octet boundaries according to the IP address class; instead, the network ID and host ID can be almost any number of bits (with the combined number of bits equaling 32, of course).

For example, a Class C address's network ID is 24 bits, and its host ID is 8 bits. Using CIDR, an address registry can assign an address with a network ID of 26 bits, leaving 6 bits for the host ID. This technique involves reallocating bits from the host portion of the address to create two or more network numbers. Put another way, one larger network is broken into two or more smaller networks, or **subnets**. **Subnetting** is the process of dividing a single network address into two or more subnetwork addresses, each with fewer available host IDs than the original network address. Subnetting allows fewer hosts on each network but results in more networks overall.

When CIDR addresses are assigned, a slash denotes the number of bits in the network ID; this format is called "CIDR notation." For example, if your company requires only 30 host addresses to assign to computers attached to the Internet, your ISP might give you the network address 192.203.187.0/27. In a 32-bit IP address, the /27 means the first 27 bits of the IP address designate the network ID, and the last 5 bits designate the host ID. The /27 following the IP address is also referred to as the **IP prefix** (or just "prefix"). CIDR notation is used in documentation and by some systems to specify IP addresses. However, the IP address configuration in most OSs uses a second dotted decimal number called a subnet mask, discussed next.

Subnet Masks

In the early days of IP, a host used the address class to determine which part of the address was its network ID. However, this method doesn't offer the flexibility that CIDR addressing requires. When an IP address is assigned to a computer or other IP device, it's always accompanied by a subnet mask. IP uses an address's **subnet mask** to determine which part of the address denotes the network portion and which part denotes the host. It's a 32-bit number in dotted decimal format consisting of a string of eight or more binary 1s followed by a string of 0s. A binary 1 in the subnet mask signifies that the corresponding bit in the IP address belongs to the network address, and a binary 0 signifies that the corresponding bit in the IP address belongs to the host ID.

Each of the three main address classes has a default subnet mask that uses the decimal number 255 for each octet corresponding to the network ID. (The number 255 is 11111111 in binary and fills all eight bit positions in an IP address octet.) Therefore, the default Class A subnet mask is 255.0.0.0, the default Class B subnet mask is 255.255.0.0, and the default Class C subnet mask is 255.255.255.0.

For example, if a computer has the IP address 153.92.100.10 and the subnet mask 255.255.0.0 (a Class B mask), the network portion is 153.92 and the host portion is 100.10. Using the same address of 153.92.100.10 but with the subnet mask 255.255.255.0, the network portion is now 153.92.100 and the host portion is 10. By altering the subnet mask, the network has been changed from one with 65,534 host addresses to one with only 254 host addresses. At first glance, diminishing the size of the address space might not seem like a good thing. However, if you consider how IANA assigns IP addresses to companies and ISPs, it makes more sense.

For example, a new business expects to have a large presence on the Internet with Web servers, DNS servers, application servers, routers, and so forth. This business contracts with an ISP for its Internet access, and the ISP allocates IP addresses the business can use. The business calculates that eventually about 400 addresses will be needed for its Web servers and other devices. The ISP has no Class C addresses to allocate but does have the Class B network 153.92.0.0; however, a Class B address provides 65,534 addresses! If the ISP allocated the company the IP address 153.92.0.0 with the 255.255.0.0 Class B subnet mask, more than 65,000 addresses would be wasted. After a network address is used in one network, no other network can use any addresses in that network. One solution is for the ISP to allocate the company two network addresses: 153.92.0.0 and 153.92.1.0, both with subnet mask 255.255.255.0. That gives the company two network addresses, each with 254 possible host addresses; more than enough for its needs. This solution leaves network addresses 153.92.2.0/24 through 153.92.255.0/24 available for assignment to other networks. (Remember that in CIDR notation, /24 means a 255.255.255.0 subnet mask.)



All devices on a single logical network (meaning each device can communicate with another device without going through a router—sometimes called a network segment) must share the same network address and, therefore, use the same subnet mask.

Table 5-1 summarizes the three classes of IP address used for IP host assignment.

Table 5-1 IP address class summary

Value of first octet	Class	Default subnet mask	Number of hosts/network
1–127	A	255.0.0.0	16,777,214
128–191	B	255.255.0.0	65,534
192–223	C	255.255.255.0	254

How Is the Subnet Mask Used? The IP address configuration of a computer that's part of an internetwork or has access to the Internet requires at least three elements: IP address, subnet mask, and default gateway. As discussed in Chapter 2, the default gateway is the IP address of the router to which the computer sends packets destined for a different

network. The default gateway address must be in the same network as the computer's IP address.

So how does a computer know from the address whether the destination computer is on a different network? Is the address 172.19.44.211 on a different network from 172.19.46.188? The only way to know is by consulting the subnet mask. Using these two addresses, take a look at the sample network shown in Figure 5-12.




Figure 5-12 Determine the destination computer's network address with the subnet mask
 Courtesy of Course Technology/Cengage Learning

Here's what happens when Computer1 has a packet to send to Computer2:

1. Computer1 must first know its network address. It determines this by doing a logical AND operation between its IP address and subnet mask. A logical AND is an operation between two binary values. AND operations can have the following results:

0 AND 0 = 0
 1 AND 0 = 0
 0 AND 1 = 0
 1 AND 1 = 1

2. The logical AND operation between Computer1's IP address and subnet mask looks like this:

```

10101100.00010011.00101110.10111100 (binary for 172.19.44.211)
      AND
11111111.11111111.11111111.00000000 (binary for 255.255.255.0)
-----
10101100.00010011.00101110.00000000 (binary for 172.19.44.0)
    
```

The resulting network address is 172.19.44.0.

3. The next step is to determine whether Computer2's address is on the same network or a different network. The same AND calculation is done between Computer2's IP address and Computer1's subnet mask. (Computer1 has no way of knowing Computer2's subnet mask.) The resulting network address is 172.19.46.0.
4. Because Computer2 is on a different network, Computer1 knows that the packet must be sent to the router, which forwards it to Computer2's network.

Why Subnet? An IP network or subnetwork can be defined as a group of computers and devices that share the same network portion of their assigned IP addresses and don't have to go through a router to communicate with one another. Another term for an IP network is broadcast domain, explained in Chapter 2. Dividing IP networks into smaller subnetworks is done for a number of reasons:

- As discussed, subnetting usually makes more efficient use of available IP addresses. To look at another example, in the older IP address class system of determining network and host IDs, the smallest network, using a Class C address, consists of 254 hosts. If an organization requires only 30 host addresses, it can be assigned a portion of a Class C address, leaving most of the remaining addresses available for assignment on other networks.
- With subnetting, a company can divide its network into logical groups. When one large network is divided into two or more smaller subnetworks, a router is needed to allow hosts on one subnetwork to communicate with hosts on another subnetwork. A router serves as a natural security barrier between the two subnetworks because access control lists can be configured on a router to restrict the type of network traffic traveling from one subnet to another. Being able to restrict access enables network administrators to, for example, place the Payroll Department computers and servers on their own subnet and disallow computers from other subnets to access any resources in the Payroll Department.
- Subnetting can make network communication more efficient. Because routers don't forward broadcast traffic from one subnet to another, broadcast traffic generated by one subnet of 125 computers is heard and processed only by those 125 computers. If you have a single network of 500 computers, the amount of broadcast traffic is often detrimental to overall network performance. Subnetting into four networks of 125 computers each results in much smaller broadcast domains, which usually improves network performance.

Binary Arithmetic

IP address subnetting and supernetting (covered later in this chapter) is a lot easier if you understand the basics of binary arithmetic. For the purposes of this book, you need to be able to convert decimal numbers to binary numbers and vice versa.

Before you tackle these calculations, you need to review how the decimal numbering system works. It's based on powers of 10 (which is where the word "decimal" comes from, with "dec" meaning "ten"). Ten different symbols, 0 through 9, are used to represent any possible number. Each place in a decimal number can have one of 10 possible values: again, 0 through 9. Furthermore, each place in a decimal number can be expressed as a power of 10. The ones place can be expressed as a number, 0 thru 9, multiplied by 10 raised to the 0 power, or 10^0 . (Any number raised to the 0 power equals 1.) The tens place can be



expressed as a number multiplied by 10 to the 1 power, or 10^1 . The hundreds place can be expressed as a number multiplied by 10^2 , and so on. For example, the decimal number 249 can be expressed as either of the following:

$$2 * 10^2 + 4 * 10^1 + 9 * 10^0 = 249$$

$$2 * 100 + 4 * 10 + 9 * 1 = 249$$

When you see the number 249, you don't think of it in these terms because you grew up using the decimal numbering system, and recognizing the hundreds place, tens place, and ones place happens without conscious effort, as does the multiplication and addition that occurs. However, take a look at this number:

379420841249

A little more thought has to go into recognizing that the 3 represents 300 billion, the 7 represents 70 billion, and so forth. The binary number system works the same way, except everything is governed by twos. Two digits, 0 and 1, represent every possible number, and each place in a binary number is 0 or 1 multiplied by a power of 2. So instead of having the ones place, the tens place, the hundreds place, and so on, you have the ones place, the twos place, the fours place, and so on, based on 2^0 , 2^1 , 2^2 , and so forth. For example, using the same method you used to solve the decimal example, you can express the binary number 101 as either of the following. The numbers in bold are the binary digits.

$$1 * 2^2 + 0 * 2^1 + 1 * 2^0 = 5$$

$$1 * 4 + 0 * 2 + 1 * 1 = 5$$

Converting Decimal to Binary One method of converting decimal to binary is with a process called “successive divisions.” With this method, you divide the decimal number by 2, write down the remainder (which must be 0 or 1), write down the dividend, and repeat until the dividend is 0.

The decimal number 125 is converted to binary in the following example:

125 divided by 2 equals 62, remainder 1
 62 divided by 2 equals 31, remainder 0
 31 divided by 2 equals 15, remainder 1
 15 divided by 2 equals 7, remainder 1
 7 divided by 2 equals 3, remainder 1
 3 divided by 2 equals 1, remainder 1
 1 divided by 2 equals 0, remainder 1

To produce the binary number corresponding to 125, you must then write the digits starting from the bottom of the remainder column and work your way up: 1111101. Now check the work involved. Because you have only seven binary digits and want to have eight total, you pad with zeros on the left, leaving you with 01111101. The exponential expansion of 01111101 is $0 \times 2^7 + 1 \times 2^6 + 1 \times 2^5 + 1 \times 2^4 + 1 \times 2^3 + 1 \times 2^2 + 0 \times 2^1 + 1 \times 2^0$.

Another way to convert from decimal to binary is shown in Table 5-2. The first two rows are the binary and exponent values of each bit position of an 8-bit number. You use 8 bits because in subnetting, most work can be done 8 bits at a time. The third row is what you complete to determine the decimal number's binary representation.

Table 5-2 Decimal-to-binary conversion table

128	64	32	16	8	4	2	1
2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
0	1	1	1	1	1	0	1

To use this method, start with the number you're trying to convert to binary: in this case, 125, which is referred to as the "test number." You compare the test number with the leftmost number in the preceding table (128). If it's equal to or greater than this number, you place a 1 in the column and subtract the number in the column from your test number; otherwise, place a 0 in the column. Remember: Eight binary places or 8 bits can represent only a value up to 255. If you're converting a number greater than 255, simply extend the table to the left (256, 512, and so on). Here's the sequence of steps:

1. 125 is less than 128, so you place a 0 in the column under the 128. The test number remains 125.
2. 125 is greater than 64, so you place a 1 in the column under the 64 and subtract 64 from 125, leaving your new test number as 61.
3. 61 is greater than 32, so you place a 1 in the column under the 32 and subtract 32 from 61, leaving your new test number as 29.
4. 29 is greater than 16, so you place a 1 in the column under the 16 and subtract 16 from 29, leaving your new test number as 13.
5. 13 is greater than 8, so you place a 1 in the column under the 8 and subtract 8 from 13, leaving your new test number as 5.
6. 5 is greater than 4, so you place a 1 in the column under the 4 and subtract 4 from 5, leaving your new test number as 1.
7. 1 is less than 2, so you place a 0 in the column under the 2.
8. 1 is equal to 1, so you place a 1 in the column under the 1 and subtract 1 from 1, leaving your new test number as 0. When your test number is 0, you're done.

Now try this with 199, 221, and 24. You should get the following results:

```
199 = 11000111
221 = 11011101
24 = 00011000
```

Converting Binary to Decimal Using 11010011 as the example, here are the steps:

1. Count the total number of digits in the number (eight digits in 11010011).
2. Subtract one from the total ($8 - 1 = 7$).
3. The result of the subtraction (7) is the power of 2 to associate with the highest exponent for 1 in the number, so in this case, the first 1 in your binary number is multiplied by 2^7 .
4. Convert to exponential notation, using all the digits as multipliers.
5. 11010011, therefore, converts to the following:

$$\begin{aligned}
 11010011 &= 1 * 2^7 + 1 * 2^6 + 0 * 2^5 + 1 * 2^4 + 0 * 2^3 \\
 &\quad + 0 * 2^2 + 1 * 2^1 + 1 * 2^0 \\
 &= 128 + 64 + 0 + 16 + 0 + 0 + 2 + 1 = 211
 \end{aligned}$$



Another way to do this conversion is to use Table 5-2, as you did for the decimal-to-binary conversion. Of course, if your binary number is more than 8 bits, you can simply extend the table to the left as many places as necessary. Here's how to do it: Write your binary number in the third row of the table. For every column with a 1 bit, write down the corresponding decimal number from the first row. For columns with a 0 bit, you can simply skip them or write down a 0. Using the binary number 11010011, you have 1 bits in the 128, 64, 16, 2, and 1 columns. Add these values together, and you get 211. Choose some numbers and practice to make sure you understand how to do this conversion.

Recognizing Bit Patterns Many of the numbers you work with when subnetting have telltale patterns. For instance, subnet masks always consist of consecutive 1s followed by consecutive 0s. You never have, for example, a subnet mask that looks like 10110001 in binary (177 in decimal). A subnet mask always consists of a series of zero or more 1s, followed by a series of zero or more 0s, as Table 5-3 shows. Work on memorizing these correlations so that you can deal with subnet masking problems when you see them later in this chapter.

Table 5-3 Subnet mask bit patterns

Binary	Decimal
00000000	0
10000000	128
11000000	192
11100000	224
11110000	240
11111000	248
11111100	252
11111110	254
11111111	255

Table 5-4 lists the binary place values from the 1s place to the 128s place. It's important to know these values so that you can calculate lists of subnets quickly after you reallocate bits for a new subnet mask.

Table 5-4 Binary place values

Binary	Place value	Calculation
00000001	1s place	2^0
00000010	2s place	2^1
00000100	4s place	2^2
00001000	8s place	2^3
00010000	16s place	2^4
00100000	32s place	2^5
01000000	64s place	2^6
10000000	128s place	2^7

Calculating a Subnet Mask

Before you start creating subnet masks, you need to know under what situations you might need to do so. If you work for an ISP, you might need to subnet the address space your company allocates to its customers, as discussed earlier. You don't want to allocate hundreds of addresses to a company that needs only a few, for example. If you're the network administrator for a large corporate internetwork, you might need to subnet the address space your ISP has assigned to you or subnet the private network addresses you're using so that you can create smaller broadcast domains or logical groupings.

There are usually two approaches to subnetting, and they depend on the answer to these questions: Am I subnetting to provide a network with a certain number of host addresses? Or am I subnetting to provide a network with a certain number of logical subnets? If you're working for an ISP, the answer is usually yes to the first question, and if you're a network administrator for a corporate network, the answer is more likely to be yes to the second question. Sometimes the answer is a combination of both.

Say you have a large internetwork and need to break an IP address space into several subnets. Follow this process:

1. First, decide how many subnets you need. You can figure out the number of subnets needed by seeing how many network cable segments are or will be connected to router interfaces. Each router interface connection indicates a required subnet.
2. Next, decide how many bits you need to meet or exceed the number of required subnets. To calculate this value, use the formula 2^n , with n representing the number of bits you must reallocate from the host ID to the network ID. For example, if your starting network number is the Class B address 172.20.0.0, its default subnet mask is 255.255.0.0, which is your starting point. The number of subnets you create is always a power of 2, so if you need 20 subnets, you must reallocate 5 bits ($2^5 = 32$) because reallocating 4 bits gives you only 24 or 16 subnets.
3. Reallocate bits from the host ID, starting from the most significant host bit (that is, from the left side of the host ID).
4. You must also ensure that you have enough host bits available to assign to computers on each subnet. To determine the number of host addresses available, use the formula $2^n - 2$, with n representing the number of host (0) bits in the subnet mask.

Here's an example to help you put this formula to work: CNT Books Inc. wants 60 subnets for its Class B address: 172.20.0.0/16. The nearest power of 2 to 60 is 64, which equals 2^6 . This means you must reallocate 6 bits from the host portion of the original subnet mask (255.255.0.0) and make them subnet bits.

Reallocating 6 bits, starting from the leftmost bit of the third octet, creates a subnet mask with the bit pattern 11111100. The decimal value for this number is $128 + 64 + 32 + 16 + 8 + 4$, or 252. This reallocating of bits changes the subnet mask from 255.255.0.0 to 255.255.252.0.

To calculate the number of host addresses for each subnet, count the number of 0 bits remaining in the subnet mask to determine the number of bits left for the host ID. In this case, that number is 10 (2 in the third octet and 8 in the fourth). Again, the formula for determining the number of host addresses is $2^n - 2$, so you have $2^{10} - 2 = 1022$. 1022 addresses per subnet should be more than enough for most networks.



Now that you have a correct subnet mask, you need to determine what network numbers can be derived from using this subnet mask. To do this, take the reallocated 6 bits, place them in the network number, and cycle the 6 bits through the possible combinations of values they represent. Table 5-5 shows the first 16 subnetwork numbers resulting from the preceding steps, with the third octet written in binary on the left and the resulting subnetwork address written in decimal on the right. The bits shown in bold are the 6 bits used to create the subnets. If you convert the third octet on the left side from binary to decimal, you'll see that it equals the third octet on the right.

Table 5-5 Subnetwork numbers and addresses

Subnetwork number in binary	Subnetwork address
172.20.00000000.0	172.20.0.0
172.20.00000100.0	172.20.4.0
172.20.00001000.0	172.20.8.0
172.20.00001100.0	172.20.12.0
172.20.00010000.0	172.20.16.0
172.20.00010100.0	172.20.20.0
172.20.00011000.0	172.20.24.0
172.20.00011100.0	172.20.28.0
172.20.00100000.0	172.20.32.0
172.20.00100100.0	172.20.36.0
172.20.00101000.0	172.20.40.0
172.20.00101100.0	172.20.44.0
172.20.00110000.0	172.20.48.0
172.20.00110100.0	172.20.52.0
172.20.00111000.0	172.20.56.0
172.20.00111100.0	172.20.60.0
.....
172.20.11111100.0	172.20.252.0

A Pattern Emerges Table 5-5 shows the first 16 of the possible 64 subnets and the last subnet created for network 172.20.0.0. As you can see, there's a pattern to the subnetwork numbers—they go in increments of 4. You can derive this pattern without having to list the subnets, however. Look at the octet where the subnet bits are reallocated, and then look at the rightmost reallocated bit. The subnet increment is determined by the binary place value of this bit (refer to Table 5-4): in this case, the 4s place.

You know when to stop counting subnets when all the subnet bits are binary 1s, as in the last entry in the table. You also know to stop counting when the subnet number equals the value of the changed octet in the subnet mask. In this case, the subnet mask 255.255.0.0

was changed to 255.255.252.0 after the bit reallocation. The 252 in the third octet of the subnet mask is the same value as the last subnet number.

Determining Host Addresses Similarly, the host addresses in each subnet can be determined by cycling through the host bits. Therefore, the subnetwork 172.20.32.0 would have host addresses from 172.20.32.0 through 172.20.35.255. However, you can't use the IP address in which all host bits are 1s because it's the broadcast address for that network, so your actual range is 172.20.32.1 through 172.20.35.254, giving you 1022 host addresses. Table 5-6 shows this for the first five subnets and the last subnet.

Table 5-6 Host addresses per subnet

Subnetwork number	Beginning and ending host addresses in binary	Beginning and ending host addresses in decimal
172.20.0.0	172.20.00000000.00000001–172.20.00000011.11111110	172.20.0.1–172.20.3.254
172.20.4.0	172.20.00000100.00000001–172.20.00000111.11111110	172.20.4.1–172.20.7.254
172.20.8.0	172.20.00001000.00000001–172.20.00001011.11111110	172.20.8.1–172.20.11.254
172.20.12.0	172.20.00001100.00000001–172.20.00010011.11111110	172.20.12.1–172.20.15.254
172.20.16.0	172.20.00010000.00000001–172.20.00010011.11111110	172.20.16.1–172.20.19.254
.....	
172.20.252.0	172.20.11111100.00000001–172.20.11111111.11111110	172.20.252.1–172.20.255.254

Another Subnet Mask Example In Figure 5-13, the network number is 192.168.100.0, which is a Class C network address with a default subnet mask of 255.255.255.0. The following steps show how to calculate a new subnet mask:

1. In this example, you can see that four cable segments are connected to router interfaces. The WAN cable segment between the two routers counts as a single cable segment and, therefore, a single subnet. You have to account for the WAN subnet even if the network has no hosts because the router interfaces require an IP address. As you can see, there are four subnetworks: Subnet A requires 43 IP addresses (40 for the Windows 7 hosts, 2 for the servers, and 1 for the router interface). Subnet B requires 53 IP addresses, subnet C requires 43 IP addresses, and subnet D requires only 2 IP addresses.
2. To accommodate the required number of subnets (4), you need a power of 2 that's equal to or greater than 4. Because $2^2 = 4$, you need to reallocate 2 bits from the host ID to the network ID.
3. Reallocating 2 bits from the leftmost part of the host portion of the original subnet mask (255.255.255.0) gives the last octet of your new subnet mask: the bit pattern 11000000. Converting to decimal and putting the entire subnet mask together yields 255.255.255.192. The 192 in the last octet is derived from adding the 128s place and the 64s place because they're the only 2 bits that are 1 in the third octet.



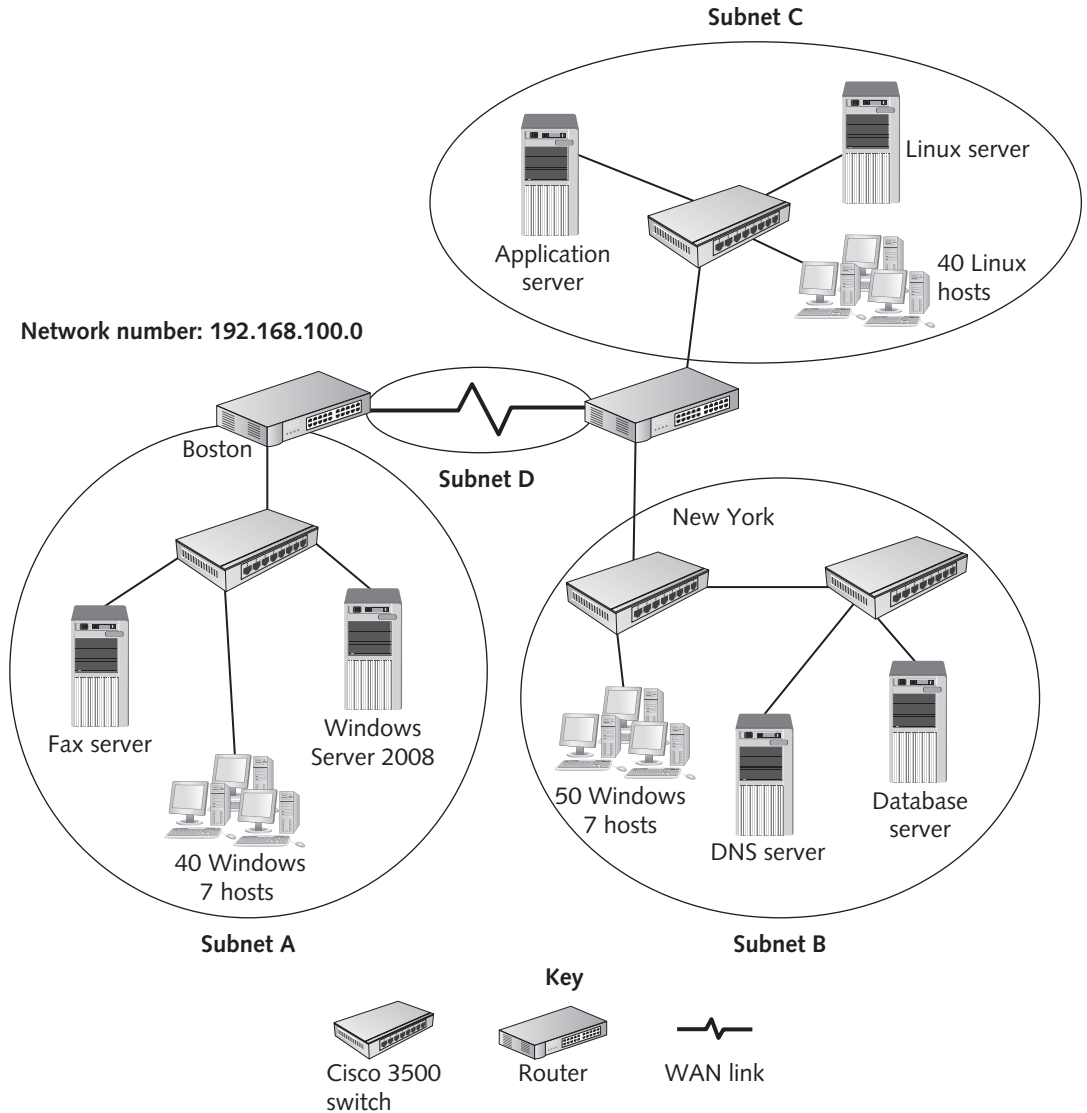


Figure 5-13 A sample network for calculating subnet mask requirements

Courtesy of Course Technology/Cengage Learning

- To be sure you have enough host bits per subnet, use the formula $2^n - 2$, where n is the number of 0 bits in the new subnet mask. The result is $2^6 - 2 = 62$. This number of host addresses satisfies your requirement of a maximum of 53 hosts per subnet.



TIP

To learn more about this topic and try an outstanding tutorial to help you calculate subnets (and supernets, for that matter), go to www.learntosubnet.com.

Supernetting

Although not as commonly practiced as subnetting, **supernetting** is sometimes necessary to solve certain network configuration problems and to make routing tables more streamlined. With routing tables, supernetting is usually referred to as “route aggregation” or “route summarization.”

Supernetting reallocates bits from the network portion of an IP address to the host portion, effectively making two or more smaller subnets a larger supernet. Supernets allow combining two or more consecutive IP network addresses and make them function as a single logical network. Here’s how it works:

1. Say you have four Class C network addresses—192.168.0.0, 192.168.1.0, 192.168.2.0, and 192.168.3.0—available for your network design. You have a total of 900 hosts on your proposed network. You don’t have four router interfaces that can use the four different network numbers, however. You can combine the four networks into one by reallocating 2 bits ($2^2 = 4$) from the network portion of the address and adding them to the host portion. You then have a network address of 192.168.0.0 with the subnet mask 255.255.252.0. The 252 in the third octet is derived from setting the last 2 bits of the original Class C subnet mask (255.255.255.0) to 0, thereby making them part of the host portion.
2. Instead of supporting only 8 bits for the host address portion, the supernet now supports 10 bits ($8 + 2$) for host addresses. This number of bits provides $2^{10} - 2$ host addresses on this supernet, or 1022, which satisfies your requirement for 900 hosts and allows you to assign all host addresses in a single network.

As mentioned, combining two or more small networks into one larger network is only one reason to supernet. Routers on the Internet can have enormous routing tables. The larger the routing table, the more work the router must do to determine where to send a packet. Route aggregation or summarization can combine multiple routing table entries into a single entry, which can drastically decrease the table’s size on Internet routers. This reduction in routing table size increases routers’ speed and efficiency. The procedure is similar to supernetting, except you configure routers.

Routing tables grow partly because routers communicate with one another by sending their routing tables to one another. If several networks can be represented by a single routing table entry, the routing tables are more efficient. Taking the previous example, suppose RouterA in a corporate network has the network addresses 192.168.0.0, 192.168.1.0, 192.168.2.0, and 192.168.3.0 in its routing table, and it communicates with RouterB (see Figure 5-14). Without supernetting/route summarization, RouterA sends all four network addresses to RouterB, each with its 255.255.255.0 subnet mask. Consequently, RouterB’s routing table expands with these four additional routes. However, because all four routes lead to the same place (RouterA), these routes can be represented by a single entry. RouterA can summarize these routes by simply sending RouterB the address 192.168.0.0 with subnet mask 255.255.252.0, which tells RouterB that the addresses 192.168.0.1 through 192.168.3.254 can be reached through RouterA.



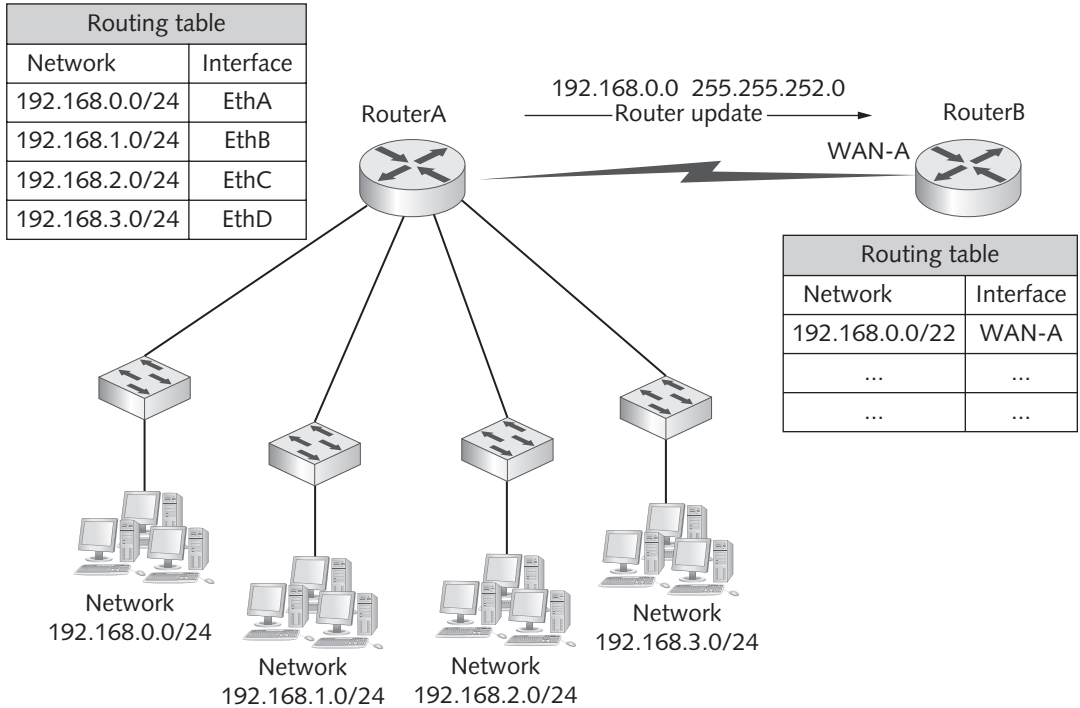


Figure 5-14 RouterA sends a summary of its routing table to RouterB

Courtesy of Course Technology/Cengage Learning

Introduction to Internet Protocol Version 6

IPv6 is the network community’s answer to resolving some problems in IPv4, including limits of the 32-bit address space, a lack of built-in security, a sometimes complicated setup, and a lack of built-in **quality of service (QoS)**. QoS describes a network’s capability to prioritize data packets based on the type of information they contain (for example, voice, video, or file data) or urgency of the information. QoS headers in IPv6 packets can identify packets that require special or priority handling, making applications such as streaming audio and video much easier to implement.

An IPv6 address is 128 bits rather than the 32 bits in an IPv4 address. This length increases the number of possible addresses from about 4 billion in IPv4 to 3.4×10^{38} addresses (that’s 34 followed by 37 zeros!) in IPv6. Unless IP addresses are assigned to every star in the universe, it’s safe to say enough IPv6 addresses will be available.

The IPSec protocol provides authentication and encryption. In IPv4, it must be added to a network, but it’s already incorporated in IPv6. Authentication ensures that the sender and receiver of data packets are known to each other and have permission to send and receive data. Encryption makes the data in packets unreadable except to the computers involved in the transmission.

IPv6 is autoconfiguring, which means there’s no IP address to assign and no subnet mask to determine. Two types of autoconfiguration are available in IPv6:

- Stateless autoconfiguration is the simplest. When a workstation starts, it listens for information broadcast by a local router and assigns itself an address based on the network configuration broadcast by the router and the station’s MAC address.

- Stateful autoconfiguration relies on a DHCP server, as with IPv4. This configuration type requires setting up and configuring a DHCPv6 server. Of the two autoconfiguration types, stateless autoconfiguration is usually the most common.

IPv6 Addresses Unlike IPv4 addresses, which are specified in dotted decimal notation in 8-bit sections, IPv6 addresses are specified in hexadecimal format in 16-bit sections separated by a colon, as in this example: 2001:1b20:302:442a:110:2fea:ac4:2b.

If one of the 16-bit numbers doesn't require four hexadecimal digits, the leading 0s are omitted. Furthermore, some IPv6 addresses contain consecutive 0s in two or more 16-bit sections, so a shorthand notation is used to eliminate consecutive 0 values. Two colons replace two or more consecutive 0 values, as the following example shows:

- Longhand notation: 2001:DB8:0:0:0:2ed3:340:ab
- Shorthand notation: 2001:DB8::2ed3:340:ab

There's actually some order to these seemingly cumbersome IPv6 addresses. In the IPv6 address space, an addressing hierarchy of three parts is used: a public topology, a site topology, and an interface identifier. In short, the first three 16-bit sections (totaling 48 bits) represent the public topology, which could be an Internet backbone service provider, for example. The next 16 bits represent the site topology, such as a business or a local ISP, and the last 64 bits (four 16-bit sections) represent the interface identifier, which is derived from the MAC address on the host's NIC. The interface identifier is the unique host address.

When IPv6 is commonplace, this hierarchical address scheme will make locating Internet resources faster and more efficient and eliminate the need for problem-prone IPv4 processes, such as NAT and complicated subnetting and supernetting.

Hexadecimal Notation Hexadecimal notation, which you have seen in MAC addresses and now IPv6 addresses, is a numbering system like decimal and binary. Hexadecimal, or just hex, is based on powers of 16 and uses 16 symbols to represent all possible numbers. Rather than invent new symbols, however, the numbers 0 to 9 are used for the first 10 symbols and the letters A to F for the remaining 6 symbols, which have the values 10 to 15 in decimal. A hex number, therefore, is expressed by using the symbols 0 to F, and each place value is based on a power of 16. For example, the hex number 4C can be converted to decimal by $4 \times 16^1 + C \times 16^0$. The symbol C represents decimal 12, so 4C in decimal is $64 + 12 = 76$.

Hexadecimal notation is often used to represent numbers in the computer world because it can be converted easily to binary, as it's based on powers of 2. For example, $2^4 = 16$, so every hex digit can be expressed as exactly 4 bits. Converting from hex to binary is just a matter of converting each digit to its 4-bit binary equivalent. For instance, AC4F in hexadecimal is expressed as 1010 1100 0100 1111 in binary.

Networking doesn't require a lot of hex-to-binary conversion until you start working with IPv6 addresses. If you're an aspiring programmer, however, understanding the hex numbering system is sure to be beneficial.

The IPv6 Host ID The host ID of an IPv6 address is typically 64 bits and uses the interface's MAC address to make up the bulk of the address. Because a MAC address is only 48 bits, the other 16 bits come from the value FF-FE inserted after the first 24 bits of the MAC address. In addition, the first two zeros composing most MAC addresses are replaced



with 02. For example, given the MAC address 00-0C-29-7C-F9-C4, the host ID of an IPv6 address is 02-0C-29-FF-FE-7C-F9-C4. This autoconfigured 64-bit host ID is referred to as an Extended Unique Identifier (EUI)-64 interface ID.

By default, Windows Vista and later don't use EUI-64 interface IDs when configuring the link-local address. Instead, they create random interface IDs when autoconfiguring an interface address. However, you can configure Windows to use the EUI-64 interface address with the following command.

```
netsh interface ipv6 set global randomizeidentifiers=disabled
```

You might wonder whether there's an equivalent for the IPv4 loopback address 127.0.0.1. There is, and it's ::1. If you expand this address, it's 0:0:0:0:0:0:0:1.



In contrast to Windows, Linux does use EUI-64 IPv6 host addressing.

Subnetting with IPv6 Although subnetting as done in IPv4 will be a thing of the past, it doesn't mean subnetting won't be used at all in IPv6 networks. Typically, ISPs allocated IPv4 addresses to businesses in groups specified by a network address and an IP prefix. ISPs try to give a business only the number of addresses it requires. However, with IPv6 having such a large address space, most address allocations will have a /48 prefix, even for small home networks. This means the network ID is 48 bits, and the network administrator has 80 bits for assigning subnets and host IDs. Because the host ID is 64 bits, 16 bits are left for creating subnets. This number of bits allows 65,536 subnets, more than enough for all but the largest organizations. Large conglomerates can get multiple /48 prefix addresses or /47 prefix addresses, which provide more than 130,000 subnets. A typical IPv6 address, then, as assigned by an ISP looks like Figure 5-15.

Global routing prefix (48 bits)	Subnet ID (16 bits)	Interface ID (64 bits)
---------------------------------	---------------------	------------------------

Figure 5-15 Structure of a typical IPv6 address

Courtesy of Course Technology/Cengage Learning

With 16 bits available to subnet, there are many strategies you can use. A small network that doesn't have multiple subnets can simply leave the subnet ID as all 0s, for example, and an address in this situation might look like this:

```
2001:DB8:A00:0000:020C:29FF:FE7C:F9C4/64
```

This address begins with 2001:DB8, which is not random. The IPv6 developers realized that people will be writing about how to work with IPv6, so instead of authors using random values for examples, they reserved 2001:DB8 for use in documentation. The A00 in the address is the last 16 bits of the network prefix and was chosen randomly for this example. The 0s following the A00 are the subnet ID, and the last 64 bits are the computer's interface ID. The /64 just indicates that the network portion of the address is the first 64 bits (network prefix plus subnet ID).

A network that does need to subnet could just take the 16 bits for the subnet ID and start counting. For example, a company could make the first three subnets as follows; the bold part of the address is the subnet ID, and the 64-bit interface ID has been omitted.

- 2001:DB8:A00:0000
- 2001:DB8:A00:0001
- 2001:DB8:A00:0002

Large organizations with multiple locations could take a more structured approach and assign each location a bank of subnets, as in the following example:

- 2001:DB8:A00:0000—Assigned to New York location
- 2001:DB8:A00:4000—Assigned to London location
- 2001:DB8:A00:8000—Assigned to Shanghai location

With this strategy, each location has 4000 hexadecimal subnet IDs to work with. For example, New York can make subnets 2001:DB8:A00:0000, 2001:DB8:A00:0001, 2001:DB8:A00:0002, and so forth, up to 2001:DB8:A00:3FFF. Put another way, each location can configure up to 16,384 subnets. As you can see, subnetting does still exist in IPv6, but it's a more straightforward process than in IPv4.



TIP

For more on IPv6, see <http://technet.microsoft.com/en-us/library/bb878121.aspx>.

This chapter has focused on network protocols, specifically those in the TCP/IP suite. Other protocol suites exist, such as IPX/SPX, NetBEUI, and AppleTalk, but they're considered legacy protocols that are no longer used in new network installations. Previous editions of this book covered these protocols, but for this edition, discussion of them has been moved to Appendix B to allow room to focus more on TCP/IP.

Chapter Summary

- TCP/IP is the main protocol suite used in networks. Like most facets of networking, TCP/IP takes a layered approach and is organized in these four layers: Application, Transport, Internetwork, and Network access.
- The Network access layer is composed of network technologies such as Ethernet and token ring. The Internetwork layer is where most network configuration occurs and is composed of IP, ICMP, and ARP. The Transport layer provides reliability and works with segments. The Application layer consists of protocols such as HTTP and DNS and provides an interface for applications to access network services.
- IP addresses are divided into a network ID and host ID. There are three primary address classes: A, B, and C. Address classes determine the default network ID and host ID portions of an IP address. Each class has a range of private IP addresses that can't be used on the Internet; they're used to address private networks. NAT must be used to access the Internet when a host is assigned a private address.
- CIDR largely replaces the IP address class system addressing; it uses a prefix number or subnet mask to determine the network and host IDs of an IP address.

- Subnetting enables an administrator to break a large network into two or more smaller networks that require a router for communication. It also allows an ISP to allocate only the number of public IP addresses a company requires instead of assigning an entire address class.
- IPv6 will eventually replace IPv4 because of its larger 128-bit address space and built-in security and QoS features. IPv6 addresses are expressed as eight four-digit hexadecimal values.

Key Terms

Address Resolution Protocol (ARP) An Internetwork-layer protocol used to resolve a host's IP address to its MAC address. ARP uses a broadcast frame containing the target host's IP address, and the host that's assigned the address responds with its MAC address.

address space The number of addresses available in an IP network number that can be assigned to hosts.

ARP cache A temporary storage location in an IP host's RAM that keeps recently learned IP address/MAC address pairs so that the ARP protocol isn't necessary for each packet sent to a host.

Automatic Private IP Addressing (APIPA) A private range of IP addresses assigned to an APIPA-enabled computer automatically when an IP address is requested via DHCP but no DHCP server responds to the request. *See also* Dynamic Host Configuration Protocol (DHCP).

Classless Interdomain Routing (CIDR) A method of IP addressing in which the network and host IDs are determined by a prefix number that specifies how many bits of the IP address are network bits; the remaining bits are host bits.

connectionless protocol A type of network communication in which data is transferred without making a connection between communicating devices first, and the receiving station gives no acknowledgement that the data was received.

Domain Name System (DNS) An Application-layer protocol that resolves computer and domain names to their IP addresses. DNS uses the UDP Transport-layer protocol. *See also* User Datagram Protocol (UDP).

dotted decimal notation The format used to express an IPv4 address; four decimal numbers separated by periods.

Dynamic Host Configuration Protocol (DHCP) An Application-layer protocol used to configure a host's IP address settings dynamically. It uses the UDP Transport-layer protocol because DHCP messages consist of a single packet and are used on the local LAN.

flow control A mechanism network protocols use to prevent a destination device from becoming overwhelmed by data from a transmitting computer, resulting in dropped packets.

fully qualified domain name (FQDN) A name that includes the hostname, subdomain names (if applicable), second-level domain name, and top-level domain name, separated by periods.

Internet Control Message Protocol (ICMP) An Internetwork-layer protocol used to send error and control messages between systems or devices. It's an encapsulated IP protocol, meaning it's wrapped in an IP header.

Internet Message Access Protocol (IMAP) An Application-layer protocol used by an e-mail client to download messages from an e-mail server; operates on TCP port 143. IMAP also provides fault-tolerance features. It downloads only message headers from the server initially, and then downloads the message body and attachments after the message is selected.

Internet Protocol Security (IPSec) An extension to IP that provides security by using authentication and encryption. It authenticates the identity of computers transmitting data with a password or some other form of credentials, and it encrypts data so that if packets are captured, the data will be unintelligible.

Internet Protocol Version 4 (IPv4) A connectionless Internetwork-layer protocol that provides source and destination addressing and routing for the TCP/IP protocol suite. Uses 32-bit dotted decimal addresses.

Internet Protocol version 6 (IPv6) A connectionless Internetwork-layer protocol that provides source and destination addressing and routing for the TCP/IP protocol suite. Uses 128-bit hexadecimal addresses and has built-in security and QoS features.

IP address A 32-bit dotted-decimal address used by IP to determine the network a host resides on and to identify hosts on the network at the Internetwork layer.

IP prefix A value used to express how many bits of an IP address are network ID bits. Usually expressed as *IPrefixNumber*—for example, 192.168.1.24/27, with 27 as the IP prefix.

localhost The name used to refer to the loopback address in an IP network. *See also* loopback address.

loopback address An address that always refers to the local computer; in IPv4, 127.0.0.1 is the loopback address.

Network Address Translation (NAT) A service that translates a private IP address to a public IP address in packets destined for the Internet, and then translates the public IP address in the reply to the private address.

octet A grouping of 8 bits, often used to identify the four 8-bit decimal numbers that compose an IP address (as in “first octet,” “second octet,” and so forth).

Port Address Translation (PAT) An extension of NAT, a service that allows several hundred workstations to access the Internet with a single public Internet address by using Transport-layer port numbers to differentiate each host conversation. *See also* Network Address Translation (NAT).

port number A field in the Transport-layer protocol header that specifies the source and destination Application-layer protocols that are used to request data and are the target of the request, respectively.

Post Office Protocol version 3 (POP3) An Application-layer protocol used by a client e-mail application to download messages from an e-mail server; operates on TCP port 110.

protocol Rules and procedures for communication and behavior. Computers must use a common protocol and agree on the rules of communication.

protocol stack *See* protocol suite.

protocol suite A set of protocols working cooperatively to provide network communication. Protocols are “stacked” in layers in which each layer performs a unique function required for successful communication. Also called a protocol stack.

quality of service (QoS) A term that describes a network’s capability to prioritize data packets based on the type of information they contain (for example, voice, video, or file data) or urgency of the information.



segment The unit of information used by the Transport layer. A segment is passed up to the Application layer as data and passed down to the Internetwork layer, where it becomes a packet.

Simple Mail Transfer Protocol (SMTP) The standard protocol for sending e-mail over the Internet.

subnet A subdivision of an IP network address space.

subnet mask A 32-bit number in dotted decimal format, consisting of a string of eight or more binary 1s followed by a string of 0s, that determines which part of an IP address is the network ID and which part is the host ID. A binary 1 in the subnet mask signifies that the corresponding bit in the IP address belongs to the network ID, and a binary 0 signifies that the corresponding bit in the IP address belongs to the host ID.

subnetting The process of dividing an IP network address into two or more subnetwork addresses. *See also* subnet.

supernetting Reallocating bits from the network portion of an IP address to the host portion, effectively combining two or more smaller subnets into a larger supernet.

three-way handshake A series of three packets used between a client and server to create a TCP connection. After the three-way handshake has been completed successfully, a connection is established between client and server applications, and data can be transferred.

Transmission Control Protocol (TCP) A connection-oriented Transport-layer protocol designed for reliable transfer of information in complex internetworks.

Transmission Control Protocol/Internet Protocol (TCP/IP) The most common protocol suite in use, TCP/IP is the default protocol in contemporary OSs and the protocol of the Internet.

User Datagram Protocol (UDP) A connectionless Transport-layer protocol designed for efficient communication of generally small amounts of data.

Review Questions

1. An IPv6 address is made up of how many bits?
 - a. 32
 - b. 48
 - c. 64
 - d. 128
 - e. 256
2. The subnet mask of an IP address does which of the following?
 - a. Provides encryption in a TCP/IP network
 - b. Defines network and host portions of an IP address
 - c. Allows automated IP address configuration
 - d. Allows users to use a computer's name rather than its address
3. If a protocol is routable, which TCP/IP layer does it operate at?
 - a. Network access

- b. Internetwork
 - c. Transport
 - d. Application
4. Which of the following is a private IP address and can't be routed across the Internet?
 - a. 192.156.90.100
 - b. 172.19.243.254
 - c. 11.200.99.180
 - d. 221.24.250.207
 - e. 12.12.12.12
 5. Which TCP/IP model layer takes a large chunk of data from the Application layer and breaks it into smaller segments?
 - a. Network access
 - b. Internetwork
 - c. Transport
 - d. Application
 6. Which of the following protocols resolves logical addresses to physical addresses?
 - a. DHCP
 - b. TCP
 - c. IP
 - d. DNS
 - e. ARP
 7. Which of the following protocols provides connectionless service? (Choose all that apply.)
 - a. IP
 - b. UDP
 - c. TCP
 - d. SMTP
 8. If you want to design an Application-layer protocol that provides fast, efficient communication and doesn't work with large amounts of data, what Transport-layer protocol would you design it to use?
 9. Which of the following is the term for identifying packets used by TCP to establish a connection?
 - a. Port number indicators
 - b. Multiwindow agreement
 - c. Three-way handshake
 - d. Sequencing establishment
 10. What is the term for each grouping of 8 bits in an IP address?



11. When using TCP/IP, which of the following must computers on the same logical network have in common? (Choose all that apply.)
 - a. Network ID
 - b. Host ID
 - c. Subnet mask
 - d. Computer name
12. Which of the following IPv6 features is an enhancement to IPv4? (Choose all that apply.)
 - a. Larger address space
 - b. Works at the Internetwork and Transport layers
 - c. Built-in security
 - d. Connectionless communication
13. Which protocol can configure a computer's IP address and subnet mask automatically?
 - a. TCP
 - b. IP
 - c. ARP
 - d. DNS
 - e. DHCP
14. How many bits must be reallocated from host ID to network ID to create 16 subnets?
15. For the Class C network address 192.168.10.0, which of the following subnet masks provides 32 subnets?
 - a. 255.255.255.252
 - b. 255.255.255.248
 - c. 255.255.255.240
 - d. 255.255.255.224
16. How many host bits are necessary to assign addresses to 62 hosts?
 - a. 6
 - b. 5
 - c. 4
 - d. 3
17. Which IP addressing process enables workstations to use private IP addresses to access the Internet?
 - a. Supernetting
 - b. NAT
 - c. DHCP
 - d. Subnetting

18. When a Windows computer is configured to use DHCP but no DHCP server is available, what type of address is configured automatically for it?
 - a. PAT
 - b. APIPA
 - c. NAT
 - d. Static
19. Which of the following represents a valid IPv6 address?
 - a. 2001:345:abcd:0:230:44
 - b. 2001:345:abcd::BEEF:44
 - c. 2001:345::abcd:0:79f::230:44
 - d. 2001:345:abcd:0:FEED:230:44
20. Which of the following is a reason to subnet? (Choose all that apply.)
 - a. Networks can be divided into logical groups.
 - b. Subnetting eliminates the need for routers.
 - c. Subnetting can decrease the size of broadcast domains.
 - d. There's no need to assign static IP addresses to each computer.
21. Which of the following Application-layer protocols typically uses the UDP Transport-layer protocol? (Choose all that apply.)
 - a. HTTP
 - b. DNS
 - c. DHCP
 - d. FTP
22. Which is the correct order of headers, from left to right, in a completed frame?
 - a. Frame, TCP, IP
 - b. UDP, frame, IP
 - c. TCP, IP, frame
 - d. Frame, IP, UDP
23. Which of the following is a task performed by the Network access layer? (Choose all that apply.)
 - a. Verifies that incoming frames have the correct destination MAC address
 - b. Defines and verifies IP addresses
 - c. Transmits and receives bit signals
 - d. Resolves MAC addresses by using IP addresses
 - e. Delivers packets efficiently
24. What field of the IP header does the Tracert program use to get the IP address of routers in the path?



- a. Version
 - b. TTL
 - c. Checksum
 - d. Protocol
25. Which of the following is *not* found in a connectionless Transport-layer protocol? (Choose all that apply.)
- a. Three-way handshake
 - b. Port numbers
 - c. Checksum
 - d. Acknowledgements

Challenge Labs



Challenge Lab 5-1: Capturing and Identifying the Three-way Handshake

Time Required: 30 minutes

Objective: Determine which packets create the three-way handshake used in establishing a communication session.

Required Tools/Equipment: Your classroom computer with Wireshark installed and access to the Internet

Description: Using Wireshark and a suitable capture filter, capture the packets involved in an HTTP session that you start by opening a Web page. Find the three packets that constitute the three-way handshake. Perform the following tasks:

- What capture filter did you use to limit Wireshark to capturing only packets related to HTTP?

- Find the three-way handshake that immediately precedes the first HTTP packet. Which Transport-layer protocol was used to create the connection?

- Find the following fields in the Transport-layer header of the first packet in the three-way handshake and write down their value:
 - Source port: _____
 - Destination port: _____
 - Sequence number: _____
 - Flags: Syn: _____
 - Window size: _____
 - Maximum segment size: _____

- Find the following fields in the Transport-layer header of the third packet in the three-way handshake. Research their meanings, and then write down their values along with brief descriptions of them:
 - Window size: _____
 - Maximum segment size: _____
- How are the sequence number and acknowledgement used to make this protocol reliable?



Challenge Lab 5-2: Creating a Subnet Mask

Time Required: 30 minutes

Objective: Create a suitable subnet mask and list the resulting networks and host address ranges.

Required Tools/Equipment: Pen and paper or a word-processing document

Description: Review the network diagram in Figure 5-16. Given this information, devise a subnet mask that works for this network. Assume the original subnet mask is the default for the network class. Write the subnet mask and number of subnets, and fill in the chart with the network numbers and host address ranges. (*Hint:* The chart has more rows than you need.)

- Number of subnets needed: _____
- Subnet mask: _____
- Number of subnets created: _____

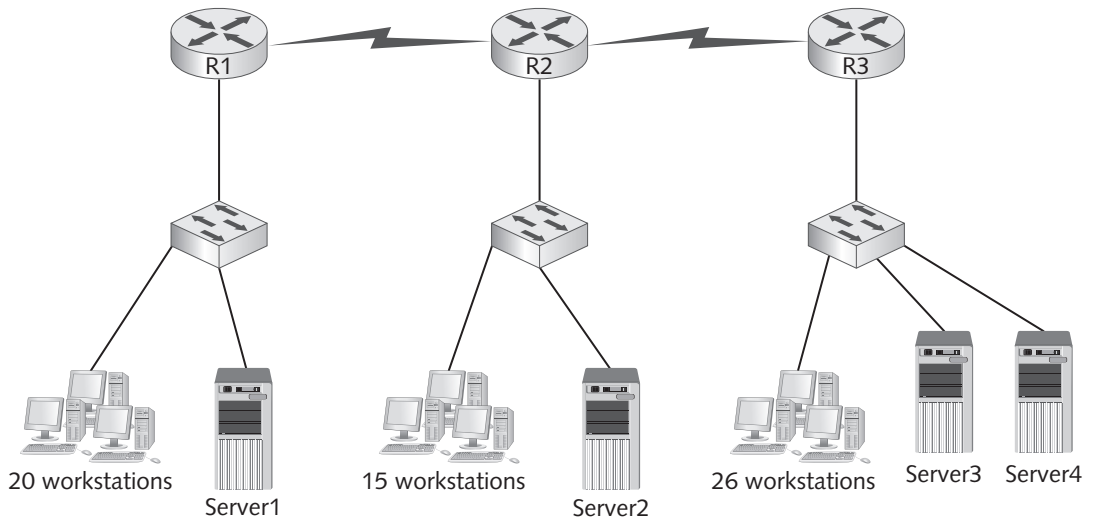


Figure 5-16 Network diagram for Challenge Lab 5-2

Courtesy of Course Technology/Cengage Learning



Network number	Host address range

Case Projects



Case Project 5-1

As the network administrator for a growing ISP, you want to make efficient use of your network addresses. One of the network addresses IANA assigned to you is a Class C network of 197.14.88.0. You have decided to use the addresses in this Class C network to satisfy the IP address requirements of 16 corporate customers who need between 10 and 14 addresses each. Without using a subnet calculator, calculate a subnet mask that meets their needs. List the subnet mask and the first four subnetwork addresses the mask will create.

Case Project 5-2

You work at a help desk and have just received a call from an employee who says she can't access network resources. You want the employee to view her IP address configuration. Write an e-mail to the employee, explaining what command-line program to use and how she can use it to find the information you need. After following your instructions, the employee tells you that her IP address is 169.254.14.11 with the subnet mask 255.255.0.0. What conclusion can you make from this information?

Case Project 5-3

You must install 125 computers for a new business that wants to run TCP/IP and have access to the Internet. The ISP in town will assign you only four public IP addresses, so you decide to assign the computers addresses in the range 172.16.1.1/16 through 172.16.1.125/16. What else must you do to allow these computers to access the Internet?

Network Reference Models and Standards

After reading this chapter and completing the exercises, you will be able to:

- Explain the OSI reference model layers and their relationship to hardware and software
- Explain the IEEE 802 networking model and related standards

The Open Systems Interconnection (OSI) reference model for networking explains how networks behave within an orderly, seven-layered model for networked communication. The OSI model isn't specific to a protocol suite and can be applied to most networking protocols past and present. Many of the networking hardware and software components discussed in this book can be identified as working within one or more of the OSI model layers.

Although the OSI model isn't specific to any one protocol suite, it's the standard model for discussing, teaching, and learning the field of computer networking. It's unlikely you'll have a course in networking that doesn't at least mention the OSI model, and some courses you take might cover it in more detail than in this chapter. Because you have already learned about the layered architecture of the TCP/IP model, some of the information in this chapter has already been introduced. However, the concept of a layered architecture in network communication is so vital to your understanding of how to configure and troubleshoot networks, the repetition should be worthwhile. In addition, descriptions of network devices refer to the OSI model rather than the TCP/IP model; for example, a switch might be called a "Layer 2 switch."

The OSI model is a general framework for how networking systems should operate, and the IEEE 802 networking standards are formal specifications for how to implement particular networking technologies. The IEEE standards are most important if you're designing network hardware or writing network drivers or protocols, as these standards define how vendors of networking products should implement particular technologies.

Introducing the OSI and IEEE 802 Networking Models

Several networking models sought to create an intellectual framework to clarify network concepts and activities, but none has been as successful as the **Open Systems Interconnection (OSI) reference model** proposed by the **International Organization for Standardization (ISO)**. This model is sometimes referred to as the ISO/OSI reference model.



ISO isn't an acronym; it comes from the Greek prefix *iso*, which means "equal" or "the same." The ISO, based in Geneva, Switzerland, is a network of national standards institutes from 140 countries. The expanded name differs from language to language. For example, in France the organization is the Organisation Internationale de Normalisation. The term ISO gives the network of institutes a common name.

The OSI reference model has become a key part of networking, in large part because it's a common framework for developers and students of networking to work with and learn from. The attempt to develop a working set of protocols and technologies based on the OSI model and put these efforts into common use never materialized, partly because existing protocols, such as TCP/IP, were already entrenched in the marketplace. However, the OSI reference model has a prominent place in networking as a model and teaching tool. This chapter covers the model's seven-layer organization, the function of each layer, and the networking devices and components operating at each layer.



The set of protocols developed to conform to the OSI model is called ISO. You can view the fruits of these labors at www.protocols.com/pbook/iso.htm.

This IEEE 802 networking model provides detailed implementation specifications for a number of networking technologies. As you learned in Chapter 3, the IEEE standards define the various Ethernet standards from 10 Mbps up to 100 Gbps. In fact, the 802 specification encompasses most types of networking and allows adding new types of networks (such as the newest 40 Gigabit and 100 Gigabit standards) as necessary. This chapter briefly discusses how the IEEE 802 standards relate to the OSI model.

Role of a Reference Model

You might wonder why a reference model for networking is needed and why the layer concept in particular is so valuable. To see the value of a layered model outside the field of networking, take a look at the complete process of a letter being created, sent, and delivered via the U.S. postal service:

1. Tom, who lives in New York, writes a letter to Cindy, who lives in San Francisco. When the letter is finished, it's ready for Cindy to read, but Tom needs to get the letter to Cindy, so he decides to use the U.S. mail.
2. Tom folds the letter and places it in an envelope, which is the container required by the U.S. mail letter-sending protocols. Tom can't send the letter yet, however; first he must address the envelope.
3. Tom addresses the envelope by putting Cindy's name and address in the middle of the front of the envelope, which is where the post office expects to find the destination address. Tom also puts his return address on the envelope's upper-left corner.
4. Before Tom can send the envelope, per post office protocol, he must place a stamp on the envelope's upper-right corner.
5. Tom then walks to the post office and drops the letter in the mailbox. At this point, Tom's job is done; it's up to the post office (the next layer) to take care of getting the letter to its destination.
6. The mail carrier picks up the mail in the mailbox at the prescribed time and brings it to the central office for sorting. The mail carrier's job is done, and now it's up to other post office workers (the next layer) to get the letter to its destination.
7. The mail is sorted according to zip code, which identifies the part of the country the mail is destined for. After sorting, the letter goes into the pile headed for the West Coast of the United States. The mail is put on a plane, and the job of the post office worker in New York is completed.
8. After the mail arrives in San Francisco, it's sorted by zip code to determine which area of San Francisco to deliver it to. After the letter has been sorted, a mail carrier takes it on his or her route.
9. The mail carrier uses the street address to determine which house to deliver the letter to, and he or she leaves the letter in Cindy's mailbox. At this point, the mail carrier's job is done.



10. Cindy receives the letter, opens the envelope, and now has exactly what Tom had in his hand before he placed the letter in the envelope. Mission accomplished.

As you can see, a number of tasks have to be completed to deliver this message. Each task is separate, but for one task to be completed, the previous task (or layer) must be completed correctly:

- The letter has to be written.
- The letter has to be placed in an envelope and addressed in the correct format.
- The local post office in New York has to sort the letter correctly and get it on the right plane to San Francisco.
- The post office in San Francisco has to sort the letter correctly for the right part of town.
- The local carrier has to deliver the letter to the correct house.
- The recipient has to receive the letter, open it, and read it.

A layered approach to what might otherwise be a daunting process reduces its complexity and turns it into a series of interconnected and well-defined tasks. Each task or activity can be handled separately and its issues solved independently, often without affecting the procedures of other tasks. This approach creates a method for solving big problems by reducing them to a series of smaller problems with separate solutions.

To further exemplify the value of layers in this analogy, consider the effect of having the mail carrier switch from walking the mail route to driving a delivery truck. In fact, the only step that's affected is the mail carrier's job—his or her job gets done faster. Addressing the envelope is still done in the same way, and post office workers still follow the same procedure to sort the mail. In short, people involved in these steps don't even have to know that the mail carrier is using a truck to get from house to house. As you can see, with a layered approach, one part of the process can change, sometimes drastically, but the rest of the process remains unchanged. Now think about what's necessary to upgrade from 100 Mbps Ethernet to 1000 Mbps Ethernet: Change the NICs and/or the switch and you're done. There's no need to change the protocols or applications. By the same token, IPv6 can replace IPv4 without having to change the Transport layer or Network access layer in the TCP/IP model.

Structure of the OSI Model

The OSI model divides network communication into the seven layers shown in Figure 6-1.



Here are two mnemonics to remember the seven layers of the OSI reference model. From the bottom up, starting with the Physical layer, the mnemonic is "People Do Not Throw Sausage Pizza Away." From the top down, starting with the Application layer, try "All People Studying This Need Drastic Psychotherapy" or "All People Seem To Need Data Processing."

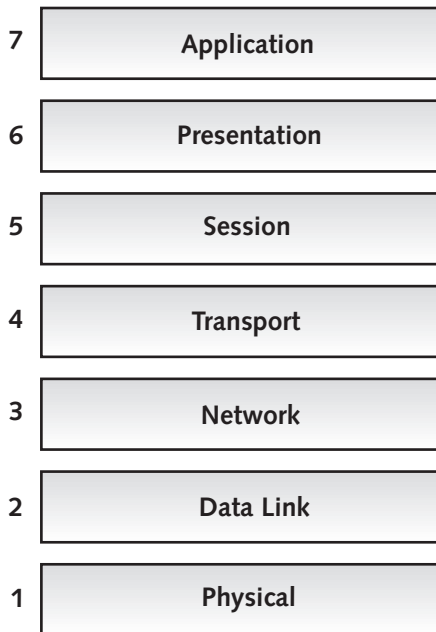


Figure 6-1 The seven layers of the OSI reference model

Courtesy of Course Technology/Cengage Learning

At the top, the Application layer provides a set of interfaces that enable user applications—such as Windows Explorer or Microsoft Word—to access network services. These user applications aren't part of the OSI model but communicate with its top layer. On the other hand, some user applications, such as Web browsers and e-mail programs, are integrated with functions of the Application layer (as well as the Presentation and Session layers).

At the bottom of the OSI model, the Physical layer is where the network medium and the signals traversing it reside. All the activities needed to handle network communication occur between the top and bottom layers. To comprehend how a network works as a whole, you simply need to understand how each layer functions, what networking components and devices operate at each layer, and how the layers interact with one another.

Each layer in the OSI model has its own set of well-defined functions, and the functions of each layer communicate and interact with the layers immediately above and below it. For example, the Transport layer works with the Network layer below it and the Session layer above it. The Physical layer doesn't have a layer below it, and the Application layer, although not having a layer above it, interacts with user applications and network services.

Because you're already familiar with the TCP/IP model, now is a good time to compare the two models. Your understanding of how the TCP/IP layers work gives you a context for the OSI model's somewhat more detailed layers. Figure 6-2 shows this relationship. Notice that both models contain an Application layer, but the TCP/IP model combines the functions of the OSI model's Application, Presentation, and Session layers. It's not that TCP/IP doesn't perform the function of these layers; it's just that a single TCP/IP Application-layer protocol performs all three functions. The Transport layer in both models is equivalent in name and function. The OSI model's Network layer is equivalent to the TCP/IP Internetwork layer.



OSI model	TCP/IP			
Application	Telnet	FTP	DHCP	TFTP
Presentation	HTTP	SMTP	DNS	SNMP
Session	Application layer			
Transport	TCP			UDP
	Transport layer			
Network	ICMP			ARP
	IP			
	Internetwork layer			
Data Link	Network access layer			
Physical				

Figure 6-2 Comparing the OSI model and the TCP/IP model

Courtesy of Course Technology/Cengage Learning

The OSI model divides the function of the Network access layer into two separate layers—Data Link and Physical—that have distinct jobs.

The familiar network connection properties in a Windows OS are used again to show the layers in an OS context (see Figure 6-3). The Ethernet adapter shown in the Connect using text box represents the model’s two bottom layers: Physical and Data Link. Internet Protocol (TCP/IP) represents the next two layers: Network and Transport. Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks represent the top three layers: Session, Presentation, and Application. All these components (layers) are required for Windows network communication to work, but any component can be replaced with a suitable substitute (for example, replacing the NIC and its driver with a different NIC and driver) without affecting the other components.

Each layer in the model provides services to the next higher layer until you get to the Application layer, which has the job of providing services to user applications. In the layered approach, each layer on one computer behaves as though it were communicating with its counterpart on the other computer. This means each layer on the receiving computer sees network data in the same format its counterpart on the sending computer did. This behavior is called **peer communication** between layers, as shown in Figure 6-4. Referring back to the

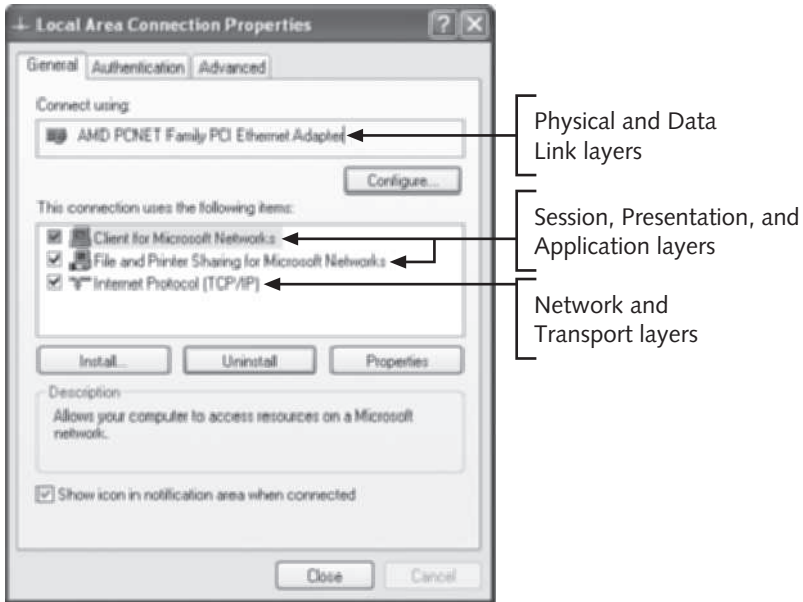


Figure 6-3 Layers of the OSI model in the Local Area Connection Properties dialog box

Courtesy of Course Technology/Cengage Learning

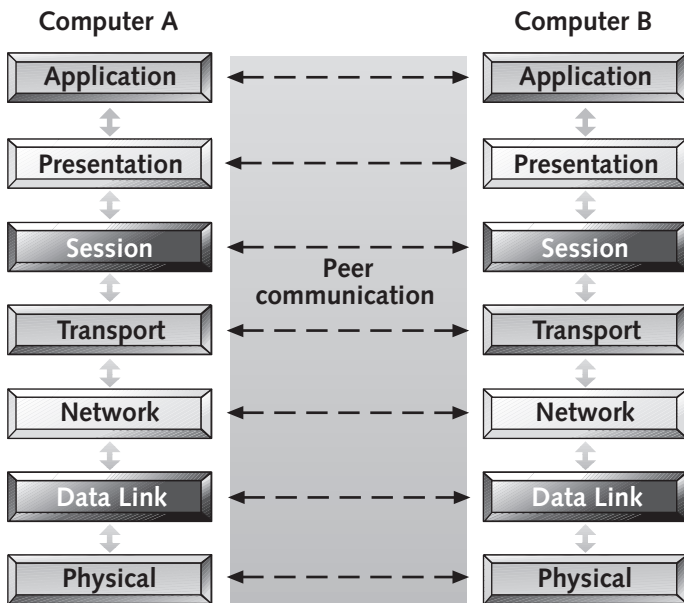


Figure 6-4 Peer communication between OSI layers

Courtesy of Course Technology/Cengage Learning



U.S. mail analogy, the receiver of the letter, Cindy, sees the letter in the same format Tom saw it before it was placed in an envelope. Likewise, the mail carrier in San Francisco saw the envelope (with the letter in it) in the same format the letter carrier in New York did.

In network communication, data passes up and down the protocol stacks on both machines. Operations occurring on the way down the stack on the transmitting machine are largely reversed on the way up the stack on the receiving machine, so data on one layer of the sender is nearly identical to data arriving on that layer for the receiver.

On data's way down the stack, it's divided into data units suitable for each layer. Each unit, called a **protocol data unit (PDU)**, is passed from one layer to another on its way up or down the protocol stack. At some layers, the software adds its own formatting or addressing to the PDU, which, as you know, is called a header. As you learned in Chapter 1, the process of adding this header is called "encapsulation." In the mail delivery analogy, the sender must put the letter (data) into an envelope (encapsulation) and address the envelope.

When data arrives at the receiving end, the packet travels up the stack from the Physical layer through the Application layer. At each layer, the software reads its PDU data and performs any additional processing that's required. It then strips its header information from the PDU (a process called **deencapsulation**) and passes the PDU to the next higher layer. When the packet leaves the Application layer, data is in a form that the receiving application can read and has been stripped of all the network addressing and packaging instructions needed to move the data from sender to receiver. Again, using the mail delivery analogy, the deencapsulation process is analogous to the letter recipient reading the envelope and verifying the address before opening and discarding the envelope to finally read the letter.

Simulation 11 shows how data generated from an application travels down through the OSI model layers, with encapsulations added where necessary, and how the process is reversed on the receiving computer.



Simulation 11: Peer communication with the OSI model

The following sections describe the layers of the OSI model and the services each one provides. After reading this material, you should have a good idea of each layer's functions, how each layer interacts with adjacent layers, and some problems that can occur at each layer. Protocols and software components are listed for each layer. When applicable, devices that operate at a layer are listed. When a device is said to operate at a layer, it means the highest layer in which the device operates. For example, a PC operates at Layer 7, so it's considered a Layer 7 device, but clearly, a PC also operates at Layers 6 through 1.

Both the name and number of the layer are listed. When discussing devices or protocols in relation to the OSI model, the OSI layer number rather than its name is often used. For example, you hear terms such as "Layer 2 switch" or "Layer 7 gateway."

Application Layer

The **Application layer** (Layer 7) provides interfaces for applications to access network services, such as file sharing, message handling, and database access. It also handles error

recovery for applications, as needed. The PDU at this layer (and the Presentation and Session layers) is referred to simply as “data.”

Generally, components at the Application layer have both a client component and a server component. An example of this client/server pairing is a Web browser (client component) that accesses a Web server (server component), both of which provide access to the Application-layer protocol HTTP. Other examples are Client for Microsoft Networks, used to access Windows network services (such as File and Printer Sharing), and the UNIX/Linux Network File System (NFS) client, which provides access to shared file resources. Common protocols found at the Application layer include HTTP, FTP, SMB/CIFS, TFTP, and SMTP. Computers with network OSs and some security devices operate at Layer 7 because they work with these Application-layer protocols.

Possible problems at this layer include missing or misconfigured client or server software and incompatible or obsolete commands used to communicate between a client and server. In addition, Application-layer protocols that use a connectionless Transport-layer protocol are more susceptible to network disruptions and must provide their own error recovery or rely on error recovery from the user application.

Presentation Layer

The **Presentation layer** (Layer 6) handles data formatting and translation. For outgoing messages, it converts data into a format specified by the Application layer, if necessary; for incoming messages, it reverses the conversion if required by the receiving application. In short, Layer 6 “presents” data in a suitable format to the Application layer. The Presentation layer handles protocol conversion, data encryption and decryption, data compression and decompression, data representation incompatibilities between OSs, and graphics commands.

An example of functionality at this level is a Web browser displaying graphics files embedded in a Web page. In this situation, the Presentation-layer component informs the Application layer what type of data or graphics format to display. Yet another example involves character conversion. For example, PCs represent the carriage return/line feed combination in text files differently than Linux and UNIX systems do. If no conversion takes place, a text file created on a Linux system looks like one long string of sentences when read by Notepad on a PC. However, if the file is transferred from Linux to a PC with a file transfer program that can convert the codes, the Presentation-layer component of the file transfer program handles the conversion. As another example, a Web browser that connects to a secure Web server with encryption protocols must encrypt data before it’s transferred to the server and decrypt data arriving from the Web server, which is a Presentation-layer function.

A software component known as a “redirector” operates at this layer. It intercepts requests for service from the computer; requests that can’t be handled locally are redirected across the network to a network resource that can handle the request. Software components operating at this layer are usually built into the Application-layer component. These components include FTP clients and servers, HTTP clients and servers, and OS-specific clients and servers, such as Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks.

Possible problems occurring at this layer include incompatible or missing translation software, in which the Presentation layer on one system doesn’t have the necessary decryption, decompression, graphics-processing, or data translation software to interpret received data correctly.



Session Layer

Layer 5, the **Session layer**, permits two computers to hold ongoing communications—called a “session”—across a network, so applications on either end of the session can exchange data for as long as the session lasts. The Session layer handles communication setup ahead of data transfers when necessary and session teardown when the session ends. Some common network functions this layer handles include name lookup and user logon and logoff. Therefore, DNS and other name resolution protocols work in part at this layer, as do the logon/logoff function and some authentication protocols built into most client software, such as FTP, Client for Microsoft Networks, and NFS.

The Session layer also manages the mechanics of ongoing conversations, such as identifying which side can transmit data when and for how long. In addition, a process called “check-pointing” is performed at this layer. Checkpointing is a synchronization process between two related streams of data, such as an audio and a video stream in a Web conferencing application. The Session layer keeps the audio in sync with the video.

Transport Layer

The **Transport layer** (Layer 4) manages data transfer from one application to another across a network. It breaks long data streams into smaller chunks called “segments.” Segmenting the data is important because every network technology has a maximum frame size called the **maximum transmission unit (MTU)**. For Ethernet, the MTU is 1518 bytes, which means segments must be small enough to allow for the Network-layer and Data Link-layer headers and still be no larger than 1518 bytes. If segmenting doesn’t occur, as with UDP, the Network layer must fragment the packets it creates, leading to inefficient and possibly unreliable communication. Figure 6-5 shows a simplified example of what the original data might look like and what each segment might look like after data is broken up into smaller pieces and the header is added.

Data created by the Application, Presentation, and Session layers:

Data data data data data data data data data data
Data data data data data data data data data data
Data data data data data data data data data data

Data is broken into smaller chunks by the Transport layer:

Transport-layer header: Segment 1	Data data data data data data data data data data
-----------------------------------	---

Transport-layer header: Segment 2	Data data data data data data data data data data
-----------------------------------	---

Transport-layer header: Segment 3	Data data data data data data data data data data
-----------------------------------	---

Figure 6-5 The Transport layer breaks data into smaller chunks called segments

Courtesy of Course Technology/Cengage Learning

To ensure reliable delivery, the Transport layer includes flow control and acknowledgements and handles resequencing segments into the original data on receipt. Flow control ensures that the recipient isn’t overwhelmed with more data than it can handle, which could result in dropped packets.

The PDU at this layer is called a segment (see Figure 6-6). The components working at this layer include TCP and UDP from the TCP/IP protocol suite. As you learned in Chapter 5, however, UDP doesn't perform all the functions expected of an OSI model Transport-layer protocol, so it's sometimes called a "pseudo-Transport-layer protocol."

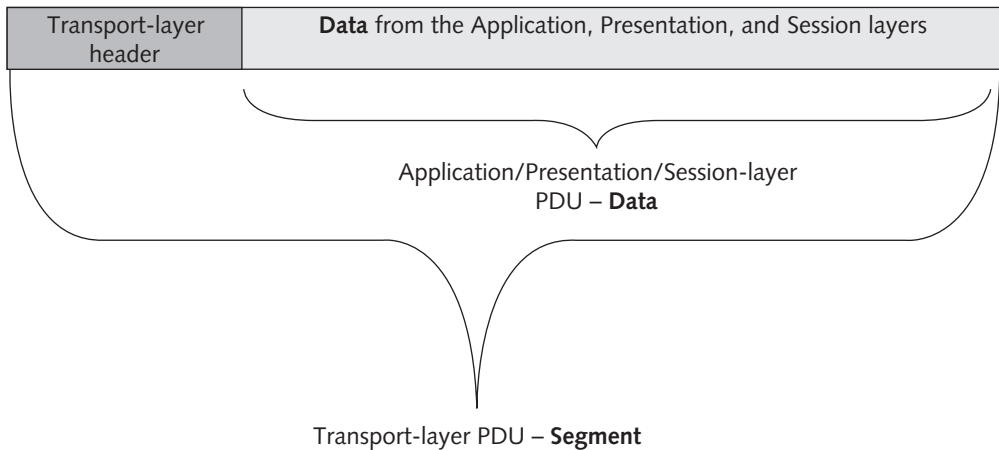


Figure 6-6 The Transport-layer PDU: a segment

Courtesy of Course Technology/Cengage Learning

Some key fields in the Transport-layer header include the following:

- *Source and destination port numbers*—As discussed in Chapter 5, port numbers identify the application or service to which the segment should be delivered. Each application or service is assigned a unique port number so that when data arrives, the Transport-layer protocol (TCP or UDP) knows to which application it should be transferred.
- *Sequence and acknowledgement numbers*—Found in TCP headers but not in UDP headers, sequence and acknowledgement numbers are used to ensure that all data sent was received; if segments arrive out of order, the sequence number is used to reorder them.
- *Window size*—This field specifies the maximum amount of data in bytes that can be transferred before the sender requires an acknowledgement. If an acknowledgement isn't received, the sender sends the data again that was sent since the last acknowledgement was received. The window size, along with acknowledgements, provides flow control because if a computer is overwhelmed with too much data, it can reduce the window size, causing the sending computer to send fewer segments before waiting for an acknowledgement.

Problems that can occur at this layer include segments that are too large for the medium between source and destination networks. This situation forces the Network layer to fragment the segments, which causes performance degradation. In addition, hackers can exploit TCP's handshaking feature with a half-open SYN attack, discussed in Chapter 10.

Network Layer

Layer 3, the **Network layer**, handles logical addressing, translates logical network addresses (IP addresses) into physical addresses (MAC addresses), and performs best path selection and routing in an internetwork. A router performs best path selection when multiple pathways, or routes, are available to reach a destination network; the router attempts to choose the best, or fastest, path.

As you can see, this layer performs the same tasks as TCP/IP's Internetwork layer. It's also the traffic cop for network activity because it provides access control. **Access control** is handled at the Network layer during the routing process; the router consults a list of rules before forwarding an incoming packet to determine whether a packet meeting certain criteria (such as source and destination address) should be permitted to reach the intended destination. This feature of routers is one reason to divide large networks into smaller subnets because by creating several logical groups of computers, you can control which users have access to which resources, using routers as gatekeeper.

The PDU at the Network layer is a packet, as shown in Figure 6-7. The software components working at this layer include IP, ARP, and ICMP from the TCP/IP protocol suite. Routers, of course, work at this layer, as do firewalls and certain remote access devices, such as virtual private network (VPN) servers. A switch with routing capabilities, called a Layer 3 switch, also works at the Network layer. Essentially, any device that mainly works with packets and the source and destination IP addresses in a packet is said to be a Network-layer device or Layer 3 device.

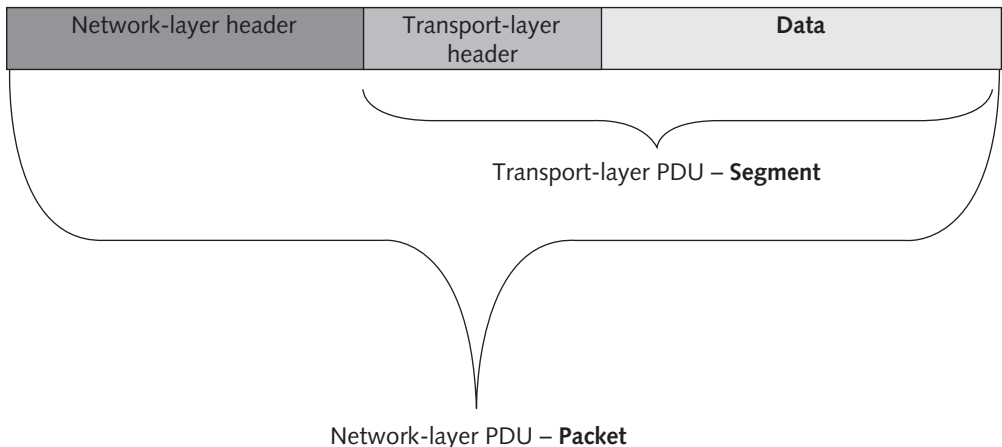


Figure 6-7 The Network-layer PDU: a packet

Courtesy of Course Technology/Cengage Learning



An infrequently used protocol called Reverse Address Resolution Protocol (RARP) also works at the Network layer. It resolves a known MAC address to an IP address.

Many problems can occur at the Network layer and often include incorrect IP addresses or subnet masks, incorrect router configuration, and router operation errors.

Data Link Layer

Layer 2, the **Data Link layer**, works with frames and is the intermediary between the Network layer and Physical layer. It defines how computers access the network medium—also called “media access control,” which is why the MAC address is defined at this layer. As you learned in Chapter 3, media access control methods include CSMA/CD and token passing, among others.

As shown in Figure 6-8, a frame consists of both a header and a trailer component. The trailer component labeled “FCS” (frame check sequence) contains the CRC error-checking code discussed in Chapter 3. The CRC value is recalculated on the receiving end, and if the sent and recalculated values agree, the assumption is that data wasn’t altered during transmission. If they differ, the frame is discarded in most networking technologies. Note that the CRC is recalculated at every intermediary device (usually a router) between the source and destination computer. For example, if the frame is delivered to a router, the router recalculates the CRC and compares it with the original to make sure the frame wasn’t damaged in transport. Next, the router changes the source MAC address to its own MAC address and the destination MAC address to the next device (which might be another router or the final destination device). Therefore, the router must recalculate the CRC and place it in the frame trailer because the frame’s contents were changed. Simulation 9, which you saw in Chapter 5, shows how a frame header changes during its journey through an internetwork.

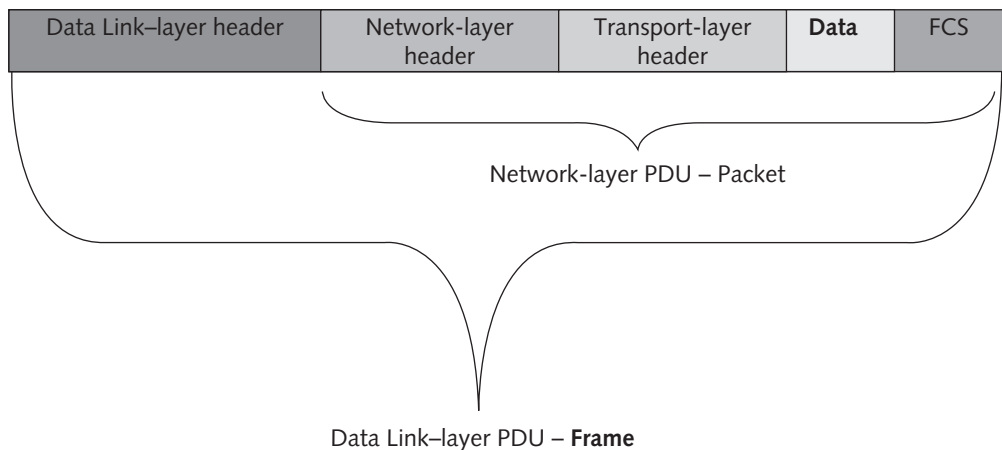


Figure 6-8 The Data Link-layer PDU: a frame

Courtesy of Course Technology/Cengage Learning



Simulation 9: The changing frame header

The Data Link header contains fields for source and destination addresses. The destination address is the hardware (MAC) address of the computer to which the frame should be

delivered or of an intermediate device, such as a router. The source address is the MAC address of the sending computer or intermediary device and tells the recipient where to send a reply.

After receiving a frame from the Physical layer and verifying the destination MAC address and the CRC, the Data Link layer strips its header and trailer information from the frame and sends the resulting packet up to the Network layer for further processing. In most networking technologies, the Data Link layer discards frames containing CRC and other frame errors. However, it's the responsibility of the upper layers (usually the Transport layer) to retransmit data that has been discarded because of errors. TCP, for example, detects missing segments caused by discarded frames because of the sequence numbers it uses to keep track of all segments.

The software component operating at this layer is the NIC driver, and the hardware components include NICs and switches. A NIC operates at this layer because it contains the MAC address and is responsible for media access control. Switches operate at this layer because they do their job by examining MAC addresses in the frame header and using this information to switch packets from incoming ports to outgoing ports. Because media access control is defined in this layer, networking technologies such as Ethernet and token ring operate at this layer. Problems occurring in the Data Link layer include collisions and invalid frames, which can be caused by collisions, poor network design, line noise, or NIC driver problems. Another problem at this layer results from trying to use incompatible network technologies, such as token ring and Ethernet, on the same LAN.

Physical Layer

Last but not least, the job of the **Physical layer** (Layer 1) is to convert bits into signals for outgoing messages and signals into bits for incoming messages. The type of signals generated depend on the medium; for example, wire media, such as twisted-pair cable, use electrical pulses, fiber-optic media use pulses of light, and wireless media use radio waves. At this layer, details for creating a physical network connection are specified, such as the type of connectors used to attach the medium to the NIC.

The Physical layer also specifies how to encode 1s and 0s. **Encoding** is representing 0s and 1s by a physical signal, such as electrical voltage or a light pulse. For example, a 1 bit might be represented on a copper wire by the transition from a 0-volt to 5-volt signal, whereas a 0 bit might be represented by the transition from a 5-volt signal to a 0-volt signal.

The network components working at the Physical layer include all the cables and connectors used on the medium plus repeaters and hubs. Problems occurring here are often related to incorrect media termination, EMI or noise that scrambles the signals, and NICs and hubs that are misconfigured or don't work correctly.

Summary of the OSI Model

The OSI model is a helpful way to categorize and compartmentalize networking activities, and most discussions of protocol suites and networking software use its terminology. Table 6-1 summarizes the actions occurring at each layer. Even though most protocol suites don't adhere strictly to this model (perhaps because so many of them were already implemented in some form before the model's development), they still incorporate its outlook on networking.

Table 6-1 OSI model summary

Layer	PDU	Protocols/software	Devices	Function
7. Application	Data	HTTP, FTP, SMTP, DHCP	Computers	Provides programs with access to network services
6. Presentation	Data	Redirectors	N/A	Handles data representation to application and data conversions, ensures that data can be read by the receiving system, and handles encryption/decryption
5. Session	Data	DNS, authentication protocols	N/A	Establishes, maintains, and coordinates communication between applications
4. Transport	Segment	TCP, UDP	N/A	Ensures reliable delivery of data, breaks data into segments, handles sequencing and acknowledgements, and provides flow control
3. Network	Packet	IP, ICMP, ARP	Routers, firewalls, Layer 3 switches	Handles packet routing, logical addressing, and access control through packet inspection
2. Data Link	Frame	Ethernet, token ring, FDDI, NIC driver	Switches, NICs	Provides physical device addressing, device-to-device delivery of frames, media access control, and MAC addresses
1. Physical	Bits	N/A	Network media, hubs/repeaters, connectors	Manages hardware connections, handles sending and receiving binary signals, and handles encoding of bits



Although not all networking protocols adhere to the OSI model, a network administrator's clear understanding of the functions at each layer is essential in troubleshooting networks and network equipment and in understanding how network devices operate.

The OSI model helps explain how data is formatted and moves up and down the protocol stack and from computer to computer. Although TCP/IP was developed long before the OSI model, you can see how the TCP/IP model adheres to the OSI concept of a layered architecture. Understanding both models and their relationship to each other is important, but when discussing devices and software that work at particular layers, the OSI model's layer names and numbers are the most pertinent.

IEEE 802 Networking Standards

The Institute of Electrical and Electronics Engineers (IEEE) defined LAN standards to ensure that network interfaces and cabling from multiple manufacturers would be compatible as long as they adhered to the same IEEE specification. This effort was called Project 802 to indicate the year (1980) and month (February) of its inception. Since then, the IEEE 802 specifications have taken firm root in the networking world. Because the OSI model wasn't standardized until 1983 to 1984, the IEEE 802 standards predate the model, as did TCP/IP.

Nevertheless, the two were developed in collaboration and are compatible with one another. (The IEEE is one of the U.S. participants in the ISO.)



TIP

For more information on the IEEE and its standards, visit www.ieee.org.

Project 802 concentrates its efforts on standards that describe a network's physical elements (the topics of Chapters 3 and 4), including NICs, cables, connectors, signaling technologies, media access control, and the like. Most of these elements reside in the lower two layers of the OSI model: Data Link and Physical. In particular, the 802 specification describes how NICs can access and transfer data across a variety of network media and what's involved in attaching, managing, and detaching these devices in a network environment.

IEEE 802 Specifications

The IEEE numbers the collection of 802 documents starting with 802.1, 802.2, and so forth. Each number after the dot represents a different technology or subset of a technology. When a technology is enhanced, such as Ethernet going from 10 Mbps to 100 Mbps, each enhancement is usually specified by letters after the number. For example, 802.3 is the original Ethernet, whereas 802.3u specifies 100BaseT Ethernet.

Table 6-2 lists the major 802 categories. For the purposes of this book, standards 802.3 and 802.11 are of the most interest because they define the most widely used technologies of Ethernet and Wi-Fi, although 802.15 (in Chapter 1's coverage of wireless PANs) and 802.16 (with Chapter 3's coverage of WiMAX) have already been discussed, too. The 802 standards aren't a static set of documents. New technologies and enhancements are added often, as with the soon-to-be-ratified 802.11n.



TIP

You can access IEEE 802 standards at www.ieee.org/publications_standards/index.html#IEEE_Standards. Most require a fee or subscription membership.

Table 6-2 IEEE 802 standards

Standard	Name	Explanation
802.1	Internetworking	Covers routing, bridging, and internetwork communication
802.2	Logical Link Control	Covers error control and flow control over data frames (inactive)
802.3	Ethernet LAN	Covers all forms of Ethernet media and interfaces, from 10 Mbps to 10 Gbps (10 Gigabit Ethernet)
802.4	Token Bus LAN	Covers all forms of token bus media and interfaces (disbanded)
802.5	Token Ring LAN	Covers all forms of token ring media and interfaces
802.6	Metropolitan Area Network	Covers MAN technologies, addressing, and services (disbanded)

(continues)

Table 6-2 IEEE 802 standards (continued)

Standard	Name	Explanation
802.7	Broadband Technical Advisory Group	Covers broadband networking media, interfaces, and other equipment (disbanded)
802.8	Fiber-Optic Technical Advisory Group	Covers use of fiber-optic media and technologies for various networking types (disbanded)
802.9	Integrated Voice/Data Networks	Covers integration of voice and data traffic over a single network medium (disbanded)
802.10	Network Security	Covers network access controls, encryption, certification, and other security topics (disbanded)
802.11	Wireless Networks	Sets standards for wireless networking for many different broadcast frequencies and techniques
802.12	High-Speed Networking	Covers a variety of 100 Mbps-plus technologies, including 100VG-AnyLAN (disbanded)
802.13	Unused	
802.14	Cable modems	Specifies data transport over cable TV (disbanded)
802.15	Wireless PAN	Covers standards for wireless personal area networks
802.16	Wireless MAN (WiMAX)	Covers wireless metropolitan area networks
802.17	Resilient Packet Ring	Covers emerging standards for very high-speed, ring-based LANs and MANs
802.18	Wireless Advisory Group	A technical advisory group that monitors radio-based wireless standards
802.19	Coexistence Advisory Group	A group that addresses issues of coexistence with current and developing standards
802.20	Mobile Broadband Wireless	A group working to enable always-on multivendor mobile broadband wireless access
802.21	Media Independent Handoff	A group working to enable handoff between wireless networks of the same or different types
802.22	Wireless Regional Area Network	Working to bring broadband access to hard-to-reach low-population areas
802.23	Emergency Services Working Group	A new group (March 2010) working to facilitate civil authority communication systems

IEEE 802 Extensions to the OSI Reference Model

The two lowest layers of the OSI model—the Physical and Data Link layers—define how computers attach to specific network media and specify how more than one computer can access the network without causing interference with other computers on the network. Project 802 took this work further to create the specifications (primarily 802.1 through 802.5) that define the most successful LAN technologies, including Ethernet and token ring, which together dominate the LAN world.



The IEEE 802 specification expanded the OSI model at the Physical and Data Link layers. Figure 6-9 shows how the 802 standards provide more detail by separating the Data Link layer into these sublayers:

- Logical Link Control (LLC) for error recovery and flow control
- Media Access Control (MAC) for access control

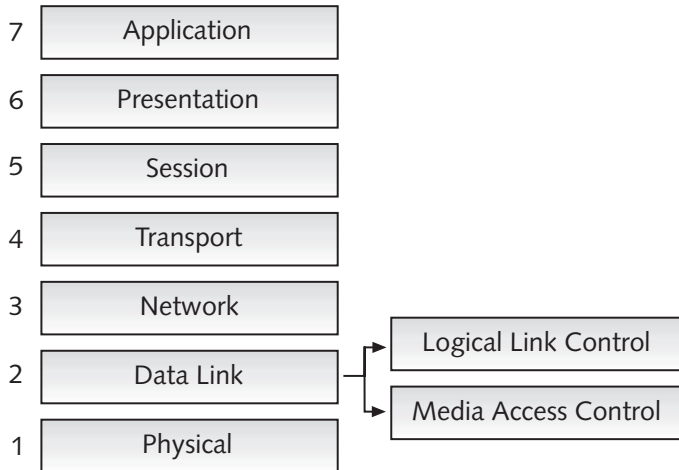


Figure 6-9 The IEEE 802 standard divides the OSI Data Link layer into two sublayers

Courtesy of Course Technology/Cengage Learning

The **Logical Link Control (LLC) sublayer** (defined by 802.2) controls data-link communication and defines the use of logical interface points, called service access points (SAPs), that other computers can use to transfer information from the LLC sublayer to the upper OSI layers. It's also responsible for error recovery in some situations and is the sublayer that communicates with the Network layer. There are several modes of LLC operation; some modes require the LLC to detect and recover from errors in transmission. This function is largely carried out in hardware on the NIC.

The **Media Access Control (MAC) sublayer** manages access to the physical medium and, therefore, communicates with the Physical layer. It communicates directly with a computer's NIC and is responsible for physical addressing. The physical address burned into every NIC is called a MAC address because it operates at this sublayer of the 802.2 specification. The MAC sublayer of the Data Link layer is where networking technologies such as Ethernet, token ring, FDDI, and so forth do their work.



Hands-On Project 6-1: Viewing Your MAC Address

Time Required: 10 minutes

Objective: Use different methods to view your MAC address.

Required Tools/Equipment: Your classroom computer

Description: In this project, you use your network connection properties and the `ipconfig` and `getmac` commands to view your MAC address.

1. Start your computer and log on as **NetAdmin**, if necessary.
 2. Open the Network and Sharing Center by clicking **Start, Control Panel**. Under Network and Internet, click **View network status and tasks**.
 3. In the left pane of the Network and Sharing Center, click **Change adapter settings**. Right-click **Local Area Connection** and click **Status** to open the Local Area Connection Properties dialog box.
 4. Click **Details**. In the Network Connection Details box, the physical address of your NIC (its MAC address) is the third item from the top. Write down the MAC address, and then click **Close** twice.
-
5. Open a command prompt window. Type `ipconfig /all` and press Enter. Your MAC address is listed under Local Area Connection as “Physical Address.” If you have more than one interface, each interface is listed along with the physical address.
 6. Type `getmac` and press Enter. Your MAC address is listed along with the Windows internal name for the interface. The `getmac` command displays the MAC address for all interfaces. Verify that the MAC address `getmac` displays is the same as in Step 4.
 7. Close all open windows, but leave your computer running for the next project.



Hands-On Project 6-2: Dragging and Dropping the OSI Model Layers

Time Required: 10 minutes

Objective: Put the OSI model layers in the correct order by using Simulation 12.

Required Tools/Equipment: Your classroom computer and Simulation 12 on the book’s CD

Description: In this project, you run Simulation 12 to place the OSI model layers in order.

1. Start your computer and log on, if necessary.
2. Insert the CD that came with the book. If the CD doesn’t autostart, navigate to the CD drive and double-click **Index.htm**.
3. Click **Simulation 12** in the menu.
4. Drag and drop the OSI model layers in their correct order.
5. Click **Menu** to go back to the simulation menu, and leave it open for the next project.



Hands-On Project 6-3: Matching OSI Model Descriptions to Layer Numbers

Time Required: 10 minutes

Objective: Match the OSI model layer descriptions to the correct layer numbers by using Simulation 13.



Required Tools/Equipment: Your classroom computer and Simulation 13 from the book's CD

Description: In this project, you run Simulation 13 to place the OSI model layer description in the correct layer number.

1. Start your computer and log on if necessary. Refer to Hands-On Project 6-2 if you closed the simulation or if the CD doesn't autostart.
2. Drag the OSI model layer description into their correct order by layer number.
3. Click **Menu** to go back to the simulation menu, and leave it open for the next project.



Hands-On Project 6-4: Creating a Frame

Time Required: 10 minutes

Objective: Create a frame by dragging and dropping frame headers in Simulation 14.

Required Tools/Equipment: Your classroom computer and Simulation 14 from the book's CD

Description: In this project, you run Simulation 14 to create a frame based on the information displayed.

1. Start your computer and log on, if necessary. Refer to Hands-On Project 6-2 if you closed the simulation or if the CD doesn't autostart.
2. Given the information in the simulation, drag frame headers to the correct places in the target frame placeholder.
3. Close all open windows.

Chapter Summary

- The OSI reference model and IEEE Project 802 define a frame of reference for networking and specify the operation of most networking technologies in current use. Together, these models describe the complex processes and operations involved in sending and receiving information across a network.
- The OSI reference model separates networking into seven layers, each with its own purposes and activities. From the top down, the layers are Application, Presentation, Session, Transport, Network, Data Link, and Physical. Most network products and technologies are also specified in terms of the layers at which they operate. The layers help describe the features and functions the products and technologies deliver.
- Following is a summary of the functions of each OSI model layer:
 - Application*—Provides access to network resources
 - Presentation*—Handles data formatting and translation
 - Session*—Manages ongoing conversations between two computers
 - Transport*—Breaks long data streams into smaller chunks (segments)
 - Network*—Provides best path selection and IP addressing

Data Link—Defines how computers access the media

Physical—Converts bits into signals and defines media and connectors

- The IEEE 802 project defines networking standards in more than 20 categories to ensure that network interfaces and cabling from different manufacturers are compatible. The IEEE 802.2 standard specifies the functions of a network's Physical and Data Link layers by dividing the Data Link layer into two sublayers: Logical Link Control (LLC) and Media Access Control (MAC). Together, these sublayers handle media access, addressing, and control (through the MAC sublayer) and provide reliable, error-free delivery of frames from one computer to another.

Key Terms

access control In the context of the Network layer and routing, the process by which a router consults a list of rules before forwarding an incoming packet. The rules determine whether a packet meeting certain criteria (such as source and destination address) should be permitted to reach the intended destination.

Application layer Layer 7 in the OSI model provides interfaces that enable applications to request and receive network services. *See also* Open Systems Interconnection (OSI) reference model.

Data Link layer Layer 2 in the OSI model is responsible for managing access to the network medium and delivery of data frames from sender to receiver or from sender to an intermediate device, such as a router. *See also* Open Systems Interconnection (OSI) reference model.

deencapsulation The process of stripping the header from a PDU as it makes its way up the communication layers before being passed to the next higher layer. *See also* protocol data unit (PDU).

encoding Representing 0s and 1s as a physical signal, such as electrical voltage or a light pulse.

International Organization for Standardization (ISO) The international standards-setting body based in Geneva, Switzerland, that sets worldwide technology standards.

Logical Link Control (LLC) sublayer The upper sublayer of the IEEE Project 802 model for the OSI model's Data Link layer. It handles error-free delivery and controls the flow of frames between sender and receiver across a network.

maximum transmission unit (MTU) The maximum frame size allowed to be transmitted across a network medium.

Media Access Control (MAC) sublayer The lower sublayer of the IEEE Project 802 model for the OSI model's Data Link layer. It handles accessing network media and mapping between logical and physical network addresses for NICs.

Network layer Layer 3 of the OSI model handles logical addressing and routing of PDUs across internetworks. *See also* Open Systems Interconnection (OSI) reference model *and* protocol data unit (PDU).

Open Systems Interconnection (OSI) reference model ISO Standard 7498 defines a frame of reference for understanding networks by dividing the process of network communication into seven layers. Each layer is defined in terms of the services and data it handles on behalf of the layer above it and the services and data it needs from the layer below it.



peer communication In the layered approach, each layer on one computer behaves as though it were communicating with its counterpart on the other computer. This means each layer on the receiving computer sees network data in the same format its counterpart on the sending computer did.

Physical layer Layer 1, the bottom layer of the OSI model, transmits and receives signals and specifies the physical details of cables, NICs, connectors, and hardware behavior. *See also* Open Systems Interconnection (OSI) reference model.

Presentation layer At Layer 6 of the OSI model, data can be encrypted and/or compressed to facilitate delivery. Platform-specific application formats are translated into generic data formats for transmission or from generic data formats into platform-specific application formats for delivery to the Application layer. *See also* Open Systems Interconnection (OSI) reference model.

protocol data unit (PDU) A unit of information passed as a self-contained data structure from one layer to another on its way up or down the network protocol stack.

Session layer Layer 5 of the OSI model is responsible for setting up, maintaining, and ending communication sequences (called sessions) across a network. *See also* Open Systems Interconnection (OSI) reference model.

Transport layer Layer 4 of the OSI model is responsible for reliable delivery of data streams across a network. Layer 4 protocols break large streams of data into smaller chunks and use sequence numbers and acknowledgements to provide communication and flow control. *See also* Open Systems Interconnection (OSI) reference model *and* protocol data unit (PDU).

Review Questions

1. The original commercial version of Ethernet supported 10 Mbps bandwidth; the version introduced in the early 1990s supports 100 Mbps; and in 1998, Gigabit Ethernet was introduced. All versions use the same data frame formats, with the same maximum PDU sizes, so they can interoperate freely. Given this information and what you know of layered technologies, which of the following statements is true? (Choose all that apply.)
 - a. Ethernet works at the Data Link and Physical layers of the OSI model, and upgrades to newer, faster versions of Ethernet can be made by changing only the components that work at these layers.
 - b. Ethernet spans several layers and requires a new protocol stack to upgrade to new versions.
 - c. Changes in technology at one layer of the OSI model don't usually affect the operation of other layers.
 - d. Ethernet isn't considered a scalable technology.
2. The addition of information to a PDU as it's passed from one layer to the next is called which of the following?
 - a. PDI transforming
 - b. Encapsulation
 - c. Deencapsulation
 - d. Converting

3. Layers acting as though they communicate directly with each other across the network are called which of the following?
 - a. Partners
 - b. Synchronous
 - c. Interchangeable
 - d. Peers
4. Place the following letters in the correct order to represent the OSI model from Layer 7 to Layer 1:
 - a. Presentation
 - b. Data Link
 - c. Session
 - d. Physical
 - e. Application
 - f. Transport
 - g. Network
5. Which OSI layer creates and processes frames?
6. Which OSI layer handles flow control, data segmentation, and reliability?
 - a. Application
 - b. Physical
 - c. Transport
 - d. Data Link
7. Which OSI layer governs how a NIC is attached to the network medium?
8. Which OSI layer determines the route a packet takes from sender to receiver?
 - a. 7
 - b. 1
 - c. 3
 - d. 4
9. Which OSI layer is responsible for setting up, maintaining, and ending ongoing information exchanges across a network?
 - a. 6
 - b. 3
 - c. 2
 - d. 5



10. Which of the following elements might the Data Link layer add to its PDU? (Choose all that apply.)
 - a. Physical addresses
 - b. Logical addresses
 - c. Data
 - d. CRC
11. When and how many times is a CRC calculated?
 - a. Once, before transmission
 - b. Once, after receipt
 - c. Twice, once before transmission and again on receipt
 - d. At the source and destination and at each intermediary device
12. Which layer of the OSI model does Project 802 divide into two sublayers?
 - a. Physical
 - b. Data Link
 - c. Network
 - d. Session
13. What are the names of the sublayers specified as part of Project 802? (Choose all that apply.)
 - a. Data Link Control (DLC)
 - b. Logical Link Control (LLC)
 - c. Carrier Sense Multiple Access/Collision Detection (CSMA/CD)
 - d. Media Access Control (MAC)
14. Which term refers to stripping header information as a PDU is passed from one layer to a higher layer?
 - a. Deencapsulation
 - b. Encapsulation
 - c. PDU stripping
 - d. Packetization
15. Which IEEE 802 standard applies to Ethernet?
 - a. 802.2
 - b. 802.3
 - c. 802.4
 - d. 802.5
 - e. 802.11

16. Which IEEE 802 standard applies to wireless LANs?
 - a. 802.2
 - b. 802.3
 - c. 802.4
 - d. 802.5
 - e. 802.11
17. What is the name of the PDU at the Transport layer?
 - a. Bit
 - b. Packet
 - c. Segment
 - d. Data
18. At which OSI layer does the PDU contain sequence and acknowledgement numbers?
 - a. Application
 - b. 4
 - c. Data Link
 - d. 6
19. Which of the following is an example of software found at the Application layer? (Choose all that apply.)
 - a. FTP
 - b. TCP
 - c. HTTP
 - d. ICMP
20. At which Data Link sublayer does the physical address reside?
 - a. Media Access Control (MAC)
 - b. Logical Link Control (LLC)
 - c. Data Access Control (DAC)
 - d. Network Access Control (NAC)
21. Which of the following problems can occur at the Physical layer?
 - a. NIC driver problems
 - b. Incorrect IP addresses
 - c. Signal errors caused by noise
 - d. Incorrect segment size



Challenge Labs



Challenge Lab 6-1: Identifying OSI Model Layers from Captured Packets

Time Required: 15 minutes

Objective: Use Wireshark to capture the packets generated from an HTTP communication session. Identify the OSI model layers represented by the headers in the captured files.

Required Tools/Equipment: Your classroom computer with Wireshark installed and Internet access

Description: Using Wireshark and an appropriate capture filter, capture the packets involved in an HTTP session that you start by opening a Web page. Select an HTTP packet, and using the headers in the middle pane, perform the following tasks:

- Map the header names in the captured packet to the layers of the OSI model.
- For each header, find two fields you can identify as pertaining to that OSI layer's function, and be prepared to explain why.
- Write the information you derived from the previous items and be prepared to turn it in to your instructor or discuss it in class.



Challenge Lab 6-2: Listing MAC Addresses in Your Network

Time Required: 15 minutes

Objective: Find MAC addresses for the computers, printers, routers, and other devices on your network.

Required Tools/Equipment: Your classroom computer

Description: Your boss has asked you to get a list of MAC addresses in your network. He has told you that the addresses of all devices are in the range 192.168.100.1 through 192.168.100.20. (If these addresses aren't the ones you're using, substitute the actual addresses.) Because people are working at their computers, he expects you to accomplish this task without leaving your computer. List the steps you took to get all the MAC addresses.

Case Projects



Case Project 6-1

The OSI model is a useful tool in troubleshooting a network because it enables you to isolate a problem to a particular software module or piece of hardware. In this project, after reading the description of a problem, identify the OSI model layer or layers that are most likely involved.

- A computer won't connect to the network. After some investigation, you find that the patch cable isn't terminated correctly.

- A computer can access resources on the local LAN but not on a different subnet. You find that the computer's default gateway isn't configured correctly.

- You can ping a computer you're trying to transfer files to via FTP, but you can't communicate by using FTP.

- All computers connected to a particular hub have lost network connectivity. You determine that the hub is the problem.

- You receive an encrypted text file, but when you open it, the text is unreadable. You determine that decryption didn't take place as it should have.

- You check some statistics generated by a network-monitoring program and discover that an abnormally high number of CRC errors were detected.

- One of your servers has been exhibiting sluggish network performance. You use a network-monitoring program to try to evaluate the problem. You find considerable TCP retries occurring because the server is being overwhelmed by data, and packets are being discarded.

- A user is trying to connect to another computer, but the logon attempt is continually rejected.

- You try to access a Linux server to share files by using NFS. You can communicate with the server, but the shared files don't appear to be available.

- You inspect a computer that isn't able to communicate with other computers. You find that IPv6 instead of IPv4 is installed on that computer.



Case Project 6-2

Your instructor might want you to organize in groups for this project. This chapter included a few real-world examples that use a layered approach to describing a process. See whether you can come up with another process that can be described in layers. You should give a presentation to the class with a detailed description of the layered process you decide on.

Case Project 6-3

You want to transfer a document from one computer to another, and *you want the document to be encrypted*. The destination computer is on another network, so you know *data has to travel through one or more routers*. The network technology on *your network is Ethernet*, but the technology on *the destination network is Wi-Fi*. From what you have learned about networking, should this document transfer work? Why or why not? Which layers of the OSI model are involved in the italicized parts of this description?

Network Hardware in Depth

After reading this chapter and completing the exercises, you will be able to:

- Describe the advanced features and operation of network switches
- Describe routing table properties and discuss routing protocols
- Explain basic and advanced wireless access point features
- Select the most suitable NIC bus and features for a computer

Network packets travel through a variety of network devices on their journey between sender and receiver. These devices vary from the simplest of hubs requiring no configuration to complex routers and switches that have a multitude of configuration settings and fine-tuning options.

This chapter begins with switches because although some high-end hubs can be configured for the purposes of network management, hubs are declining in use to the point of near obsolescence. Routers and routing protocols are discussed next, and you learn more about a routing table and the protocols used to build it. In the section on wireless access points, you learn about some configuration and security options available for wireless networks. Finally, the discussion of NICs focuses on some high-performance and enterprise features you might see on server NICs and workstation NICs in a large internetwork.

Network Switches in Depth

Network switches at their simplest are plug-and-play devices. Apply power to a switch, and it's ready to move frames from one device to another. However, advanced features on some switches enable administrators to fine-tune their networks for optimal operation. Before getting into these features, reviewing some properties common to all switches is helpful:

- Switches work at the Data Link layer (Layer 2) of the OSI model. At the Data Link layer, physical addresses are defined, and the PDU at this layer is the frame. Switches receive frames on one port and forward the frames out the port where the destination device can be found.
- Switches **flood** broadcast frames out all ports. When a switch receives a broadcast frame (defined as a destination address of all binary 1s, or FF:FF:FF:FF:FF:FF), it forwards the incoming frame out all connected ports except the port where the frame was received. When a switch forwards a frame out all ports, it's referred to as "flooding the frame."
- Each switch port is considered a collision domain (see the circled areas in Figure 7-1). If a collision occurs between the switch port and the devices connected to a switch port, the switch doesn't forward collision information to any of its other ports. This behavior limits a collision's effects to only the devices connected to the port, so each port is called a collision domain.
- Switch ports can operate in full-duplex mode, allowing connected devices that also support full-duplex to transmit and receive simultaneously, thereby eliminating the possibility of a collision. However, even though a collision can't occur on a port in full-duplex mode, each port is still referred to as a collision domain.



With switches, collisions can occur only if a switch port is connected to a computer running in half-duplex mode or if the switch port is connected to a hub.

Now that you have the basics of switch operation, the following sections cover a few properties of switches in more detail.




Figure 7-1 Each switch connection is a collision domain

Courtesy of Course Technology/Cengage Learning

Switch Port Modes of Operation

Most switches have the capability to work in multiple modes. Ports on a switch, for instance, can usually operate at multiple speeds and multiple duplex modes. For example, ports on a typical 10/100 Mbps switch can usually operate in these modes:

- 10 Mbps half-duplex
- 100 Mbps half-duplex
- 10 Mbps full-duplex
- 100 Mbps full-duplex

Most inexpensive switches run in **auto-negotiate mode**, which means the switch exchanges information with a device connected to a port and attempts to set the port's operating mode to the highest performance setting the device supports. If the device is set to operate at a particular speed and duplex mode (a setting configured in a NIC's properties, for example), the switch sets the port to match the connected device's settings. Occasionally, auto-negotiate fails and a link is never achieved. In this case, one or both devices should be set to a particular mode instead of relying on auto-negotiate. Mid-range and high-end switches allow configuring each port to the mode that works best for the connected device.

Another option you sometimes find is **auto-MDIX** (media-dependent interface crossed), in which the switch port detects the type of device and cable it's connected to. If necessary, the port swaps its transmit and receive pins, which enables you to use a straight-through or crossover cable regardless of the type of device you're connecting to the port. If each port on the switch can be configured separately, the auto-MDIX feature can usually be enabled or disabled.

If port configuration on switches is important in your environment, you need to invest more money than if you simply want to rely on switches' capability to configure their ports automatically.



Creating the Switching Table

Chapter 2 discussed the basics of a switching table. This section explains more details of creating and maintaining the table. A switching table is composed primarily of MAC address/port pairs that tell the switch where to forward a frame, based on the frame's destination MAC address. When a switch is first powered on, however, the switching table is empty because the switch hasn't yet learned which devices are connected to which ports.

As network devices begin to send frames throughout the network, the switch reads each frame's source address and adds it to the switching table along with the port it was received from. Each frame's destination address is searched for in the switching table and, if found, is forwarded out the corresponding port. However, what if the frame's destination address isn't found in the switching table? The switch does the reasonable thing and floods the frame.

The switching table isn't limited to a single MAC address per port. The technical specifications for most switches usually include the number of MAC addresses the switch supports. This number is usually in the thousands and is often expressed with K, as in "8K MAC addresses supported." For example, if your network looks something like Figure 7-2, Switch1's switching table would be similar to the table in this figure. In the figure, the two switches are connected, so Switch1 must forward frames destined for any computers connected to Switch2 out the port where Switch1 is connected to Switch2 (in this case, port 1).

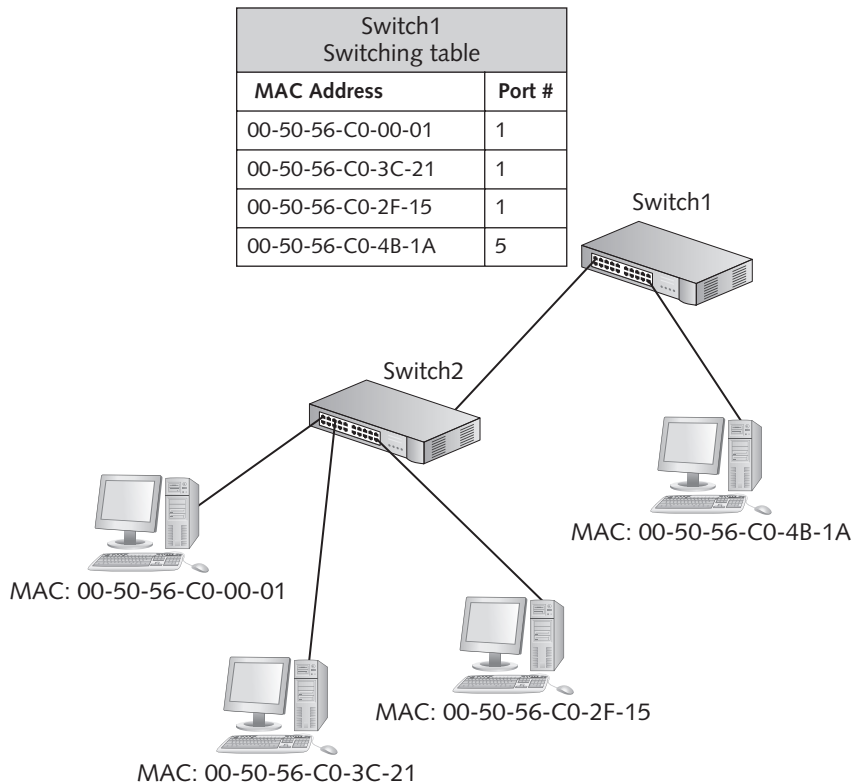


Figure 7-2 Switching tables can contain multiple MAC addresses per port

Courtesy of Course Technology/Cengage Learning

As you learned in Chapter 2, a switching table prevents stale entries by including a timestamp when a switching table entry is created. When a switch receives a frame, the entry for the frame's source address is updated with a new timestamp. Periodically, the switch inspects its switching table for expired entries. Switches typically maintain a MAC address for a period called the **aging time**, which is usually several minutes. If an entry's timestamp isn't updated within the aging time, it expires and is removed from the table.

Frame Forwarding Methods

Switches receive a frame on one port and forward it out another port with a variety of methods. The simplest and fastest is **cut-through switching**, in which the switch reads only enough of the incoming frame (in Ethernet, 12 bytes) to determine the frame's source and destination addresses. After the forwarding location is determined, the frame is switched internally from the incoming port to the outgoing port, and the switch is free to handle additional frames. The benefit of cut-through switching is speed. A typical Ethernet frame can be up to 1518 bytes. With cut-through switching, the switch reads only a small portion of the frame's contents before sending the frame on its way. The disadvantage of this switching method is that the switch indiscriminately forwards frames containing errors, so it ties up bandwidth needlessly with frames that will be discarded.

On the other hand, **store-and-forward switching** requires that the switch read the entire frame into its buffers before forwarding it. The switch first examines the frame's frame check sequence (FCS) field to make sure it contains no errors before it's forwarded. If an error is found, the switch discards the frame. This method has the advantage of conserving bandwidth when many frames contain errors. The disadvantage is that the entire frame must be read, stored in memory, and examined before it can be forwarded. This process takes time and slows the network slightly.

A third popular switching method is **fragment-free switching**, in which the switch reads enough of the frame to guarantee that it's at least the minimum size for the network type. For Ethernet, it's 64 bytes. One type of frame error that can occur in a network is a **frame fragment**, meaning the frame is damaged because of a collision or a malfunctioning device, such as a NIC or hub. When this type of damage occurs, the frame might be truncated to less than the minimum allowable size. A switch operating in fragment-free mode detects this problem and discards the frame without forwarding it. Table 7-1 summarizes these switching methods.

Table 7-1 Switching method summary

Switching method	Switching performance	Errors forwarded
Cut-through	Fastest	All errors forwarded
Fragment-free	Medium	All errors except frame fragments forwarded
Store-and-forward	Slowest	No error frames forwarded

High-end switches can combine the best features of these switching methods. For example, they can initially operate in cut-through mode for the best performance. However, if they detect frequent errors, they can change to store-and-forward mode, thereby decreasing the number of propagated errors. If the error rate decreases enough, the switch can be put in cut-through mode again.



Advanced Switch Features

All switches have a main objective: to receive frames on one port and forward them out another port to arrive at the destination device eventually. As you have learned, there are different methods for performing this task and different speeds at which it's accomplished. However, high-end switches, often referred to as “smart switches” and “managed switches,” can offer more features to help you design an efficient, reliable network. There are too many advanced options on **managed switches** to cover all of them in this book, but this section gives you an overview of the most common features in configurable **smart switches**:

- Multicast processing
- Spanning Tree Protocol
- Virtual local area networks
- Port security

Multicast Processing You know how a switch handles unicast and broadcast frames, but how does it handle a multicast frame? A multicast frame contains as its destination a special address that signifies one or more computers or devices. The application waiting for the frame determines this address. For example, some disk-imaging programs can use multicast frames. When a computer classroom or lab is configured, one computer can be configured with all the applications the classroom needs. Then an exact copy of the disk is made and stored on a server. Next, the image is transferred by using multicast frames to only the computers that require it and are running the application that's “listening” for a particular multicast address. The image is sent only once, and only the computers running the disk-imaging program receive the frames.

Now that you have an idea how multicast frames can be used, there are two ways switches can process them:

- By treating them as broadcasts and flooding them to all ports
- By forwarding the frames only to ports that have registered the multicast address

The first method is used by low-end switches that don't have specific multicast support and by switches that support multicast but haven't been configured for it. The second method is used by switches that support Internet Group Management Protocol (IGMP), specifically IGMP “snooping.” Multicast MAC addresses always begin with 01:00:5E, leaving the rest of the MAC address to identify a particular multicast application. When a switch sees MAC addresses containing 01:00:5e as the first 24 bits arriving on a port, it registers that port as belonging to the multicast group (specified by the remaining 24 bits of the multicast MAC address). When the switch sees a multicast frame, it forwards it out only registered ports. The details of IGMP are beyond the scope of this book, but you should know that the protocol allows a computer to “unregister” the multicast application so that the switch can remove it from the switching table.

Spanning Tree Protocol Some switches, unlike hubs, are designed to accommodate redundancy, or multiple paths, in the network. However, using redundant switches can also cause a network administrator's worst nightmare: a **switching loop** (also referred to as a “bridging loop”). A switching loop occurs when switches are connected in such a way that it's possible for frames to be forwarded endlessly from switch to switch in an infinite loop. Figure 7-3 shows a network configuration in which this problem could happen.




Figure 7-3 Switches with redundant paths can create a switching loop

Courtesy of Course Technology/Cengage Learning

Because a switch forwards broadcast frames out every port except the port the frame was received on, a broadcast frame originating from a computer attached to Switch A is forwarded to Switches B and C. Switches B and C forward the frame to Switch D. Switch D forwards the frame received from Switch B to Switch C and the frame received from Switch C to Switch B. Switches B and C then forward the frame to Switch A, which starts the process again. It continues until the switches are just forwarding the same broadcast frame repeatedly while causing every computer in the network to process the frame. If this type of loop occurs with a broadcast frame, it's referred to as a **broadcast storm**.

Luckily, this problem doesn't usually happen. IEEE 802.1D specifies the Spanning Tree Algorithm to prevent this behavior. It requires that switches communicate with one another. The protocol used to communicate between switches is **Spanning Tree Protocol (STP)**, which enables switches to detect when there's a potential for a loop. When this happens, one of the switch ports goes into **blocking mode**, preventing it from forwarding frames that would create a loop. If the loop configuration is broken, perhaps because of a switch failing or the connection between two switches failing, the switch that was in blocking mode resumes forwarding frames. In this way, redundancy is achieved, allowing frames to reach their destination in the event of a switch or media failure but preventing the disastrous affect of switching loops. Simulation 15 shows how a switch uses STP to prevent a switching loop.



Simulation 15: STP prevents switching loops

STP is an integral part of most mid-range to high-end switches. Beware of SOHO (small office/home office) products, as most of these lower-end switches don't support STP. Before using a switch in a configuration that could form a loop, be certain that all involved switches support the 802.1D standard for STP.

One side effect of STP is that devices take a bit longer to create a link with a switch that runs the protocol. The reason is that after the switch detects that a new device is plugged into a port, STP begins transferring packets to determine whether the new device is another switch that could cause a loop. When the switch determines that the new device won't cause a switching loop, the link can be established. This process usually takes several seconds. With many mid-range to high-end switches supporting STP, administrators can disable the protocol on specific ports, so if you know that certain ports will never be used to connect another switch, it's safe to turn STP off on these ports.

Virtual Local Area Networks Switches that support **virtual local area networks (VLANs)** enable you to configure one or more switch ports into separate broadcast domains. A switch with two or more VLANs configured is effectively divided into logically disconnected networks. In other words, it's like separating the switch into two or more switches that aren't connected to one another. So how do you communicate between these virtual networks, or broadcast domains? A router is needed to communicate between VLANs.

A switch that enables you to create multiple broadcast domains can offer many benefits for your network. A switch with VLAN capability can optimize your network configuration by creating broadcast domains without having to add switches. This capability can be an advantage for your network, as it improves management and security of the network and gives you more control of broadcast frames.

Because a VLAN divides the network into more broadcast domains, devices on switch ports belonging to different VLANs must have logical addresses (IP addresses, for example) on different networks. Furthermore, if devices on one VLAN are to communicate with devices on another VLAN, a Layer 3 device (typically a router) is required to route packets between VLANs. Because routers are slower devices than switches or hubs, you should plan your network so that most resource accesses occur within a VLAN. Figure 7-4 shows a network divided into two VLANs, with a router communicating between them. Notice that there's a server on each VLAN so that workstation traffic doesn't need to cross the router for access to a server. You must consider the ramifications of moving a workstation from one switch port to another. When adding a workstation to your network, you need to know which VLAN the workstation should belong in before choosing a switch port to connect the station and a logical address to assign to the workstation.

In addition to dividing a network into multiple broadcast domains, VLANs make it possible for network administrators to group users and resources logically instead of by physical location. With conventional networks, a user or resource's location dictates which network it's assigned to. This limitation sometimes makes resource sharing inefficient because ideally, users are assigned to the same network as the resources they access most often.

A switch supporting VLANs allows assigning any switch port or group of ports to a VLAN. Suppose you have a group of employees from different departments working on a long-term project. A new server has been allocated for this project, but the employees working on the




Figure 7-4 A network using VLANs

Courtesy of Course Technology/Cengage Learning

project are scattered in different buildings. To solve this problem, you can assign switch ports in each building to the same VLAN in which the server is configured. In this way, these employees and the resources they share, although physically separated, are logically grouped by using VLANs. Figure 7-5 shows how users and resources from different physical locations can be assigned to the same VLAN.



TIP

Although the details of implementing VLANs are beyond the scope of this book, you can read a good overview on the subject at <http://computer.howstuffworks.com/lan-switch16.htm>.

VLAN Trunks If you study Figure 7-5, you realize there must be a way to communicate between devices in the same VLAN. There are two possible configurations for doing so. One method is connecting a cable from a port in each VLAN on one switch to a port in each VLAN on another switch. This method is impractical because of the amount of cabling and the number of ports needed to make the VLAN connections. VLAN trunk ports take care of this problem. A **trunk port** is a switch port configured to carry traffic from all VLANs to

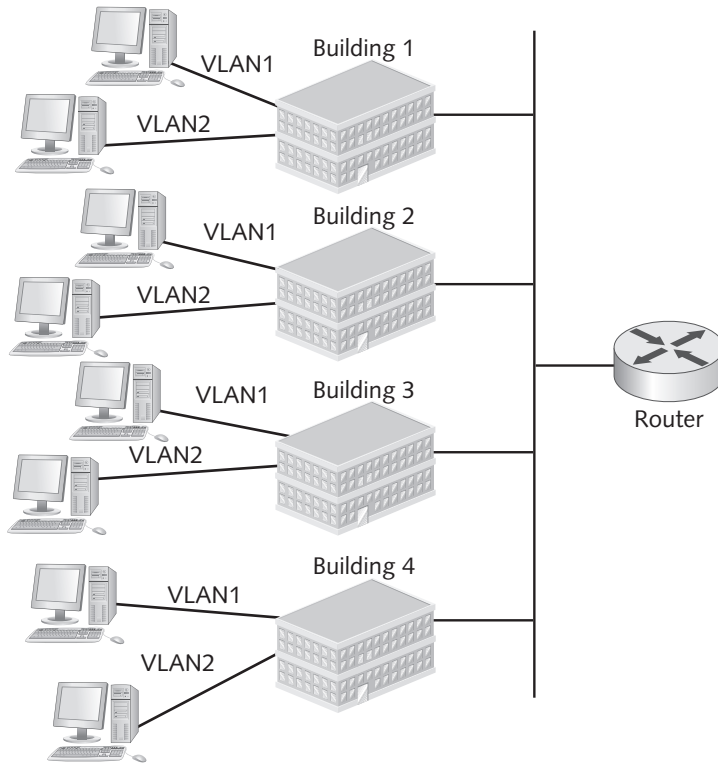


Figure 7-5 VLANs logically group users and resources from different physical locations

Courtesy of Course Technology/Cengage Learning

another switch or router. The switch or router port must also be configured as a trunk port. IEEE 802.1Q defines how to configure a trunk port, which involves the switch adding a tag to each frame that must traverse the trunk port. The VLAN tag identifies which VLAN the traffic originated from, so the traffic stays in that VLAN unless it's being routed. After the switch or router connected to the other side of the trunk cable receives the frame, the tag is removed from the frame before it's forwarded. Simulation 16 demonstrates how trunk ports work to allow VLAN traffic to travel from switch to switch.



Simulation 16: How switches use trunk ports with VLANs

Factors to Consider Before Using VLANs You might be tempted to use VLANs because the prospect of optimizing your network is enticing. However, the overuse of VLANs can end up costing you more than it benefits you. In addition, because more VLANs mean more logical networks, your network will be more complex. Furthermore, they can actually slow down your network when your intention is to increase performance.

Because VLANs need a router for communication, every VLAN you create requires a corresponding router interface, which usually means more expense. More router interfaces mean

additional IP networks, which is likely to require subnetting your existing network, and this task can become quite complex. If you need to change your existing network addressing scheme, you'll probably have to reconfigure many devices to reflect the new addressing scheme, which can take a lot of time. Finally, having more smaller VLANs can slow your network unless workstations and the network resources they access most of the time are in the same VLAN. So although VLANs can help optimize and organize your network, make sure you have a carefully planned network design before using them.



Remember that for a workstation to communicate outside its VLAN, it must go through a router, and communication through a router is always slower than through a switch.

Switch Port Security In some public buildings, such as libraries and schools, controlling network access is difficult sometimes. In particular, network jacks with connections to switches are often available to public users who can easily plug in a laptop computer that could contain viruses, hacker tools, and other malware. A switch with port security features can help prevent this type of connection.

Port security on most switches enables an administrator to limit how many and which MAC addresses can connect to a port. If an unauthorized computer with an unauthorized MAC address attempts to connect to the switch port, the port can be disabled and a message can be generated to indicate the violation. Alternatively, administrators can disable switch ports entirely until they're needed, giving them complete control over when a port can be used and by whom.



Hands-On Project 7-1: Observing a Switching Loop

Time Required: 20 minutes

Objective: Observe what happens in a network where two switches are connected in a way that creates a switching loop.

Required Tools/Equipment: Two computers (one with Wireshark installed), two switches that don't have STP enabled, two patch cables, and two crossover cables (or four patch cables if your switches support auto-MDIX)

Description: In this project, you connect two computers to two separate switches. First, you connect the two switches with a single cable and verify that you can ping from one computer to the other. Next, you connect the switches with a second cable, which creates the switching loop. Then you clear the ARP cache so that an ARP broadcast is created and ping from one computer to another while capturing ARP packets. This project can be done as a demonstration by the instructor or in groups.

1. Configure the network as shown in Figure 7-6. Use a crossover cable between switches unless the switches support auto-MDIX, in which case a regular patch cable works. Designate one computer as Computer1 and the other as Computer2. Start both computers and log on with an administrator account.
2. Configure Computer1 with IP address 192.168.1.1/24 and Computer2 with IP address 192.168.1.2/24. No default gateway or DNS server address is required.
3. If you're using Windows 7, the network is configured as a Public network by default, which means the firewall blocks ping packets. On both computers, turn off the firewall.




Figure 7-6 Network setup for Hands-On Project 7-1

Courtesy of Course Technology/Cengage Learning

4. On Computer1, open a command prompt window as administrator. To do so, click **Start**, point to **All Programs** and **Accessories**, right-click **Command Prompt**, and click **Run as administrator**.
5. Ping Computer2 (192.168.1.2) to verify that your network is working correctly. If the ping isn't successful, check all connections and settings and verify that the firewall is turned off for the Public profile on both computers. After you have pinged Computer2 successfully, continue to Step 6.
6. On Computer1, type `arp -d` at the command prompt and press **Enter** to delete the current ARP cache.
7. Using the second crossover cable (or patch cable), connect the two switches a second time to create a switching loop on switches that don't support (or haven't enabled) STP.
8. On Computer1, ping Computer2. You should see Wireshark begin to capture packets—many of the same packets over and over. Your computer might even freeze. What you're seeing is a broadcast storm caused by a switching loop. Depending on your switches, you might see the number of ARP packets captured slow or stop, but you'll likely see hundreds or thousands of ARP packets caused by the loop. The ping probably won't be successful.
9. Accidental switching loops can be created by cabling errors, and if they do occur, you can see how this condition can bring a network down or slow it dramatically. If you need redundant switch paths, using switches that support STP is critical, but even if you don't need redundant paths, switches equipped with STP can prevent an accidental loop. Disassemble the network and turn off the computers.

Routers in Depth

As you know, routers operate at the Network layer (Layer 3) and work with packets as the PDU. These advanced devices connect separate logical networks to form an internetwork. A router can be used to facilitate exchanging data between networks, but broadcast frames are

kept in their respective networks. The Internet uses routers to interconnect thousands of networks around the world. If it interconnected networks with bridges or switches, which forward broadcast frames, any broadcast frame generated by any computer connected to the Internet would be forwarded to and processed by every other computer on the Internet. If this happened, the only traffic flowing on the Internet would be broadcasts!

In addition to dividing large networks into smaller broadcast domains, routers are used to create complex internetworks so that LANs in a large international corporation, for example, can communicate efficiently. As shown in Figure 7-7, you can use routers to create complex internetworks with multiple paths between networks; these multiple paths are used for fault tolerance and load sharing. If a network link goes down, an alternate path can be chosen to get a packet to its destination. Furthermore, if a path becomes congested, additional paths can be used to ease congestion.

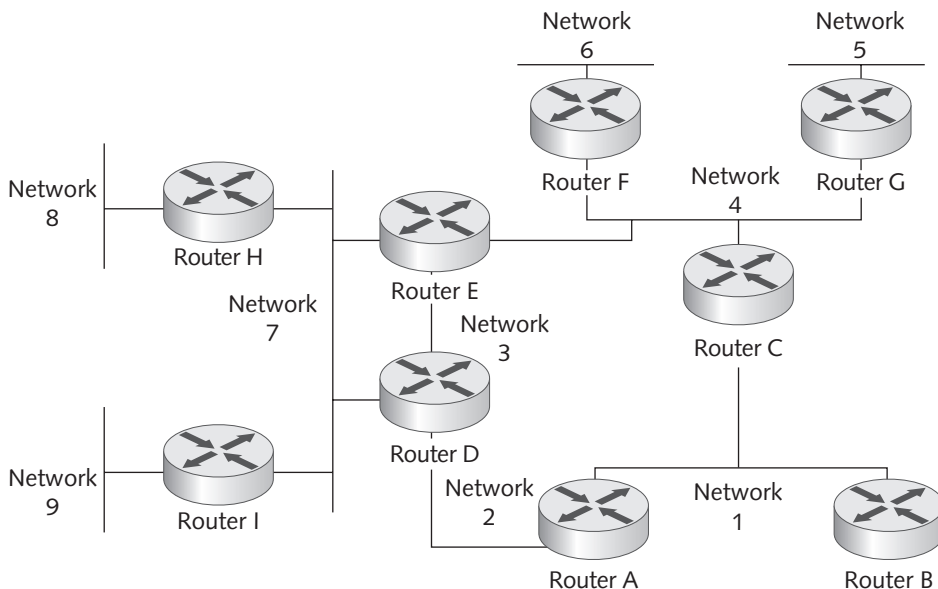


Figure 7-7 Routers can connect networks with many different paths between them

Courtesy of Course Technology/Cengage Learning

Routers are also used to control network access by inspecting source and destination address information of packets they handle. Based on rules an administrator defines, a router can forward a packet or discard it. All processing done by routers depends on the following features found on most routers:

- Router interfaces
- Routing tables
- Routing protocols
- Access control lists

The following sections describe these features and how routers use them to create effective internetworks.

Router Interfaces

Routers must have two or more interfaces, or ports, to be able to take packets coming from one network and forward them to another network. Each interface on a router has full Layer 3 functionality, including both an IP address and a MAC address. In fact, you can look at a router interface as just a NIC with the IP protocol bound to it.

When a router interface receives a frame, it performs the Data Link layer function of comparing the destination MAC address with the interface’s MAC address. If they match, the router reads the frame and strips the frame header and trailer; if they don’t match, the router discards the frame. Next, the router checks the resulting packet’s destination IP address. If this address matches the IP address of the interface on which the packet was received, the packet was intended for the router, which simply processes the packet. If the destination IP address’s network ID doesn’t match the interface address’s network ID, the router knows the packet should be routed to another network.

It then consults its routing table to determine how to get the packet to its destination and moves the packet from the incoming interface to the interface that will get the packet to its destination—the outgoing interface. The process of moving a packet from the incoming interface to the outgoing interface is called **packet forwarding**, or just “forwarding.” Before the packet can be sent out the outgoing interface, however, it must be encapsulated in a new frame header and trailer. The new frame header contains the outgoing interface’s MAC address as the source and the MAC address of the destination computer or the next router in the path as the destination. In addition, a new CRC is calculated and placed in the FCS field of the frame trailer. This process is shown in Figure 7-8.

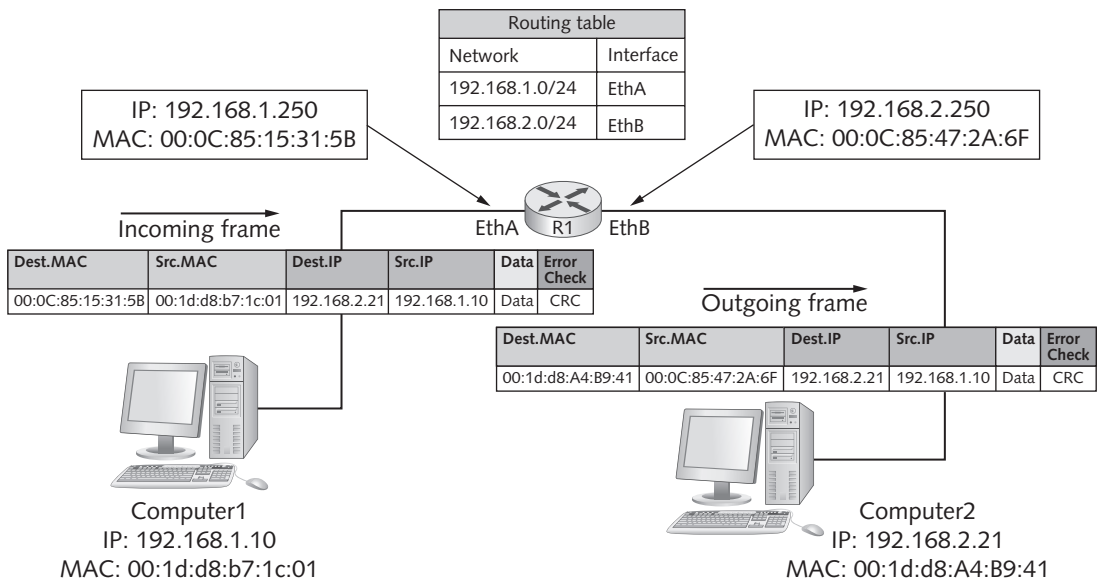


Figure 7-8 Packets are forwarded from one network to another

Courtesy of Course Technology/Cengage Learning

The following steps summarize how a router uses its interface to forward packets from one network to another:

1. Router receives a frame on an interface.
2. Router checks the frame's destination MAC address.
3. If the destination MAC address matches the interface's address, the router reads the frame; otherwise, the frame is discarded.
4. The frame header and trailer are stripped to create a packet.
5. The destination IP address is checked.
6. If the IP address's network ID is different from the interface's network ID, the packet should be routed.
7. The router consults the routing table to determine to which of its interfaces the packet should be forwarded.
8. The packet is encapsulated in a new frame header and trailer.
9. The packet is forwarded to the destination computer or the next router in the path.

**TIP**

A configuration called "a router on a stick" is sometimes used when a router is routing between VLANs on a switch. In this configuration, a single router interface is used to connect to a switch trunk port that carries traffic going from one VLAN to another. The router interface is configured with multiple IP addresses, each in a different network. In this setup, the router has multiple logical interfaces, but the incoming and outgoing interfaces are physically the same.

**NOTE**

Simulation 9: The changing frame header, first referenced in Chapter 5, also shows how a packet is encapsulated with a new MAC address at each router hop.

Routing Tables

As discussed in Chapter 2, routing tables are composed mainly of network address and interface pairs that tell the router which interface a packet should be forwarded to so that it gets to its destination network. Of course, there's more to the story than the overview in Chapter 2. The routing table in most routers contains the following information for each table entry:

- *Destination network*—The network address of a network to which the router can forward packets is called the **destination network**. It's usually expressed in CIDR notation, such as 172.16.0.0/16. When a router receives a packet on one of its interfaces, it compares the packet's destination address with the list of destination networks in its routing table. If it finds a match, the packet is forwarded as specified by the information in the next hop field.
- *Next hop*—The **next hop** (or, as specified in Windows routing tables, the gateway) indicates an interface name or the address of the next router in the path to the destination. If an interface name is specified, such as Ethernet 0 or Fast Ethernet 1, the



destination network is usually connected directly to the router. In this case, the router gets the destination device's MAC address from its ARP cache or an ARP request broadcast. After the destination MAC address is retrieved, the frame is delivered. If an address is specified, indicating that the packet must be forwarded to the next router in the path, the router retrieves the MAC address of the next router and forwards the frame on its way. When a packet must be sent to a router to get to its destination, it's called a **hop**. The total number of routers a packet must travel through is called the **hop count**.

- *Metric*—The **metric** is a numeric value that tells the router how “far away” the destination network is. Other terms for metric are cost and distance. The metric doesn't have anything to do with actual distance measured in feet or miles. It can be composed of a number of values, including the bandwidth of links between the source and destination, the hop count, the link's reliability, and so forth. If the destination network is connected directly, the metric is usually 0. The values used to determine the metric depend on how routes get into the table, discussed next.
- *How the route is derived*—This field tells you how the route gets into the routing table. A route is added to the routing table in three main ways: The destination network is connected directly; an administrator enters the route information manually (called a **static route**); or the route information is entered dynamically, via a routing protocol. The first two are somewhat self-explanatory, and the third is discussed later in “Routing Protocols.”
- *Timestamp*—Just as switching table entries need timestamps, so do routing table entries, but only those that are created dynamically. A timestamp tells the router how long it has been since the routing protocol updated the dynamic route. Not all routing protocols require a timestamp.

Figure 7-9 shows a network of several Cisco routers running a routing protocol. The table shown is the actual routing table of RouterB. The top part of the table lists codes that indicate how a route is derived. Take a look at the first row of the routing table, starting from the leftmost column:

- R—The route derived from the RIP routing protocol
- 192.168.8.0/24—The destination network
- [120/2]—The 120 indicates an administrative value for RIP; the 2 indicates a metric of 2, meaning the route can be reached by traversing through two routers (in this case, RouterC and RouterD)
- via 192.168.5.2—The next hop address; the address of the next router (RouterC) in the path
- 00:00:24—The age of the route; because RIP updates every 30 seconds, this value is usually less than that unless there's a problem
- Serial0—The interface the packet should be sent out to reach the next hop

In the network configuration shown in Figure 7-9, there's only one path from any network to any other network. To provide fault tolerance and load balancing, network administrators often design networks with multiple paths to critical networks. If the network is designed purely for fault tolerance, a router keeps only the shortest path to a destination network in its routing table. However, if the shortest path route goes down, the router can use the next

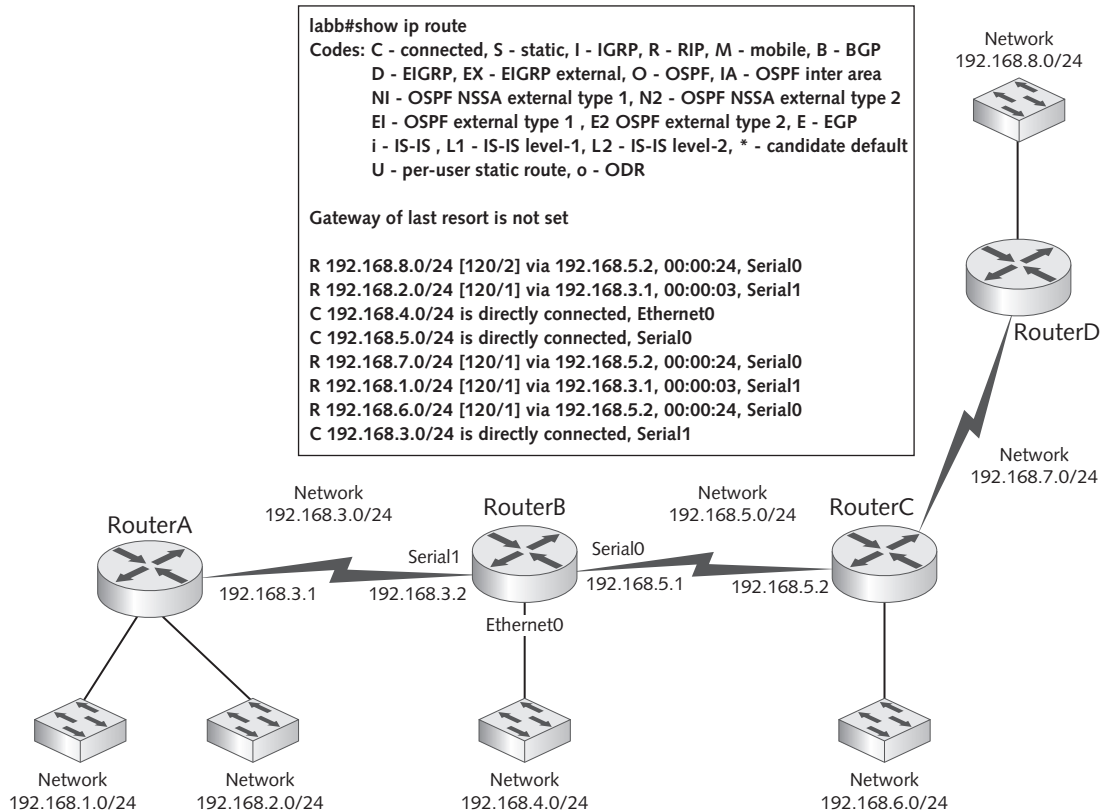


Figure 7-9 An internetwork and a router's routing table

Courtesy of Course Technology/Cengage Learning

shortest path as a backup. In other cases, a router can use two or more paths to a network to spread the traffic load over multiple network links. Simulation 17 shows what happens when a link goes down in an internetwork and the router automatically chooses an alternate path to get to the destination network.



Simulation 17: Routers use multiple paths in an internetwork

Routing Protocols

As mentioned, routing tables can be populated in three ways: directly connected networks, manually added static routes, and dynamically added routes via routing protocols. A **routing protocol** is a set of rules that routers use to exchange information so that all routers have accurate information about an internetwork to populate their routing tables. When an internetwork changes because of new networks coming online, network addresses changing, and networks going offline, routers affected by the changes pass the information on to other

routers so that all routers have an up-to-date picture of the entire internetwork. By having accurate status information on the internetwork, routers can choose the best path for routing packets they receive.

There are two main types of routing protocols; the type of protocol, in part, determines the algorithm a router uses to choose the best path to a destination when more than one exists:

- Distance-vector protocols** share information about an internetwork's status by copying a router's routing table to other routers with which they share a network. Routers sharing a network are called **neighbors**. In a large internetwork, changes to the network are passed from one router to another until all routers have received them. Distance-vector protocols use metrics based on factors such as hop count, bandwidth of the links between networks, network congestion, and delays. The best path is determined by identifying the route with the lowest metric. **Routing Information Protocol (RIP)** and **Routing Information Protocol version 2 (RIPv2)** are the most well-known distance-vector routing protocols. RIP and RIPv2 consider only hop count in path selection, which works fine if the speeds of all internetwork links are equivalent. However, if multiple paths exist, a path that uses slower links but has a lower hop count is selected over a path with faster links but a higher hop count. Figure 7-10 shows this setup. Two paths from RouterB to network 192.168.9.0 are available. One path goes directly to RouterD through network 192.168.7.0, and the other goes

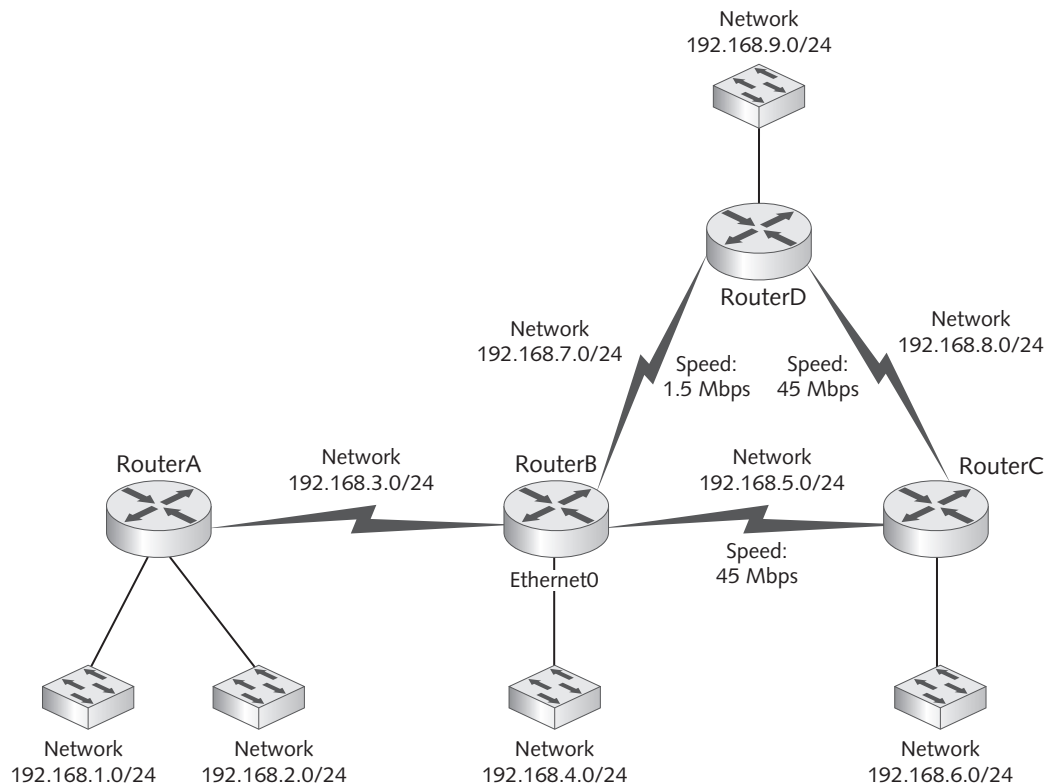


Figure 7-10 When using RIP, a router always chooses the path with the lower hop count

Courtesy of Course Technology/Cengage Learning

through RouterC via networks 192.168.5.0 and 192.168.8.0. However, each link between RouterB and RouterC and RouterC and RouterD is 45 Mbps, whereas the link between RouterB and RouterD is only 1.5 Mbps. Clearly, the path with the faster links is the better choice, but RIP and RIPv2 nonetheless choose the path from RouterB to RouterD based on a hop count of 1 versus 2.

- **Link-state protocols** share information with other routers by sending the status of all their interface links to other routers in the internetwork. The status includes link speed, whether the link is up or down, and the link's network number. This exchange of information takes place only when a change occurs in the network. When a router receives information from other routers, an algorithm is run on the information gathered to determine the best path to all networks in the internetwork. This algorithm requires more processing power than a distance-vector protocol does, but because the metric is based mainly on link speed, better choices are made in a complex internetwork. Open Shortest Path First (OSPF) is the most common link-state routing protocol. Referring again to Figure 7-10, if OSPF were the routing protocol, RouterB would choose the route through RouterC to get to network 192.168.9.0.

Table 7-2 summarizes the differences between distance-vector and link-state routing protocols.

Table 7-2 Distance-vector versus link-state routing protocols

Protocol type	CPU use	Network use	Memory use	Speed of convergence	Size of network	When is routing data transferred?
Distance-vector	Lower	Higher	Lower	Slower	Small	Periodically
Link-state	Higher	Lower	Higher	Faster	Large	Only when a change occurs

A note of explanation for one column in Table 7-2: The speed of **convergence** refers to how fast the routing tables of all routers in an internetwork are updated with accurate information when a change in the network occurs. Because distance-vector protocols pass routing tables from one router to the next, based on a periodic timer, convergence can take quite a bit of time. However, when a router is running a link-state protocol and a network change occurs, information about the change is sent immediately to all routers in the internetwork.



A third type of routing protocol called a “hybrid protocol” combines some features of distance-vector and link-state protocols. The most common example of a hybrid protocol is Cisco’s proprietary Enhanced Interior Gateway Routing Protocol (EIGRP).

Routing Protocol or Static Routes? You might wonder whether your network should run a routing protocol at all. Routing protocols aren’t a necessity in all situations, as static routes can be entered in a routing table manually. Some factors to consider when deciding whether to use a routing protocol or static routes include the following:

- Does the network change often, with new networks being added or networks going offline frequently? If so, a routing protocol is probably a good choice. If not, static routes should suffice.
- Are there several alternate paths to many of the networks in the internetwork? If so, a routing protocol can reroute around down links or congested routes automatically,

but with static routes, an administrator must change the routing table manually after the problem is discovered.

- Is the internetwork large, with many networks and remote sites? If so, keeping up with the status of all networks might prove to be more work than you bargained for if you're using static routes. A routing protocol builds and maintains routing tables automatically, freeing you to do more important tasks.

Keep in mind that there's no reason you can't combine static routing with routing protocols in your internetwork. Some areas of the internetwork might be simple and straightforward and, therefore, be suitable for static routes, and the more complex areas might benefit from using routing protocols.

Access Control Lists

One advantage of using routers in an internetwork is that you can group users who need access to common resources into subnets. A router facilitates users in one subnet being able to access resources on another subnet, but it can also be used to block access. Routers use access control lists to determine which network traffic passes through and which traffic doesn't. An **access control list (ACL)** is a set of rules configured on a router's interface for specifying which addresses and which protocols can pass through the interface and to which destinations. When an access control list blocks a packet, it's called **packet filtering**. ACLs are usually configured to filter traffic based on the following:

- Source address
- Destination address
- Protocol

The source and destination addresses can be specific IP addresses or network numbers. Filtering can be done on just the source address, just the destination address, or both. The ACL can specify anything from the entire IP protocol suite to a particular Transport-layer protocol to a specific TCP or UDP port.

ACLs as defined here can usually be configured only on mid-range to high-end routers. These routers are shipped without any configuration, and an administrator must configure all aspects of the router, including its interfaces, routing protocols, and ACLs, before the router is operational.

SOHO routers, used mainly to give a small group of computers Internet access, are usually configured already. The default access control configuration on these routers typically allows all network traffic from the private network to pass to the Internet and allows only network traffic from the Internet that was initiated from the private network to pass to the private network. An example of traffic allowed from the Internet is a Web server's response to a Web browser's request for a page. An example of traffic that's blocked is any packet coming from the Internet to the private network that the private network didn't request.



Hands-On Project 7-2: Viewing and Changing Your Computer's Routing Table

Time Required: 5 minutes

Objective: View your computer's routing table with the `route` command.

Required Tools/Equipment: Your classroom computer

Description: In this project, you use the `route` command to view and change your computer's internal routing table. Even though your computer isn't a router, it maintains an internal routing table with entries for the network interface network, the loopback network, and a variety of other internal networks.

1. Start your computer and log on as **NetAdmin**, if necessary.
2. Open a command prompt window as an administrator.
3. To view your routing table, type `route print` and press **Enter**. The `| more` after the `route print` command causes output to be displayed one screen at a time. Your output should look similar to Figure 7-11.

```

=====
Interface List
11...00 0c 29 f5 3e 97 .....Intel(R) PRO/1000 MT Network Connection
1 .....Software Loopback Interface 1
13...00 00 00 00 00 00 e0 Microsoft ISATAP Adapter
12...00 00 00 00 00 00 e0 Teredo Tunneling Pseudo-Interface
=====

IPv4 Route Table
=====
Active Routes:
Network Destination        Netmask          Gateway          Interface        Metric
0.0.0.0                    0.0.0.0          192.168.100.1    192.168.100.100  266
127.0.0.0                  255.0.0.0        On-link          127.0.0.1        306
127.0.0.1                  255.255.255.255 On-link          127.0.0.1        306
127.255.255.255            255.255.255.255 On-link          127.0.0.1        306
192.168.100.0              255.255.255.0    On-link          192.168.100.100  266
192.168.100.100            255.255.255.255 On-link          192.168.100.100  266
192.168.100.255            255.255.255.255 On-link          192.168.100.100  266
224.0.0.0                  248.0.0.0        On-link          127.0.0.1        306
224.0.0.8                  248.0.0.8        On-link          192.168.100.100  266
255.255.255.255            255.255.255.255 On-link          127.0.0.1        306
255.255.255.255            255.255.255.255 On-link          192.168.100.100  266
=====
-- More --

```

Figure 7-11 Output from the `route print` command

Courtesy of Course Technology/Cengage Learning

4. Next, examine the output of the `route print` command. Your computer's network interfaces are listed at the top, and the IPv4 Route Table lists entries in the routing table, which has five columns:
 - *Network Destination*—The network destination your computer compares with the destination IP address of outgoing packets to determine where to send them.
 - *Netmask*—The subnet mask of the network destination. A value of 255.255.255.255 indicates that the address in the Network Destination column is a specific IP address rather than a network address; it's referred to as a "host route." A value of 0.0.0.0 is used when the network destination is 0.0.0.0, indicating the default route or gateway.

- *Gateway*—The next hop address or the on-link, which means the network destination is connected directly to an interface. Make a note of the value in this column for the 0.0.0.0 network destination, as you need it later:

-
- *Interface*—The address of the interface Windows uses to send the packet to the network destination.
 - *Metric*—The metric assigned to the route. If there are two entries for the same network destination, the lower metric is the route chosen.

Press the **spacebar** one or more times to display the rest of the output. You'll see a row of output labeled Persistent Routes. If you create a route manually and want it to stay in the table between reboots, it's listed here. You'll also see your default route listed under Persistent Routes in the IPv4 section of the output.

5. To verify that you can communicate with the Internet, type **ping www.course.com** and press **Enter**. If the ping is successful, your default network is working correctly.
6. Type **route delete 0.0.0.0** and press **Enter** to delete your default route. Try to ping **www.course.com** again. The ping will fail. Type **route print | more** and press **Enter** to verify that the 0.0.0.0 network destination is no longer in the table. Press the **spacebar** one or more times to display the rest of the output.
7. To create the default route entry, type **route add -p 0.0.0.0 mask 0.0.0.0 default-gateway** and press **Enter** (replacing *default-gateway* with the address you noted in Step 4).
8. Display the routing table again to verify that your default route is in the table. Ping **www.course.com** to verify that you can do so again.
9. Close all open windows, and stay logged on if you're going on to the next project.

Wireless Access Points in Depth

You learned about the basic operation of wireless access points (APs) in Chapter 2 and Wi-Fi technology in Chapter 3. This section discusses some configuration options available on most wireless APs and wireless routers, regardless of the 802.11 networking standard in use. First you look at wireless network configuration options, and in the next section, you explore options available on wireless routers. Remember that a wireless router, as it's usually marketed, is actually three devices in one: a wireless AP, a router, and a switch. So when you examine the features of an AP, the discussion also applies to the AP built into a wireless router. The following sections explain these AP options and settings: basic wireless settings, wireless security options, and advanced wireless settings.

Basic Wireless Settings

Basic wireless settings on most APs define the settings a client wireless device needs to connect to an AP:

- Wireless network mode
- Wireless network name (SSID)
- Wireless channel
- SSID broadcast status

Figure 7-12 is an example of where these settings can be configured.



Figure 7-12 Basic wireless settings

Courtesy of Course Technology/Cengage Learning

The wireless network mode allows you to choose which 802.11 standard the AP should operate under. If an AP supports multiple modes, it can generally operate in all or some combination of the supported modes. For example, if the AP supports 802.11n, 802.11g, and 802.11b, the options are as follows:

- *Mixed*—In mixed mode, the AP supports client connections by using any of the three standards.
- *BG-mixed*—Both 802.11b and 802.11g client connections are supported.
- *N only*—Only 802.11n client connections are possible.
- *G only*—Only 802.11g client connections are possible.
- *B only*—Only 802.11b client connections are possible.
- *Disabled*—On wireless routers, you can disable the AP portion if you have only wired devices. If you don't have wireless devices, you should use this setting.

APs generally work best when only one standard is enabled. So if your wireless network has only clients with 802.11n NICs, choose N only mode. Mixed mode is useful when you're

transitioning clients from an older standard, such as 802.11b, to a newer and faster standard, such as 802.11n.

The wireless network name refers to the service set identifier (SSID) discussed in Chapter 2. When an AP is shipped, the SSID is set to a default value, such as “default” or “linksys” for a Linksys AP. You can leave the name as is, but doing so is a sure giveaway to wardrivers that your network is probably not secure. An SSID can be up to 32 characters. If you’re running a wireless network in your home in an area with other wireless networks operating, it’s best not to set the SSID to a name that obviously identifies you. The more hackers know about a network and its users, the more easily they can infiltrate it.



A wardriver is someone who drives around neighborhoods with a wireless scanning device looking for unsecured wireless networks. After finding one, a wardriver might simply use the network for free Internet access—or worse, break into the computers connected to the network.

Wireless channels were discussed in Chapter 3. An AP in the United States can usually be set to operate in channels 1 through 11. Recall that the channels overlap, so if you set your AP to channel 3, data is being transferred over channels 1 to 5. For optimal operation when multiple APs are in use, choose channels that are five channels apart, such as 1, 6, and 11.

The option to enable or disable the SSID broadcast is a low-end security option. By default, APs are configured to transmit the SSID so that any wireless device in range can see the network. If the SSID isn’t broadcast, wireless devices can still connect to the wireless network by entering the SSID manually, if it’s known. Disabling SSID broadcasts thwarts only the least sophisticated wardrivers because wireless packet-capturing programs can still intercept and view SSIDs transmitted between clients and the AP.

Wireless Security Options

Enabling wireless security options on your AP is a critical step of setting up a wireless LAN. Most APs offer the following security options:

- Encryption
- Authentication
- MAC filtering
- AP isolation

All private wireless networks should use encryption at a minimum to secure communication. With encryption enabled, even if a hacker captures transmitted data, it’s unintelligible. Encryption protocols vary in strengths, and the most common protocols are listed from weaker to stronger:

- Wired Equivalent Privacy (WEP)
- Wi-Fi Protected Access (WPA)
- Wi-Fi Protected Access 2 (WPA2)

These protocols are described in more detail in Chapter 10. What you need to know in practice is that you should use the highest level of security your systems support. Be aware that

older APs and wireless NICs might not support the stronger WPA and WPA2 protocols. Because all devices must use the same protocol, you have to set your wireless security to a protocol supported by all your wireless devices.

If authentication is used on an AP, users must enter a username and password to access the wireless network. Authentication is most likely to be used on wireless networks that allow limited public access, such as at colleges, libraries, and other organizations that allow wireless access by members who can be identified with a username or some other credential. APs that support authentication usually support the Remote Authentication Dial-In User Service (RADIUS) protocol, in which the AP contacts your network servers that store user account information to provide authentication. Often, if a wireless network is protected by authentication, it's also protected by encryption.

MAC filtering, a feature available on most APs, enables you to restrict which devices can connect to your AP. To use MAC filtering, you add the MAC addresses of the wireless devices allowed to access your network to a list on the AP. After this configuration, only computers with wireless NICs that have a MAC address in the list can connect to the wireless network. MAC filtering should be used *only* in combination with encryption because a sophisticated user could capture packets that include the MAC addresses of authorized devices. These captured MAC addresses can then be used to impersonate authorized computers.

AP isolation mode creates a separate virtual network for each client connection. If this mode is enabled, clients connecting to the AP can access the Internet but can't communicate with each other. This mode is a good option for Internet cafés and other establishments that offer wireless Internet access.

Advanced Wireless Settings

Many advanced settings are available on high-end APs to help you manage a wireless network. The following are some common settings:

- *Adjustable transmit power*—This setting lets you control the power and, therefore, the range of the wireless network signal. For example, if you place your AP near the center of your building, you can adjust the power so that clients inside the building can connect, but clients outside the building can't.
- *Multiple SSIDs*—Two or more wireless networks can be created with different security settings, such as when you want to create a private network and a guest network.
- *VLAN support*—Enable VLAN support to assign wireless networks to wired VLANs.
- *Traffic priority*—If your AP is configured for multiple SSIDs and, therefore, multiple wireless networks, you can assign a priority to packets coming from each network.
- *Wi-Fi Multimedia*—Defined by 802.11e, this standard provides quality of service (QoS) settings for multimedia traffic, giving priority to streaming audio or video, for example.
- *AP modes*—An AP can be set to operate as a traditional access point, a repeater, or a wireless bridge. In repeater mode, the AP is used to extend the range of an existing AP, making two or more APs essentially act as a single AP. A repeater configuration also enables users to roam the wireless network. When a client passes out of one AP's



range, it connects to another AP automatically. Bridge mode is used to connect physically separate wired networks by using two or more APs. For example, wired networks in separate buildings can be connected by using APs in bridge mode so that wired clients in each building can communicate with one another.

Network Interface Cards in Depth

A NIC makes the connection between a computer and the network medium, so the performance and reliability of a computer's NIC are crucial to the computer's network performance. Chapter 2 introduced a NIC's operation and basic factors for choosing a NIC. This section covers the PC bus options and advanced features to look for in a NIC when purchasing one for a workstation or server.

PC Bus Options

A bus makes the connections between a computer's vital components, such as the CPU, RAM, and I/O devices. The faster the bus, the faster data can be transferred between these components, which makes for a faster overall system. NICs are considered I/O devices, and whether they're built into the motherboard or added as an expansion card, they still communicate with the rest of the components via the bus. PC bus options have changed over the years with older, and usually slower, technologies fading into obsolescence and being replaced by newer, faster technologies. The following list describes the most common PC bus architectures in current use:

- *PCI*—Several local bus standards appeared in the early 1990s as computers became faster, but by 1995, Intel's **Peripheral Component Interconnect (PCI)** bus became the default bus standard. PCI 2.x moved PCI from being a local bus tied to the CPU to a microprocessor-independent bus that can be used with any CPU. Most PCI implementations are 32-bit and operate at 33 MHz with a maximum data transfer rate of 133 MBps. The 64-bit PCI, running at 66 MHz with a maximum data transfer rate of about 533 MBps, wasn't widely implemented. PCI supports bus mastering and was the first bus to accommodate the Microsoft Plug-and-Play architecture. Finally, PCI supports interrupt sharing on a PC, so any PCI adapters can share a single interrupt request (IRQ) line without requiring a unique IRQ for each adapter. (Therefore, only a single free IRQ is required for all PCI cards on a PC.) PCs are still built with PCI (or PCI-X) slots but this bus, at more than 15 years old, is coming to the end of its lifetime.
- *PCI-X*—**Peripheral Component Interconnect-Extended (PCI-X)** is backward-compatible with PCI but supports speeds from 66 MHz to 533 MHz, providing data transfers from 500 MBps to more than 4 GBps. The original PCI-X 1.0 specification topped out at 133 MHz with a transfer rate of 1066 MBps. Like PCI, PCI-X supports 32-bit or 64-bit bus widths. These speeds reduced the system bus as a performance bottleneck and accommodated developments in faster network cards (such as Gigabit and 10 Gigabit Ethernet) and disk controllers (such as Ultra3 SCSI and 10Gb Fibre Channel). As of this writing, both PCI and PCI-X NICs support Gigabit Ethernet running on fiber-optic, Cat 5e, or Cat 6 cabling. Figure 7-13 shows a PCI/PCI-X bus NIC.




Figure 7-13 A PCI/PCI-X NIC

Courtesy of Course Technology/Cengage Learning

- **PCIe—PCI Express (PCIe)**, formerly named 3GIO, uses a high-speed serial communication protocol of one or more lines or lanes. Each lane of PCIe 1.0 can operate at 250 MBps in each direction, and PCIe 2.0 provides data rates of 500 MBps. Because PCIe can be set up in lanes, several lanes can be combined, resulting in tremendous transfer speeds up to 16 GBps. Although PCI-X bus slots are backward-compatible with existing PCI boards, PCIe hardware maintains backward-compatibility with PCI only in board design; the expansion slot required for PCIe is vastly different from PCI's. PCIe boards are specified with notations such as x1, x4, x8, x16, and x32. The number following the “x” is the number of lanes. The more lanes, the higher the bandwidth. For example, a PCIe 1.0 x1 board supports data transfer rates up to 250 MBps, and a PCIe 2.0 x32 board supports transfer rates up to 16 GBps. The PCIe 3.0 specification, released in November 2010, doubles the bandwidth of PCIe 2.0. PCIe is now the dominant bus technology, particularly for high-bandwidth devices such as NICs, disk controllers, and video cards. Figure 7-14 shows a PCIe x1 NIC.
- **PCMCIA—PCMCIA cards** are credit card-sized expansion cards used mainly to add functionality to laptop computers. Two primary standards are in common use: Cardbus and ExpressCard. Cardbus is the more mature standard, having been around since the mid-1990s. It operates at 33 MHz and supports a 32-bit bus, providing up to 132 MBps data transfer rates. ExpressCard was developed as computer users' thirst for faster data transfer speeds continued to grow. It uses PCIe technology to provide data transfer rates up to 500 MBps, with future versions expected to reach 4 GBps. A variety of NICs are available in these formats, including wireless NICs. Figure 7-15 shows a Cardbus NIC with an RJ-45 connector.






Figure 7-14 A PCIe x1 NIC

Courtesy of Course Technology/Cengage Learning



Figure 7-15 A Cardbus NIC

Courtesy of Course Technology/Cengage Learning

- **USB—Universal Serial Bus (USB)** comes in three versions: USB 1.0, USB 2.0, and USB 3.0. USB 1.0 is a low-speed serial interface operating at a maximum bandwidth of 12 Mbps. It's used mostly for low-speed peripheral devices, such as mice, keyboards, and joysticks, but can also be used to connect printers, scanners, phones, and some video devices to a computer. USB is now a standard interface on both PCs and Macintoshes for these uses. In networking, USB is usually used as an interface for wireless network adapters or as an attachment for cable or DSL modems. Because USB is an external interface on PCs, it offers the advantage of being able to add a NIC to a system without opening the computer case or even shutting down the computer. USB 2.0 can operate at up to 480 Mbps (60 MBps)—40 times faster than USB 1.0. This interface is used for external hard

drives, CD/DVD burners, flash memory card readers, high-speed scanners, wired and wireless network adapters, and digital cameras. USB 3.0 (or SuperSpeed USB) is a new USB specification, with products first coming on the market in 2010, and promises speeds up to 3.2 Gbps or 400 MBps.

Table 7-3 is a quick reference of several bus types used for networking.

Table 7-3 Common bus types

Bus type	Maximum transfer rate	Bus size	Use
PCI	533 MBps	32 and 64 bits	Declining; absent on most servers
PCI-X	4 GBps	32 and 64 bits	Declining use on servers and high-end workstations
PCI Express (PCIe)	16 GBps	1 to 32 serial lanes	Standard on all types of computers
PCMCIA	500 MBps	1 serial lane	Laptops
USB 2.0	60 MBps	N/A	Laptops and easy addition of NICs to desktops

Advanced Features of NICs

Because NICs are the focus of network traffic on workstations and large volumes of traffic on network servers (even those with more than one network interface), they can have a major influence on network performance. If a NIC is slow, it can limit network performance. Particularly on networks with shared media, slow NICs anywhere on the network can decrease performance for all users.

When selecting a network adapter, first identify the physical characteristics the card must match. They include the type of bus the card will connect with (PCI or PCMCIA, for example), the type of network technology in use, and the kind of connector or physical attachment the adapter must accommodate. After you determine these basic characteristics, it's equally important to consider purchasing other options that can seriously affect a card's speed and data-handling capabilities. Some of these options suit servers better, and others work equally well for servers and clients; all help improve overall network performance. These hardware-enhancement options include the following:

- **Shared adapter memory** means the adapter's buffers map directly to RAM on the computer. A computer actually writes to buffers on the NIC instead of writing to its own memory. In this instance, the computer treats adapter RAM as its own RAM.
- **Shared system memory** means a NIC's onboard processor selects a region of RAM on the computer and writes to it as though it were buffer space on the adapter. In this instance, the adapter treats computer RAM as its own RAM.
- **Bus mastering** permits a network adapter to take control of the computer's bus to initiate and manage data transfers to and from the computer's memory, independent of the CPU. This feature lets the CPU concentrate on other tasks and can improve network performance.
- **RAM buffering** means a NIC includes additional memory to provide temporary storage for incoming and outgoing data that arrives at the NIC faster than it can be shipped out. This option speeds overall performance because it lets the NIC process data as quickly as it can, without having to pause occasionally to grab (or send) more data.



- **Onboard co-processors** included on most NICs enable the card to process incoming and outgoing network data without requiring service from the CPU.
- Security features available on some high-end NICs allow them to handle several protocol functions, including IP Security (IPSec) and other encryption services related to authentication and payload protection. IPSec is a secure transport mechanism that protects network traffic from unwanted snooping.
- Quality of service (QoS) allows prioritizing time-sensitive data, such as streaming video and voice.
- **Automatic link aggregation** enables you to install multiple NICs in one computer and aggregate the bandwidth so that, for example, you can install two 1 Gbps NICs and have a total bandwidth of 2 Gbps to and from that computer. This feature is found most commonly on NICs designed for servers.
- Improved **fault tolerance**, in the form of redundant NICs with failover capabilities, is available on some high-end NICs. By installing a second NIC in a PC, failure of the primary NIC shifts network traffic to the second NIC instead of cutting off the PC from the network. Hot-pluggable NICs are also an option for fault tolerance because a NIC can be installed or removed without turning off the server. NICs with dual ports provide added bandwidth and fault tolerance. These NICs have two media connectors, both of which can be active, which doubles bandwidth and fault tolerance in case one media connection fails.
- Advanced Configuration Power Management Interface (ACPI)–compliant cards provide features such as wake-on-LAN, which allows an administrator to power on a PC remotely by accessing the NIC through the network. In addition, Simple Network Management Protocol (SNMP) is built in on some NICs to allow remote configuration and management.
- Preboot Execution Environment (PXE)–compliant adapters allow a computer to download an OS instead of booting it from a local hard drive. This feature is used on diskless workstations (“thin clients”) that don’t store an OS locally. PXE-compliant adapters are also useful when network administrators use disk-imaging software to install the OS and applications on a number of computers on the network. A computer configured to boot from its PXE-enabled NIC announces to the network that it’s trying to boot, and a server responds by sending a disk image to the PC. After the image is written to the hard drive, the PC boots from the local OS.

For a typical desktop computer, a standard PCI or PCIe NIC with basic features is usually adequate. Servers, however, warrant some of the high-end features discussed in the preceding list. Servers get network requests from multiple clients, so they benefit from having buffer space to hold incoming frames temporarily while other frames are processed. Virtualized environments benefit from NICs with multiple ports so that virtual machines don’t have to share bandwidth with their host computer. You’ll find when you investigate purchasing a NIC that prices can vary wildly from a few dollars to several hundred dollars. A \$3.99 NIC on clearance might work fine in your desktop computer but definitely shouldn’t be installed in mission-critical servers.

Chapter Summary

- Network switches use auto-negotiate mode to determine the link speed and duplex mode. Auto-MDIX allows a switch to swap transmit and receive pins negating the need for crossover cables.
- Switching tables can hold many more MAC addresses than ports. Several MAC addresses can be mapped to a single port. The aging time prevents stale switching table entries.
- Switches forward frames by using a variety of methods. In order of fastest to slowest and least reliable to most reliable, they're cut-through, fragment-free, and store-and-forward. Some switches can use a combination of these methods; fragment-free is used until a number of errors occur, and then store-and-forward is used until the errors decrease.
- Advanced features, such as VLANs, STP, multicast support, and port security, are found on smart or managed switches. VLANs create multiple broadcast domains but require a router to communicate between them. STP can prevent broadcast storms in a redundant configuration. Multicast support prevents the switch from simply flooding multicast frames. Port security prevents unauthorized MAC addresses from connecting to the switch.
- Routing tables contain destination networks, next hop addresses, metrics, methods used to derive routes, and timestamps. They can be populated statically or with a routing protocol.
- Routing protocols populate routing tables dynamically. The most common types of routing protocols are distance-vector and link-state. RIP, an example of a distance-vector protocol, is best used in smaller networks without redundant links. OSPF, an example of a link-state protocol, is suitable for large complex networks.
- Access points have the following basic settings: wireless mode, SSID, and wireless channel. APs support security protocols that encrypt wireless data. From least secure to most secure, these protocols are WEP, WPA, and WPA2. Other security features include authentication, MAC filtering, and AP isolation.
- Higher-end APs can support advanced features, such as multiple SSIDs, adjustable transmit power, VLANs, QoS, and repeater and bridge modes.
- NIC selection includes the PC bus. PCI, PCI-X, and PCIe are the most common for internal NICs; USB can be used for external NICs. PCIe provides the fastest bus interface and is becoming the standard PC bus.
- Some advanced NIC features to consider include RAM buffering, onboard co-processors, automatic link aggregation, multiple ports for fault tolerance, ACPI, and PXE.

Key Terms

access control list (ACL) A set of rules configured on a router's interface for specifying which addresses and protocols can pass through the interface and to which destinations.



aging time The amount of time a switch maintains a switching table entry that hasn't been updated.

automatic link aggregation A feature that enables you to install multiple NICs in one computer and aggregate the bandwidth so that, for example, you can install two 1 Gbps NICs and have a total bandwidth of 2 Gbps to and from that computer.

auto-MDIX A switch port option used to detect the type of device and cable the switch port is connected to; if necessary, the port swaps its transmit and receive pins, which enables you to use a straight-through or crossover cable regardless of the type of device you're connecting to the port.

auto-negotiate mode Communication between a switch and a device connected to a switch port, in which the switch attempts to set the port's operating mode to the highest performance setting the device supports.

blocking mode A mode on a switch port that prevents the switch from forwarding frames out the blocked port, thereby preventing a switching loop. *See also* switching loop.

broadcast storm A condition that occurs when a broadcast frame is forwarded endlessly in a switching loop. *See also* switching loop.

bus mastering A feature that allows a network adapter to take control of the computer's bus to initiate and manage data transfers to and from the computer's memory, independent of the CPU.

cut-through switching With this switching method, the switch reads only enough of the incoming frame to determine its source and destination addresses. After the forwarding location is determined, the frame is switched internally from the incoming port to the outgoing port, and the switch is free to handle additional frames.

destination network The network address of a network to which the router can forward packets.

distance-vector protocol A routing protocol that routers use to share information about an internetwork's status by copying their routing table to other routers with which they share a network.

fault tolerance A feature available on some high-end NICs. By installing a second NIC in a PC, failure of the primary NIC shifts network traffic to the second NIC instead of cutting off the PC from the network.

flood The process whereby a switch forwards a frame out all connected ports.

fragment-free switching With this switching method, the switch reads enough of the frame to guarantee that it's at least the minimum size for the network type, reducing the possibility that the switch will forward a frame fragment.

frame fragment An invalid frame that's damaged because of a collision or a malfunctioning device.

hop Each router a packet must go through to get to the destination network.

hop count The total number of routers a packet must travel through to get to its destination network.

link-state protocol A routing protocol that a router uses to share information with other routers by sending the status of all its interface links to all other routers in the internetwork. The status includes link speed, whether the link is up or down, and the link's network number.

managed switch A high-end switch with many advanced features that can be configured.

metric A numeric value that tells the router how “far away” the destination network is. It can be composed of values such as the bandwidth of links between the source and destination, the hop count, and the link’s reliability.

neighbor In an internetwork, routers sharing a common network.

next hop An interface name or the address of the next router in the path to the destination network.

onboard co-processors A feature included on most NICs that enables the card to process incoming and outgoing network data without requiring service from the CPU.

packet filtering A process whereby a router blocks a packet from being forwarded based on rules specified by an access control list. *See also* access control list (ACL).

packet forwarding The process of a router receiving a packet on one port and forwarding it out another port based on the packet’s destination network address and information in the routing table.

PCI Express (PCIe) A bus standard that uses a high-speed serial communication protocol of one or more lines or lanes. Each lane of PCIe 1.0 can operate at 250 MBps in each direction. *See also* Peripheral Component Interconnect (PCI).

PCMCIA cards Credit card–sized expansion cards used mainly to add functionality to laptop computers. The main standards are Cardbus and ExpressCard. Cardbus operates at 33 MHz and supports a 32-bit bus; ExpressCard uses PCIe technology to provide data transfer speeds up to 500 MBps.

Peripheral Component Interconnect (PCI) A bus standard used to connect I/O devices to the memory and CPU of a PC motherboard. PCI is implemented in both 32-bit and 64-bit versions at speeds of 33 and 66 MHz, respectively.

Peripheral Component Interconnect-Extended (PCI-X) A bus standard that’s backward-compatible with PCI and supports speeds of 66 to 533 MHz with 32-bit or 64-bit bus widths. *See also* Peripheral Component Interconnect (PCI).

RAM buffering A NIC feature for including additional memory to provide temporary storage for incoming and outgoing data.

Routing Information Protocol (RIP) A distance-vector protocol that uses hop count as the metric to determine the best path to a destination network.

Routing Information Protocol version 2 (RIPv2) A newer version of RIP that supports a more complex IP addressing scheme and uses multicast packets rather than broadcasts to transmit routing table updates. *See also* Routing Information Protocol (RIP).

routing protocol A set of rules routers use to exchange information so that all routers have accurate information about an internetwork to populate their routing tables.

shared adapter memory A feature on some NICs in which the NIC’s buffers map directly to RAM on the computer. A computer actually writes to buffers on the NIC instead of writing to its own memory.

shared system memory A feature on some NICs in which a NIC’s onboard processor selects a region of RAM on the computer and writes to it as though it were buffer space on the adapter.

smart switch A mid-range switch with some advanced features, typically multicast processing, Spanning Tree Protocol, VLANs, and port security. *See also* Spanning Tree Protocol (STP) and virtual local area networks (VLANs).



Spanning Tree Protocol (STP) A communication protocol switches use to ensure that they aren't connected in a way that creates a switching loop. *See also* switching loop.

static route A routing table entry that's entered manually by an administrator.

store-and-forward switching This switching method requires the switch to read the entire frame into its buffers before forwarding it. It examines the frame check sequence (FCS) field to be sure the frame contains no errors before it's forwarded.

switching loop A condition that occurs when switches are connected in such a way that frames can be forwarded endlessly from switch to switch in an infinite loop.

trunk port A switch port configured to carry traffic from all VLANs to another switch or router. *See also* virtual local area networks (VLANs).

Universal Serial Bus (USB) An external PC bus interface for connecting I/O devices. Speeds range from 12 Mbps in USB 1.0 to 3.2 Gbps in USB 3.0.

virtual local area networks (VLANs) A feature on some switches that allows configuring one or more switch ports into separate broadcast domains.

Review Questions

1. When a switch receives a frame on a port and floods the frame, what does it do with the frame?
 - a. Discards it
 - b. Changes the destination address to FF:FF:FF:FF:FF:FF
 - c. Forwards it out all other connected ports
 - d. Clears the switching table and adds the frame source address to the table
2. You have two eight-port switches. On each switch, seven stations are connected to ports, and the two switches are connected with the eighth port. How many collision domains are there?
 - a. 16
 - b. 15
 - c. 14
 - d. 8
 - e. 1
3. Which of the following is considered a Layer 2 device?
 - a. Computer
 - b. Switch
 - c. Router
 - d. Hub
4. You just purchased some new switches for your company's network. Your junior technicians are doing most of the work connecting switches to workstations and to each other, and you don't want to confuse them by requiring them to use both patch cables and crossover cables. How can you test the switches to determine whether you need

- both types of cable, and what's the feature for using only one type of cable for all connections?
- Connect the switch to a PC NIC and configure different speeds on the NIC by using the NIC driver. You're okay if the switch links at all speeds. It's called auto-MDIX.
 - Connect two switches by using a crossover cable. If the connection works, the switch supports auto-negotiate.
 - Connect the switch to a PC NIC and configure different speeds on the NIC by using the NIC driver. You're okay if the switch links at all speeds. It's called auto-negotiate.
 - Connect two switches by using a patch cable. If the connection works, the switch supports auto-MDIX.
5. What feature of a switch keeps switching table entries from becoming stale?
6. Which is the fastest switching method?
- Store-and-forward
 - Fragment-free
 - Cut-through
 - Forward-free
7. There can be only one MAC address per port in a switching table. True or False?
8. What does it mean if the first 24 bits of a MAC address are 01:00:5E?
- The NIC was manufactured by Intel.
 - It's a multicast frame.
 - It's an invalid CRC.
 - The frame will be flooded.
9. What feature should you look for in switches if your network is cabled like the one in Figure 7-16?




Figure 7-16 Network diagram for review question 9

Courtesy of Course Technology/Cengage Learning



- a. VLANs
 - b. Auto-negotiate
 - c. STP
 - d. Auto-MDIX
10. What should you configure on a switch that's connected to three broadcast domains?
- a. IGMP
 - b. VLANs
 - c. Port security
 - d. STP
11. Which of the following is a Layer 3 device?
- a. Router
 - b. NIC
 - c. Switch
 - d. Computer
12. What does a router do after receiving a frame on one of its interfaces? (Choose all that apply.)
- a. Deencapsulates the frame to create a packet
 - b. Deencapsulates the packet to create a segment
 - c. Encapsulates the frame to create a new packet
 - d. Encapsulates the packet to create a new frame
13. Which of the following is found in a routing table? (Choose all that apply.)
- a. Destination MAC address
 - b. Port number
 - c. Metric
 - d. Next hop
 - e. Domain name
14. Which of the following accurately describes a distance-vector routing protocol?
- a. OSPF is an example.
 - b. It learns from its neighbors.
 - c. It sends the status of its interface links to other routers.
 - d. It converges the fastest.
15. Which of the following is a characteristic of routing protocols? (Choose all that apply.)
- a. They populate routing tables statically.
 - b. Network changes are reflected in the routing table automatically.
 - c. They're not a good solution with redundant routes.
 - d. They add routing table entries dynamically.

16. Which of the following is the best routing solution for a network that includes redundant links?
 - a. RIP
 - b. STP
 - c. OSPF
 - d. Static
17. If you don't want wireless clients to view the name of your wireless network, what feature should you use?
 - a. WEP
 - b. Disabling SSID broadcasts
 - c. MAC filtering
 - d. AP isolation
18. To prevent a wardriver from being able to interpret captured wireless network data, you should enable which of the following?
 - a. MAC filtering
 - b. AP isolation
 - c. WPA or WPA2
 - d. Repeater mode
19. What feature can you use to wirelessly connect the wired networks in two buildings?
 - a. Repeater mode
 - b. AP isolation
 - c. Bridge mode
 - d. VLAN mode
20. Which AP feature is useful when you have many guests accessing your network and you don't want them to be able to access the computers of other guests?
 - a. MAC filtering
 - b. AP isolation
 - c. Bridge mode
 - d. VLAN mode
21. Which PC bus uses up to 32 lanes to achieve very high data transfer rates?
 - a. PCI
 - b. PCI-X
 - c. USB
 - d. PCIe



22. Which PC bus allows you to connect a NIC to your computer easily without powering off?
 - a. PCI
 - b. PCI-X
 - c. USB
 - d. PCIe
23. Which NIC feature do you need to configure on a thin client?
 - a. QoS
 - b. PXE
 - c. IPSec
 - d. ACPI
24. Which device is used to communicate between broadcast domains?
 - a. Repeater
 - b. Switch with VLANs
 - c. Router
 - d. Switch with STP
25. What feature should you configure to prevent users on one subnet from accessing the Web server on another subnet?
 - a. MAC filtering
 - b. Access control lists
 - c. Dynamic routing
 - d. Spanning Tree Protocol

Challenge Labs



Challenge Lab 7-1: Configuring a Wireless Network for Security

Time Required: 40 minutes

Objective: Install a wireless NIC (if necessary), connect to an access point, and configure security options.

Required Tools/Equipment: Two or more computers with 802.11 wireless NICs installed; USB wireless NICs work well because they don't require opening the computer case, and laptops with built-in wireless NICs will also do. One wireless AP or wireless router configured with the SSID NetEss. The 802.11 standard supported doesn't matter as long as the AP is compatible with the NICs. Windows 7 is the preferred OS, but some steps can be changed to accommodate other OSs. The computers shouldn't be connected to a hub or switch.

Description: In this lab, you connect to a wireless access point and then configure security and other options. You can work in groups or have the instructor configure the AP and the students configure the NICs. If you work in groups, each group should choose a different SSID and channel. Perform the following tasks:

- Install the wireless NIC in each computer, if necessary.
- Configure the AP SSID as NetEss, if necessary.
- Connect to the AP from each computer.
- Configure the AP for the following security protocols, if available: WEP, WPA, and WPA2. After you configure the AP for a security protocol, change the wireless NIC configuration to use the same security protocol. Use as many of these protocols as your hardware supports.
- Change the AP's configuration to not use a security protocol and verify that all computers can connect before continuing.
- Configure these options on the AP to see how they affect your client connections; reset each option after you have tested it before moving on to the next option:
 - Disable SSID broadcasts.
 - MAC filtering: Configure filtering to allow only one or two computers to connect, and verify that the others can't connect.
- Answer the following questions:
 - Which security protocol was the easiest to configure? Which protocol is the most secure?

 - How did disabling the SSID broadcast affect your ability to connect to the AP?

 - How did MAC filtering affect your ability to connect to the AP?



Challenge Lab 7-2: Filling in Routing Tables

Time Required: 45 minutes

Objective: Fill in routing tables for the routers in Figure 7-17.

Required Tools/Equipment: None

Description: In this lab, you fill out routing tables for the network shown in Figure 7-17, using the following charts. The network is running the RIP routing protocol. RouterA's table has been started for you. Assume all networks are Class C networks.




Figure 7-17 Network diagram for Challenge Lab 7-2

Courtesy of Course Technology/Cengage Learning

Key for the How column: C = directly connected network; R = RIP.

Router A

How	Network	Via/next hop	Metric	Interface
C	192.168.1.0/24	Connected	0	Ethernet0
	192.168.2.0/24			
	192.168.3.0/24			
	192.168.4.0/24			
	192.168.5.0/24			
R	192.168.6.0/24	192.168.3.2	2	Serial0
	192.168.7.0/24			
R	192.168.8.0/24	192.168.3.2		
	192.168.9.0/24			

Router B

How	Network	Via/next hop	Metric	Interface
	192.168.1.0/24			
	192.168.2.0/24			
	192.168.3.0/24			
	192.168.4.0/24			
	192.168.5.0/24			
	192.168.6.0/24			
	192.168.7.0/24			
	192.168.8.0/24			
	192.168.9.0/24			

Router C

How	Network	Via/next hop	Metric	Interface
	192.168.1.0/24			
	192.168.2.0/24			
	192.168.3.0/24			
	192.168.4.0/24			
	192.168.5.0/24			
	192.168.6.0/24			
	192.168.7.0/24			
	192.168.8.0/24			
	192.168.9.0/24			

Router D

How	Network	Via/next hop	Metric	Interface
	192.168.1.0/24			
	192.168.2.0/24			
	192.168.3.0/24			
	192.168.4.0/24			
	192.168.5.0/24			
	192.168.6.0/24			
	192.168.7.0/24			
	192.168.8.0/24			
	192.168.9.0/24			



Case Projects



Case Project 7-1

You have been called in to consult on a new network design for CNT Books. The requirements of this design are summarized as follows:

- The building has three floors.
- There are 300 user workstations and 10 servers.
- Users must be grouped according to the projects they're working on, but the users for each project are located on all three floors.
- There must be fault tolerance for communicating between the floors.

What features would you look for on the switches you purchase for this design? Explain why you would want each feature. Do you need to include any other devices in this design? Write a memo to your instructor with the answers to these questions or be prepared to discuss your answers.

Case Project 7-2

The CNT Books network described in Case Project 7-1 is expanding. There are 200 more user stations in the building, and a total of five floors are in use by the network. You have kept up with the design so far with a network of five subnets, each with its own router. The company leased a building across the street. You expect at least four subnets to be added to the design. The owner is concerned about how to connect to the building across the street, as he thinks the cost of contracting with a communications provider is too expensive for such a short distance. What solution can you suggest for connecting the building across the street with the existing building?

Case Project 7-3

Three years later, you're still consulting for CNT Books. The network has more than 15 subnets and 10 routers in several buildings and locations. You have been keeping up with the network by configuring the routers statically. However, users have had problems with downtime in the past year because of network links going offline, as there's only one route to reach every subnet. The owner wants fault tolerance built into the network to include backup links in case a primary link goes offline. You're concerned that the current router configuration method will still cause some downtime, even if the backup links operate correctly. Why might there be downtime if a primary link goes offline but the backup link is okay? What can you do to reduce the possibility of downtime?

Network Operating System Fundamentals

After reading this chapter and completing the exercises, you will be able to:

- Explain the major components of an OS, including the file system, processes, and the kernel
- Describe network operating systems and compare client and server OSs
- Describe the components of virtualization and virtualization products
- Plan for an OS installation and perform postinstallation tasks

From a user's standpoint, the network operating system is the focal point of a network. A computer's operating system is what users interact with when accessing a network's resources. Indeed, today's operating systems hide the details of network access so well that users often realize they're connected to a network only when access to network resources fails. As a network administrator, you know the network operating system (NOS) is only one piece of the network puzzle, but you're likely to spend quite a bit of time with this piece because of all the network-specific services you must install and configure.

This chapter discusses components common to almost every OS, networked or not, and then describes specific network services that network operating systems provide. In addition, virtualization, a technology that's integral to most medium and large networks and even many small networks, is introduced. Finally, specifics of installing Windows Server 2008 and Linux are described.

Operating System Fundamentals

A computer's OS provides a convenient interface for users and applications to access computer hardware components. It controls access to memory, CPU, storage devices, and external input/output (I/O) devices (such as printers, webcams, scanners, and so forth). To fully understand the network services and tasks a contemporary NOS provides, you need a solid grasp of how an NOS manages its local resources. The following sections expand on some OS concepts introduced in Chapter 1, specifically:

- File systems
- Processes and services
- Kernel

The File System

A **file system** is the method by which an OS stores and organizes files and manages access to files on a storage device, such as a hard drive. File systems differ in how they allocate space for files, how files are located on a disk, what level of fault tolerance is built into the system, and how access to files is secured. Regardless of how these tasks are accomplished, contemporary file systems have the following objectives:

- Provide a convenient interface for users and applications to open and save files.
- Provide an efficient method to organize space on a drive.
- Provide a hierarchical filing method to store files.
- Provide an indexing system for fast retrieval of files.
- Provide secure access to files for authorized users.

The first objective was described in Chapter 1 in enough detail for the purposes of this book, so the next sections focus on the other four objectives.

Disk Drive Space Organization The storage space on a disk drive is divided into manageable chunks called sectors. On most disk drives, each sector is 512 bytes. To make storage processing more efficient, one or more sectors are grouped to make a cluster or

block. A cluster is the smallest amount of space that can be occupied by a file stored on the disk. For example, if you have a file system that groups four sectors to make a cluster, each cluster is 2048 (2K) bytes. So if you store a file that's 148 bytes, it occupies one cluster of 2048 bytes, which wastes 1900 bytes of storage. The waste occurs because no other file can occupy any part of a cluster already occupied by another file. If you store a file that's 10,000 bytes, it occupies five clusters, with about 240 bytes of unused, wasted space.

You might think that having a smaller cluster size is the optimal way to organize your disk. This is true in some cases but only when mostly small files are stored on the disk because data is read from and written to the disk in cluster-sized chunks. So the smaller the clusters, the more read/write operations are required when using a file. Each read/write operation takes time, so the more operations required, the slower the system runs. If you store mostly large files on the disk, a larger cluster size usually results in better performance because fewer read/write operations need to be performed. In addition, smaller cluster sizes can lead to a fragmented disk, in which files are spread out all over the disk instead of being stored in consecutive locations. Fragmentation causes many more disk seek operations, which slows file access. Recall from Chapter 1 that the disk seek time is the amount of time required to move a drive's read/write heads to the correct position on disk platters to read or write clusters.

A disk's cluster size is selected when the disk is formatted. Most OSs set the cluster size to a medium value by default; for example, Windows sets the cluster size on NTFS-formatted disks to 4K. However, if you know you're going to be storing many files under 2K bytes, choose a smaller cluster size when you format to reduce wasted space. If you know you're going to store mostly files larger than 16K, choose a larger cluster size.

The formatting process groups sectors into clusters and maps all disk clusters for fast access. In addition, clusters are marked as unused. When you format a disk containing files, the data is actually still there, but the file system can no longer access the data because it doesn't know how to find the file data. Third-party disk recovery programs can often recover data from a formatted disk by bypassing the file system and reading the data in each cluster.

Hierarchical Filing Method Most file systems organize files in a hierarchy of folders or directories; the top of the hierarchy is called the "root" of the file system. ("Directory" is an older term for folder but is still used, particularly when discussing Linux file systems; however, the term "folder" is generally used in this book.) The root of the file system often represents a disk drive or other mass storage device, such as a flash drive. Off the root of the file system can be files and folders, with folders containing files and additional folders usually referred to as "subfolders." To navigate the file system with a GUI tool, such as Windows Explorer, users simply double-click folders to open them and view their contents. They can then navigate the file hierarchy by double-clicking subfolders. Figure 8-1 shows a logical diagram of a typical Windows file system.

Navigating the file system from a command prompt is a different proposition. Because users must use the exact syntax with no typos, those who began using computers when a GUI was the standard user interface are often frustrated when trying to navigate the file system at the command prompt. However, mastering the command prompt (or shell prompt in Linux) is a valuable skill to have as a Windows administrator and essential for Linux administrators.






Figure 8-1 A hierarchical file system

Courtesy of Course Technology/Cengage Learning



Commercial routers and switches are often configured with their own command-line user interfaces.

File Indexing System With today's large disk drives, more files can be stored on a hard drive, making it harder for users to find the files they need. To help solve this problem, most file systems include an indexing system that enables users to search for a file based on all or part of its filename or its contents. Users don't need to know where in the file system's hierarchy the file they need is located; they just need to know some or all of the filename or some keywords in the file. The indexing system maintains a database that's updated as files are created, modified, and deleted.

Secure Access to Files Computers are often shared at home or in the workplace. Each user might want to maintain a separate set of files and documents that other users can't access. In addition, files on a computer that's part of a network are potentially accessible to all users on the network. A file system's access controls, or permissions, can be used to allow only authorized users to access certain files or folders. In addition, access controls can be used to secure OS files from accidental corruption or deletion. Not all file systems have access controls, but the ones installed by default on current OSs do. Notably, the NTFS file system on Windows supports file and folder permissions, as do the Ext2 and Ext3 file systems common in Linux. The older DOS and Windows FAT16 and FAT32 don't support file and folder permissions, so any user logged on to the system console has full access to all files. The details of how these file systems work are beyond the scope of this book, but you learn more about using permissions in Windows and Linux in Chapter 9.



Hands-On Project 8-1: Navigating the Windows File System with the Command Line

Time Required: 20 minutes

Objective: Navigate the Windows file system with the command line.

Required Tools/Equipment: Your classroom computer with Windows installed

Description: In this project, you open a command prompt window and practice navigating the Windows file system from the command prompt.

1. Log on to your computer as an administrator.
2. Open a command prompt window. The prompt indicates your current file system context. Usually, when you open a command prompt window, you're placed in the file system context of your user profile. For example, if you log on as Administrator, your prompt is `C:\Users\Administrator>`. In this prompt, `C:` indicates the drive letter, and `\Users\Administrator` indicates the path. This file system context is the same as having the Administrator folder open in Windows Explorer (see Figure 8-2). In Windows file systems, the backslash (`\`) has two meanings. At the beginning of a path, it indicates the root or top of the file system. Anywhere else, it's used as a separator between folders, subfolders, and files. The forward slash (`/`) is used in many command-line programs to denote options for the command.

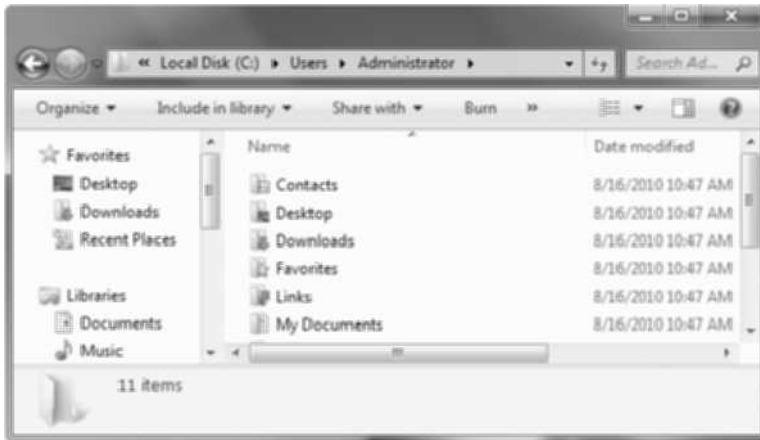


Figure 8-2 Windows Explorer with the `C:\Users\Administrator` folder open

Courtesy of Course Technology/Cengage Learning

3. At the command prompt, type `dir` and press **Enter**. The `dir` command means “directory.” You see a list of files and subfolders in the folder. As in Windows Explorer, the `dir` command doesn't display hidden files. To see hidden files, type `dir /ah` and press **Enter**. The `/ah` option tells `dir` to display files with the hidden attribute set. Type `dir /a` and press **Enter** to see all files. To see more options for the `dir` command, type `dir /?` and press **Enter**.



TIP

Remember to enter a space before any options you add to a command. Although not all commands require a space, many do, so it's best to get in the habit of entering one after the command.



4. To move to the root of the file system, type `cd \` and press **Enter**. The `cd` command means “change directory.” Your prompt should now be `C:\>`. Type `dir` and press **Enter**. To go to `C:\Windows\System32`, type `cd \windows\system32` and press **Enter**. Notice that the prompt changes to `C:\Windows\System32>` as a result. Type `dir` and press **Enter**. Several files scroll by quickly, but you can view them page by page by typing `dir /p` and pressing **Enter**. The `/p` option paginates the output. Press any key to see the next page of files, or press **Ctrl+C** to terminate the output if you don’t want to page through all the files.
5. Navigate back to the root of the file system. If you have more than one drive, you can switch drives by typing the drive letter and a colon and pressing **Enter**. For example, if you have a D drive, type `D:` and press **Enter**. The prompt changes to `D:\>`. If you don’t have a D drive, you get an error stating that the drive can’t be found. Type `C:` and press **Enter** to get back to the C drive, if necessary.
6. Next, create a new folder by typing `mkdir TestDocs` and pressing **Enter**. The `mkdir` command means “make directory.” To verify that the folder was created, type `dir` and press **Enter**, and then go to the new folder by typing `cd TestDocs` and pressing **Enter**. (Note: In Windows file systems, capitalization of filenames is ignored, so `TestDocs` is the same as `testdocs`.)
7. To create a subfolder, type `mkdir SubDocs1` and press **Enter**. Change to this subfolder by typing `cd SubDocs1` and pressing **Enter**. To go back to the `TestDocs` folder, type `cd \TestDocs` and press **Enter**. You must include the `\` character because you’re telling the file system that `TestDocs` is located directly under the root. To navigate to the `SubDocs1` subfolder, type `cd SubDocs1` and press **Enter**. You don’t use the `\` character in this command because `SubDocs1` is located directly under your current location. To go up one level in the file hierarchy, use the `..` notation: Type `cd ..` and press **Enter**, which takes you to the `TestDocs` folder. Type `cd ..` and press **Enter** again to get to the root.
8. Sometimes folder names are long and easy to mistype, so using a shortcut can be handy. Type `cd test` and press **Tab**. If `TestDocs` is the only folder name starting with “Test,” the command prompt fills in the rest of the name for you. If more than one folder begins with “Test,” the command prompt displays the first one in alphabetical order. Pressing **Tab** repeatedly cycles through all folders beginning with “Test.” Press **Enter**.
9. The command prompt maintains a history of commands you’ve used since the window has been open. If you have been entering long commands that you need to repeat, you can scroll through the history by pressing the up arrow. Press the **up arrow** repeatedly to scroll through your recent commands. Press **Esc** when you’re finished to cancel the command.
10. Type `mkdir ARealLongFolderName` and press **Enter**. Next, you make a mistake on purpose: misspelling the folder name. Type `cd ARealLongFoldName` (omitting the “er” in “Folder”) and press **Enter**. To correct this error, press the **up arrow**. Press the **left arrow** four times until the cursor is under the “N” in `Name`. Type `er` and press **Enter**. Making a correction in this fashion is called command-line editing.
11. If you want to create a text file in your current folder, type `notepad myfile.txt` and press **Enter**. When prompted to create the file, click **Yes**. Type whatever you like in the file,

- and then click **File, Exit** from the menu. When prompted to save the file, click **Yes**. Type **dir** and press **Enter** to verify that the file exists. To rename it, type **ren myfile.txt newfile.txt** and press **Enter**. To copy the file, type **copy newfile.txt newfile1.txt** and press **Enter**. Type **ren newfile.txt newfile.old** and press **Enter**.
12. To view only files with a .txt extension, type **dir *.txt** and press **Enter**. To see all files starting with “new,” type **dir new*** and press **Enter**. To delete newfile.old, type **del newfile.old** and press **Enter**. To delete all files in the ARealLongFolderName folder, type **del *** and press **Enter**. Press **y** and **Enter** when prompted.
 13. This project has shown you the basics of using the command line to navigate the file system in Windows. As a future network administrator, you’ll find yourself using the command line often. Close the command prompt window and log off Windows for the next project.



Hands-On Project 8-2: Navigating the Linux File System

Time Required: 20 minutes

Objective: Learn basic navigation of the Linux file system with the GUI and command line.

Required Tools/Equipment: A computer with Linux installed or a Linux Live CD; this project uses the Ubuntu 10.4 installation DVD, which contains a Live CD option, but others can be used.

Description: In this project, you use the Linux GUI and command line to navigate the file system.

1. Start Linux. If you’re using the Ubuntu 10.4 installation DVD as a Live CD, start the computer and press **Enter** when you see the language menu. Press **Enter** when the option “Try Ubuntu without any change to your computer” is selected.
2. To begin navigating the file system, click **Places, Computer** from the menu at the top to open File Browser, which looks similar to Windows Explorer. Double-click **File System** on the left (see Figure 8-3).
3. Remember that the file system in Linux doesn’t use drive letters. All available drives are accessed as folders off the root of the file system, which in Linux is designated with a forward slash (/). In the File Browser title bar, you see the / character, which indicates you’re currently at the root.
4. To access the contents of your CD-ROM drive, double-click the **cdrom** folder. (If there’s no CD in the drive, the folder is empty, but if you do have a CD in the drive, you see its contents.) Click **Back** on the File Browser toolbar to go back to the root.
5. Double-click the **home** folder. Home folders for all users on the system are in this folder. If you’re using the Ubuntu Live CD, you see a folder named Ubuntu, which is the user you’re currently logged on as. Double-click **Ubuntu** (or the home folder for your installation and logon account). You see folders similar to the ones created for each user in Windows. Exit File Browser.



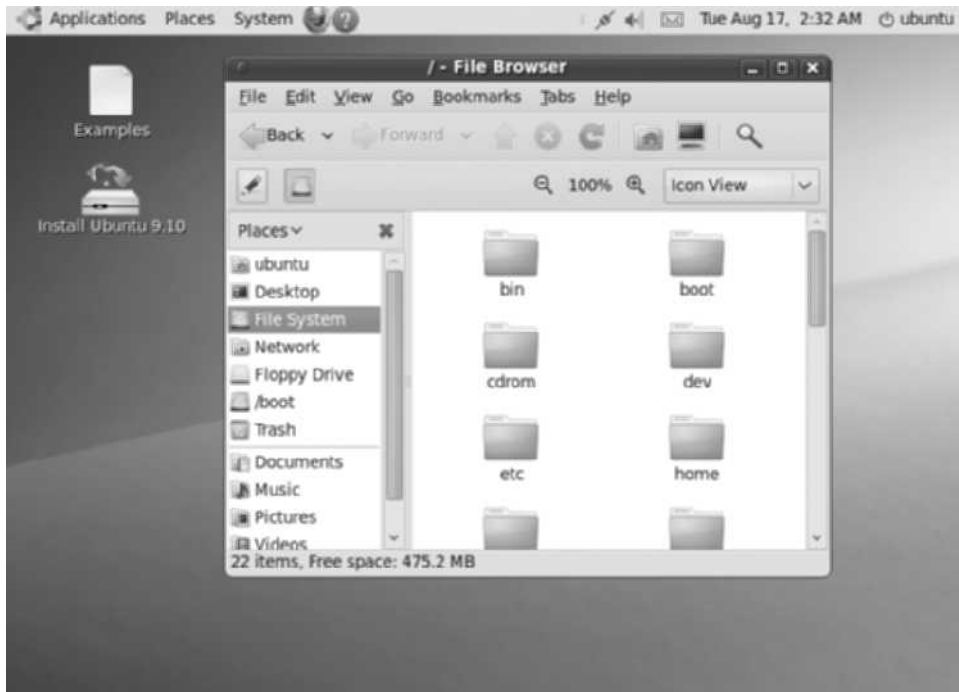


Figure 8-3 The Ubuntu 10.4 desktop with File Browser open to the root (/)

Courtesy of Course Technology/Cengage Learning

6. Open a command prompt window, called a terminal window in Linux, by clicking **Applications**, pointing to **Accessories**, and clicking **Terminal**. The prompt in a terminal window is different from the Windows command prompt. In Figure 8-4, the prompt is `ubuntu@ubuntu:~$`. The first `ubuntu` is the username, and the `ubuntu` after the `@` is the computer name. The tilde (`~`) following the colon indicates the user's home folder, and `$` is the end of the prompt. When you open a Linux terminal window, you're usually placed in your home folder.

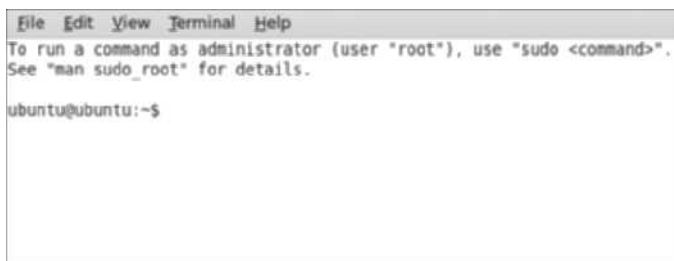


Figure 8-4 A Linux terminal window

Courtesy of Course Technology/Cengage Learning

7. Besides the terminal prompt, there are other ways to see the logged-on user's name and the computer name. To see your username, type `whoami` and press **Enter**. To view the computer name, type `hostname` and press **Enter**.

8. To move to the root of the file system, type `cd /` and press **Enter**. Remember that Linux uses forward slashes, and Windows uses backslashes. If you type a backslash accidentally, Linux thinks you're continuing the command on the next line and displays a `>` prompt. To get back to the normal prompt, you press **Enter**.
9. To view a list of files in the root, type `ls` (for "list") and press **Enter**. You can get back to your home folder in three ways. Type `cd /home/ubuntu` and press **Enter** (replacing `ubuntu` with your logon name if you aren't using Ubuntu), and then type `cd /` and press **Enter** again. A second way to get to the home folder is to type `cd` and press **Enter**, and a third way is to type `cd ~` and press **Enter**.
10. Create a new folder in your home folder by typing `mkdir newfolder` and pressing **Enter**. Type `ls` and press **Enter** to verify that the new folder has been created.
11. You see other folders in your home folder, including one named Documents. Type `cd documents` and press **Enter**. Typing "documents" with a lowercase "d" causes an error because the Linux file system is case sensitive. Press the **up arrow** to repeat the command. Press the **left arrow** until your cursor is over the "o" in documents and press **Backspace**. Type `D` and press **Enter**.
12. To create a new empty file, type `touch newfile` and press **Enter**. Verify that the file was created.
13. Type `gedit newfile` and press **Enter** to open newfile in gedit, a Notepad-like editor. Type whatever you like, click **Save**, and then close the gedit window.
14. To view newfile's contents from the command line, type `cat newfile` and press **Enter**. With a long file, you can use the `more` command to paginate the output. Type `more newfile` and press **Enter**.
15. If you can't remember the complete name of a command, type one or more letters in it and press **Tab**. If the command can be found by using the letters you typed, Linux completes the command. If there's more than one match, press **Tab** a second time to display all matches. Type `ge` and press **Tab** twice. A list of command beginning with "ge" is displayed. Type `cd ..` and press **Enter** to move back one folder. Type `ls Doc` and press **Tab**. Linux completes the command. Press **Enter**. You see newfile and newfile~, which is a backup of newfile made automatically when you changed it with gedit.
16. Type `cd Doc` and press **Tab** and then **Enter**. The `mv` command is used to rename or move files. To rename the backup file that was created, type `mv newfile~ newfile.bak` and press **Enter**.
17. The `rm` command means "remove." To remove the file you just renamed, type `rm newfile.bak` and press **Enter**.
18. To shut down Linux from the terminal, type `shutdown -h now` and press **Enter**. The `-h` option tells Linux to shut down and halt, and the `now` option means, well, now! (You use `-r` to restart the computer.) You can also specify a number of minutes to delay before the system shuts down by using `+m` instead of `now`, replacing `m` with the number of minutes to delay. If you get the message "Shutdown: Need to be root," you have to be an administrator to shut down the computer. The root user is the default administrator account in Linux. Type `sudo shutdown -h now` and press **Enter**. The `sudo` command means "do this as superuser."



Processes and Services

A **process** is a program that's loaded into memory and run by the CPU. It can be an application a user interacts with, such as a word-processing program or a Web browser, or a program with no user interface that communicates with and provides services to other processes. This type of process is usually called a **service** in Windows and a "daemon" in Linux and is said to run in the background because there's no user interface. Examples of services include Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks, which provide the client and server sides of Windows file sharing. Many TCP/IP Application-layer protocols, such as DNS and DHCP, also run as services.



Some OSs refer to processes as "tasks."

TIP

Network services are important because they allow your computer and applications to perform tasks they otherwise couldn't or would need additional built-in functionality to handle. For example, a Web browser is designed to request Web pages from a Web server and display them. However, because most people use the Web server's name rather than its address, a name lookup is required before a Web browser can do its main job. If it weren't for the DNS client service running on the computer, the Web browser would have to know how to perform DNS functions. Instead, as you learned in Chapter 5, the Web browser simply sends the DNS service a request for a name lookup, and DNS returns the IP address to the Web browser. The same is true of almost every network application. Most network access is initiated by using the server name, and DNS is always running as a process to provide the name lookup service, so the application is free to do what it was designed to do. In Windows 7, you can use a handy tool called Task Manager to see processes and services running on your computer, check how much CPU time and memory each process is using, and stop a process from running, if necessary. In Linux, you use the System Monitor application for these tasks.

An OS can run 2, 10, 100, or more processes seemingly at once by using multitasking. Whether a computer has one or multiple CPUs, it multitasks by using a method called **time slicing**, which occurs when a CPU's computing cycles are divided between more than one process. Each process receives a limited number of processor cycles before the OS suspends it and activates the next process. The act of changing to another process is called **context switching**.

When a process has work to do, as when a user types at the keyboard or a Web browser submits a request for a Web page, the CPU is notified. If it already has other processes waiting, the new request is put into a queue. If this process has a higher priority than others in the queue, it jumps to the front of the line. Because a CPU can execute many billions of instructions per second, processes waiting in the queue are usually scheduled to run quickly. This activity is perceived as many applications operating simultaneously because each time slice is a very short period. People can't distinguish instances of such a brief time period, so it creates the illusion that the CPU and OS are performing several tasks at once. There are two types of multitasking:

- **Preemptive**—With **preemptive multitasking**, the OS controls which process gets access to the CPU and for how long; when the assigned time slice expires or a higher priority task has work to do, the current process is suspended, and the next process gets access to the CPU.

- *Cooperative*—With **cooperative multitasking**, the OS can't stop a process; when a process gets control of the CPU, it maintains control until it satisfies its computing needs. No other process can access the CPU until the current process releases it.

Cooperative multitasking was used on older OSs, such as Windows 3.1. An application that stopped working because of an infinite loop could bring the entire system to a screeching halt because it never gave up control of the CPU. Thankfully, all current OSs use preemptive multitasking, so the OS or the user can terminate misbehaving applications.

Many applications are now designed so that different parts can be scheduled to run separately, almost as though they were different processes. Each part that can be scheduled to run is called a **thread**, which is the smallest unit of software that can be scheduled. A **multi-threaded application** has two or more threads that can be scheduled separately for execution by the CPU. For example, a multithreaded word-processing program might have one thread that waits for keyboard entry and then formats and displays the characters as they're typed and another thread that checks the spelling of each word as it's typed.

A multithreaded application benefits most when the OS and hardware support **multiprocessing**, which allows performing multiple tasks or threads simultaneously, each by a different CPU or CPU core. All current OSs support multiprocessing. Windows 7 supports up to two physical CPUs; a physical CPU is a chip installed in a socket on the motherboard. This means Windows 7 supports two CPUs, but each CPU can have one, two, four, or more cores. Windows Server 2008 supports up to 64 CPUs, depending on the edition. Most Linux OSs can support up to 32 or more CPUs.



The Kernel

If the CPU is the brain of a computer, the kernel is the office manager of the OS. Just as an office manager schedules everyone and everything and manages office resources, the kernel schedules processes to run, making sure high-priority processes are taken care of first; manages memory to ensure that two applications don't attempt to use the same memory space; and makes sure I/O devices are accessed by only one process at a time, in addition to other tasks. Because the kernel performs these important tasks, its efficiency and reliability are paramount to the OS's overall efficiency and reliability. Of course, the kernel is a process like any application, but it has the highest priority of any process, so when it needs to run, it takes precedence. You can't view it in Task Manager, and you certainly can't stop it, or the whole system would come crashing down.

Operating systems are designed in layers, as network protocols are, and the kernel is usually shown as the layer just above the hardware. This structure means nothing goes in or out without passing through the kernel—or at least without the kernel's approval. Figure 8-5 is a simplified illustration of the Windows OS structure, with the kernel near the bottom of the stack and above the hardware.

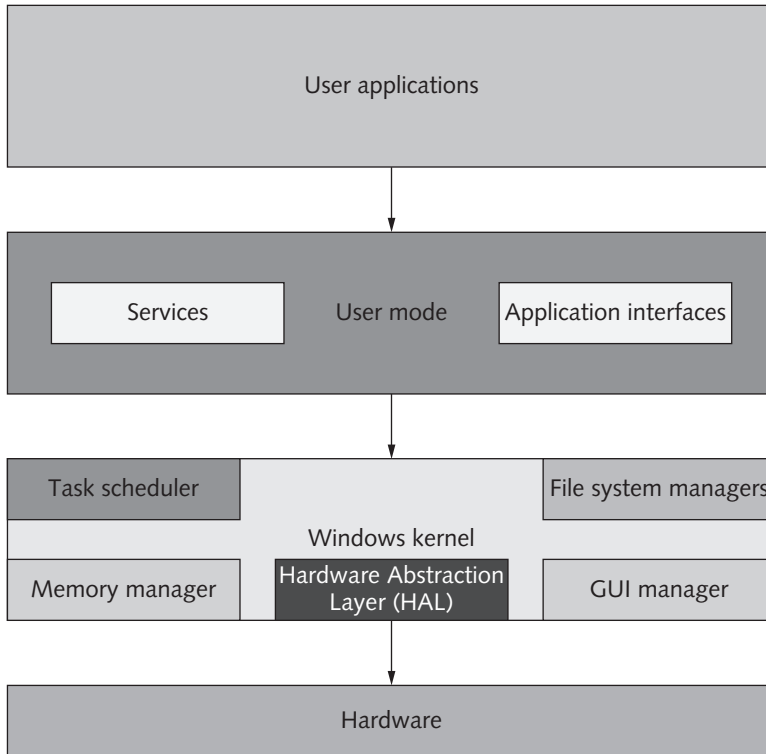


Figure 8-5 The Windows OS structure

Courtesy of Course Technology/Cengage Learning



Hands-On Project 8-3: Using Windows Task Manager

Time Required: 10 minutes

Objective: Use Windows Task Manager to view running processes, services, and real-time performance.

Required Tools/Equipment: Your classroom computer with Windows installed

Description: In this project, you use Task Manager to view processes and services.

1. Log on to your computer as an administrator.
2. Start Task Manager by right-clicking the **taskbar** and clicking **Start Task Manager**. If necessary, click the **Applications** tab. If you have no other windows open, no tasks are listed. Click **Start**, type **notepad** in the Start Search text box, and press **Enter** to start Notepad. Switch to the Task Manager window to see Notepad listed in the Applications tab.
3. Right-click **Notepad** in Task Manager and click **Go To Process**. The Processes tab opens and the notepad.exe process is selected. By default, Task Manager shows only processes started by the current logged-on user. To see all processes, click **Show processes from all users** (see Figure 8-6).




Figure 8-6 Windows Task Manager

Courtesy of Course Technology/Cengage Learning

- To sort running processes by the percentage of CPU time they're using, click the **CPU** column. You'll probably see System Idle Process jump to the top, and its CPU percentage will be in the high 90s. When a Windows OS has no real work to do, the kernel schedules the System Idle Process with CPU time. When the CPU percentage for System Idle Process is a high value, the computer isn't very busy; if it's a low value, other processes are using the CPU a lot.
- Click **View, Select Columns** from the menu, and then click the **CPU Time** check box. This column shows you the total CPU time a process has used since it was started. Next, scroll down and click **Command Line**. This column shows you the name of the actual command that loaded the process. Click **OK**.
- Click the **CPU Time** column to sort entries by overall amount of CPU time in *hours:minutes:seconds* format. Notice that several processes are named `svchost.exe`, which is used to start many services in Windows. The Command Line column gives you a better idea of which service each `svchost.exe` entry refers to and can also help you track down what application a process belongs to, based on the path to the application.
- Find one of the `svchost.exe` entries ending with "LocalServiceNetworkRestricted" in the Command Line column. Right-click the `svchost.exe` entry and click **Go to Service(s)**. The Services tab opens and the services started by that process are highlighted. Scroll through the Services tab to see all the highlighted services.
- In the Services tab, click the **Status** column to sort services by status. Running services are listed first, and stopped services are listed next. Scroll through the services to see how many are running. Clearly, a lot is going on behind the scenes on your computer.
- Click and then right-click the **Dhcp** service and click **Go to Process**. The Processes tab opens with the `svchost.exe` process that started the Dhcp service highlighted.



- Click the **Performance** tab. The CPU Usage box at the top left shows the total CPU % currently being used, and at the top right, the CPU Usage History box shows a one-minute history of usage in line graphs for each CPU or CPU core on your system. Figure 8-7 shows the Performance tab on a system with four CPU cores. Boxes for current memory usage and memory usage history are also displayed.




Figure 8-7 The Performance tab

Courtesy of Course Technology/Cengage Learning

- Click the **Processes** tab. Right-click **notepad.exe** and click **End Process**. Click **End process** when prompted, and Notepad closes. You can use this feature to exit an application that no longer responds to the mouse or keyboard.
- Close Task Manager, and log off Windows for the next project.



Hands-On Project 8-4: Displaying Linux Processes

Time Required: 10 minutes

Objective: View processes running in Linux.

Required Tools/Equipment: A computer with Linux installed or a Linux Live CD; this project uses the Ubuntu 10.4 installation DVD, which contains a Live CD option, but others can be used.

Description: In this project, you use Linux System Monitor and the `ps` command to view processes and performance information.

1. Start Linux. If you're using the Ubuntu 10.4 installation DVD as a Live CD, start the computer and press **Enter** when you see the language menu. Press **Enter** when the option "Try Ubuntu without any change to your computer" is selected.
2. Start System Monitor by clicking **System**, **System Monitor** from the menu at the top. Click the **Processes** tab, and then click the **% CPU** column twice to sort from highest to lowest CPU %. Linux doesn't show a system idle process. System Monitor itself is probably the biggest user of the CPU now. Most processes are listed as Sleeping, which means they're loaded into memory but don't have work to do.
3. Click **View**, **Active Processes** from the menu to see only processes with a status of Running. Periodically, you might see other processes that begin running listed here. Open some other windows if you want to see other processes listed.
4. Click the **Resources** tab, which shows information similar to the Performance tab in Task Manager.
5. Click the **Firefox** icon (orange and blue globe) at the top of the desktop. You should see the CPU utilization spike higher as well as network utilization. Close Firefox and System Monitor.
6. Click **Applications**, point to **Accessories**, and click **Terminal**. To see all currently running processes, type `ps -A` and press **Enter**. Type `ps -A | less` and press **Enter** to paginate the output. You can use the up and down arrows and Page Up and Page Down keys to scroll through the output. Press `q` to quit.
7. Type `top` and press **Enter** to see a real-time view of the top CPU users and other statistics. Press `q` to quit.
8. To shut down Linux, type `sudo shutdown -h now` and press **Enter**.



Network Operating System Overview

Desktop OSs now include many features once reserved for a server OS, such as file and printer sharing and file system security. Indeed, a desktop OS is classified as an NOS because it has these networking features built in, but many features and functions are still reserved for NOSs designed to be installed on a server. The determining factor of whether you need a server NOS or a client NOS is what role the computer will play in your network.



Because the distinction between an OS and NOS is largely moot with contemporary OSs, this book uses the term OS rather than NOS. When a distinction is made, it's between a client or desktop OS and a server OS.

As you know, computers in a network usually play one of two roles: a client or a server. Although contemporary OSs allow servers to perform client tasks and clients to perform server tasks, most vendors have specific versions of their OSs to fulfill these roles. The client version generally comes configured with client software, such as Web browsers, DNS and DHCP clients, and file-sharing clients. Most server versions also include client software but have server components, such as Web servers, DNS and DHCP servers, and file-sharing servers. In addition, advanced server OSs usually include directory services, additional authentication options, fault-tolerance features, and virtualization.

The Role of a Client Operating System

The client OS is where network users spend all their time. Its purpose is to run applications, which often access network resources. Most desktop computers run a client OS equipped with the following network client software:

- DHCP client
- DNS client
- HTTP client (Web browser)
- File-sharing client
- E-mail client

Other client software can be installed on a network client, such as the client side of a client/server database application, but these specialized applications are beyond the scope of this book. The preceding list of client software is installed on most OSs and used by most users.

DHCP Client The DHCP protocol was covered in Chapter 5, so this chapter examines its role in the context of a client OS. As you know, a computer can be assigned an IP address statically or dynamically with DHCP. DHCP is usually preferred because it allows centrally managing the assignment of IP addresses throughout an organization, which means fewer chances of misconfiguring devices by assigning the wrong address.

When an OS is first installed, IP address assignment is done through DHCP by default, so if a DHCP server is running on the network, the client OS gets an address and is up and running. If no DHCP server is running, the client assigns itself an IP address with APIPA.

Computers need more than just an IP address and subnet mask to operate in most networks. They need a default gateway if they access computers on other networks, including the Internet, and the address of a DNS server that can be queried to resolve computer and domain names to IP addresses. DHCP servers are configured to provide these additional addresses when a client requests an IP address. When a computer requests its IP address configuration, the process involves the following broadcast packets if the computer has no address assigned or its address lease has expired:

- *DHCPDiscover*—The client announces to the network that it's looking for a DHCP server from which to lease IP address settings.
- *DHCPOffer*—The server replies and offers the client an IP address for lease.
- *DHCPRequest*—The client wants the offered IP address.
- *DHCPAck*—The server acknowledges the transaction, and the client can now use the IP address.

After a client has the IP address configuration, it can begin using TCP/IP. The IP address is just a lease that must be periodically renewed. When half the lease is over, the client sends a unicast DHCPRequest packet to the server that leased it the address. The server sends a unicast DHCPAck packet to indicate that the address has been renewed.

When IP addresses are assigned with DHCP, a station's address can change periodically, especially if it's turned off when the lease time expires. You might want to manage IP address configurations with DHCP but still assign certain devices, such as network printers and some

workstations, addresses that don't change. To do this, you configure a reservation address on the DHCP server (discussed later in "The Role of a Server Operating System").

The DHCP client software runs as a service that starts when the computer starts. You can stop, start, restart, and view the status of the DHCP Client service in the Services control panel in Windows (see Figure 8-8). This service runs even if your IP address is assigned statically. To prevent it from running, you can disable it in the Services control panel or from the command line with the `net` command.



Figure 8-8 The DHCP client service

Courtesy of Course Technology/Cengage Learning

DNS Client The DNS client is responsible for communicating with a DNS server to resolve computer and domain names to IP addresses, so it's referred to as a "resolver." As discussed in Chapter 5, DNS resolvers maintain a local cache of the results of recent DNS lookups. The resolver cache speeds communication because it eliminates the need to communicate with a DNS server for recently looked up records.

An OS must be configured to use DNS. At the very least, a client computer needs one address of a DNS server it can query. In Windows, the first DNS server configured is called the preferred DNS server, and the second one is the alternate DNS server (see Figure 8-9).





Figure 8-9 Preferred and alternate DNS servers in Windows

Courtesy of Course Technology/Cengage Learning

When a client computer tries to resolve a computer name to an address, the DNS resolver attempts to append a domain name to the computer name because DNS servers require a domain name in addition to a computer name. In Windows, the default domain appended to DNS lookups is called the primary DNS suffix. This value is set when a computer is added as a member of a Windows domain, or it can be set manually. To view it, go to the Computer Name tab in the System Properties dialog box (see Figure 8-10). In this figure, for example, if a user attempts to contact server1, the DNS resolver sends the query to the DNS server as server1.mydomain.local.



Figure 8-10 Viewing the primary DNS suffix

Courtesy of Course Technology/Cengage Learning

Some environments are more complicated, with multiple domains that are accessed frequently. If users should be able to access computers in different domains with only their names, the DNS resolver can append a list of domains, or DNS suffixes, to computer names automatically. If the first query isn't successful, the next suffix is tried, and so forth. You can create the list of DNS suffixes the DNS resolver uses in the DNS tab of the Advanced TCP/IP Settings dialog box (see Figure 8-11).



Figure 8-11 DNS suffixes used to resolve names to addresses

Courtesy of Course Technology/Cengage Learning

In Figure 8-11, notice the “Register this connection’s addresses in DNS” check box. Windows supports Dynamic DNS (DDNS), which allows computers and other devices to contact their primary DNS server whenever their name or address changes. If the contacted DNS server allows DDNS, the server creates or updates the DNS host record in its database automatically.

Like the DHCP client, the DNS Client runs as a service that can be configured in the Services control panel in Windows.



Windows 9x doesn't use DNS to resolve Windows computer names; instead, the older NetBIOS or WINS protocols are used. Although these protocols are still supported in current Windows OSs, they're approaching obsolescence. DNS is used by Windows 9x computers only when resolving Internet names.



Hands-On Project 8-5: Viewing DHCP Client and DNS Client Status

Time Required: 10 minutes

Objective: Use the Services control panel to view DHCP Client and DNS Client status.

Required Tools/Equipment: Your classroom computer with Windows installed

Description: In this project, you use the Services control panel to view the status of DNS Client and DHCP Client services, and then use the command line to view the same information.

1. Log on to your computer as an administrator.
2. Click **Start**, type `services.msc`, and press **Enter** to start the Services control panel. The Services control panel is somewhat different from the Services tab in Task Manager, which enables you to stop or start a service and view its status but not change other properties, such as its startup type or how it logs on to the system.
3. Scroll down until you find the DHCP Client service. Notice that its status is Started. Double-click **DHCP Client** to open its properties (see Figure 8-12).




Figure 8-12 The DHCP Client Properties dialog box

Courtesy of Course Technology/Cengage Learning

4. Click the **Startup type** list arrow to view the available options. You shouldn't disable or stop (unless you restart it again) the DHCP Client because it's used to register and update your computer's DNS record. So even if you aren't getting an IP address via DHCP, the DHCP Client should remain running. Above the Startup type list, notice the

path to the executable file. When looking for the DHCP Client in the Processes tab in Task Manager, you should look for this path.

5. Click the **Log On** tab. Most services are started by using a special account: Local Service, Local System, or Network Service. The password is automatically changed periodically for security reasons, so you shouldn't have to change it.
6. Click the **Recovery** tab. You use this tab to specify what should happen if the service fails. In most cases, the service attempts to restart twice. You can specify actions for the computer to take (such as restarting) if the service encounters errors when trying to start.
7. Click the **Dependencies** tab, where you can view other processes or services this service depends on to run and other processes or services that depend on this service. Before stopping a service you think you don't need, check the dependencies to make sure another service you do need isn't affected. Click **Cancel**.
8. Next, examine the DNS Client properties. They're largely the same as for the DHCP Client, except that a different svchost.exe command is used to start DNS. Note also that the service's name is shown as Dnscache. When you're finished, close the Services control panel.
9. Open a command prompt window. To view the status of services from the command line, you use the `sc` command. Type `sc query` and press **Enter** to view the status of all running services. To view the status of DHCP, type `sc query dhcp` and press **Enter**, and to view the status of DNS, type `sc query dnscache` and press **Enter**.
10. Sometimes restarting a service is necessary. A computer restart does this, but you can also do it in the Services control panel or from the command line. Type `sc stop dhcp` and press **Enter**. The status of DHCP is displayed as STOP_PENDING. (If you see the message "A stop control has been sent to a service that other running services are dependent on," try the `sc stop dhcp` command again until it's successful.) Type `sc query dhcp` and press **Enter** to see that DHCP has been stopped. Type `sc start dhcp` and press **Enter** to start the service again.
11. If you're going on to the next project, leave the command prompt window open.



Hands-On Project 8-6: Using the DNS Command-Line Program Nslookup

Time Required: 10 minutes

Objective: Use Nslookup to work with DNS.

Required Tools/Equipment: Your classroom computer with Windows installed

Description: In this project, you use the Nslookup command-line program to work with DNS.

1. If necessary, log on to your computer as an administrator, and open a command prompt window.
2. Nslookup can be used by entering a single command and seeing the results or interactively, in which you start the Nslookup program and enter commands until you exit. Type `nslookup www.course.com` and press **Enter**.
3. The first two lines of output show the name and the address of the DNS server Nslookup is using for the DNS query. The next lines of output show the computer's fully qualified



domain name (FQDN) and IP address. Some names have more than one IP address. For example, type `nslookup www.google.com` and press **Enter** to see multiple IP addresses.



Several IP addresses are shown for *www.google.com* because many servers can respond to this name. If you go to the Google.com Web page, your DNS client queries the DNS server for its IP address. The list of addresses you see is returned to the client. The client typically chooses to request the page from the first address in the list. If you run the same query again, addresses are returned in a different order, which ensures that access to *www.google.com* servers is load-balanced among all servers.

4. To enter interactive mode, type `nslookup` and press **Enter**. The default server is displayed. To do a reverse lookup with Nslookup, type `198.60.123.100` and press **Enter**.
5. You can change which server Nslookup queries by default, which is useful if your DNS server doesn't seem to be responding or you're concerned that it's returning inaccurate results. Type `server 8.8.8.8` and press **Enter**. The name returned for the DNS server is a public server provided by Google.com. This command doesn't change your DNS server settings outside Nslookup.
6. Type `?` and press **Enter** to get a list of commands you can enter in Nslookup. Be aware that Nslookup doesn't use the local resolver cache, so all queries are sent to the DNS server.
7. Exit Nslookup by typing `exit` and pressing **Enter**. Close the command prompt window, and leave Windows running for the next project.

HTTP Client The HTTP client software doesn't run as a service as DNS and DHCP do; instead, it's built into programs that use it, such as Web browsers. The HTTP client requests Web pages from Web servers and can also transfer large files; largely replacing File Transfer Protocol (FTP) for that purpose. FTP lacks an option for secure communication, but HTTP has the advantage of being able to create secure connections by using HTTPS. The "S" designates the use of Secure Sockets Layer, a protocol that authenticates the Web server to the client, encrypts data before it's transferred, and decrypts it on receipt. For normal, unencrypted connections, HTTP uses TCP port 80 by default. When HTTPS is specified in the URL, port 443 is used instead. Web servers supporting SSL respond to requests on port 443, and your Web browser indicates that you have a secure connection by displaying a locked padlock icon.

File-Sharing Client A file-sharing client allows the computer to access files and printers on the network. When a user or an application requests a resource—such as a printer or a data file—a **redirector** intercepts the request and then examines it to determine whether the resource is local (on the computer) or remote (on the network). If the resource is local, the redirector sends the request to the local software component for processing. If the resource is remote, the redirector sends the request over the network to the server hosting the resource.

With redirectors, network resources can be accessed as though they were local. For example, a user or user application doesn't distinguish between a printer connected to a local USB port and one connected to the network. In addition, with drive mapping, shared network folders are accessed just like a drive that's physically attached to the system—at least from the user's point of view. In Windows, the redirector component is part of Client for Microsoft Networks (listed

in the network connection properties). This client software is designed to access shared folders and files on servers by using the Server Message Block (SMB) protocol. In Windows, the two most common ways to access a shared resource are using the UNC path or mapping a drive.

In Chapter 1, you used the UNC path to access a shared folder, which has the syntax `\\server-name\sharename`. The *server-name* is the name of the computer where the shared resource resides. You can also use the server's IP address in place of its name. The *sharename* is the name given to the folder or printer when it was shared. (Sharing folders and printers is discussed later in “The Role of a Server Operating System.”) Typing a UNC path in the Start Search text box opens an Explorer window showing the shared folder's contents. You can access a subfolder or file in the share directly by continuing the UNC path, as in `\\server-name\sharename\subfolder\file.extension`.



TIP

Linux systems also use the UNC path to access shared resources, but on Linux systems, forward slashes are used in place of backslashes.

You can use the UNC path to access shared folders and printers, but you must type the path every time you need it or create a shortcut with the UNC path as the target. One widely used method of making access to shared files easier (particularly those that are used often) is drive mapping, which associates a drive letter with the UNC path to a shared folder. Drives are usually mapped by using Windows Explorer or the `net` command. To use Windows Explorer, simply type the server portion of the UNC path in the Start Search text box. A list of shared folders and printers the server is hosting is displayed. Right-click a shared folder and click Map network drive, as shown in Figure 8-13. You then have the opportunity to pick a drive



Figure 8-13 Mapping a drive in Windows Explorer

Courtesy of Course Technology/Cengage Learning

letter (one that's not already in use) and can choose to have Windows reconnect to the share with the same drive letter every time you log on.

Another method of mapping a drive is using the `net` command. This method is often used by administrators in a logon script, a set of commands that run when a user logs on to a Windows domain. The command to map a drive with the `net` command is `net use drive-letter: \\server-name\sharename`.

The *drive-letter* is an unused drive letter and must be followed by a colon (:). The command can be entered at the command prompt or placed in a batch file. A **batch file** is a text file containing a list of commands you ordinarily type at the command prompt. To run a batch file, enter its name at the command prompt or double-click the file in Windows Explorer. Batch files are useful for storing long complex commands that are used often or a series of commands that are always used together.

For the sake of comparison, Linux doesn't use drive letters at all. Instead, Linux file systems are based on the concept of a file system root designator, which is simply the `/` character. All local and network drives and folder are accessed from the root as folders (or directories, as most Linux users call them). A drive or network share is mounted into an empty directory so that it becomes part of the file system hierarchy. So to access a shared folder in Linux, you create a new directory at the root of the file system or in a subdirectory, and then mount the shared folder in the new directory.

The protocol used in Windows to share files and printers is SMB, also known as Common Internet File System (CIFS). Aside from file and printer sharing, SMB also provides a mechanism for interprocess communication between computers. Interprocess communication allows processes running on computers to communicate with one another for configuring and administering a computer over the network, for example.

Linux also supports SMB implemented as an installation option called Samba, but the native file-sharing protocol in the Linux environment is Network File System (NFS). NFS works much like SMB, in that NFS clients mount the shared folder into their local file system so that it appears as a local resource to both users and applications accessing it.

Using shared printers in Windows is even easier than mapping a drive. Simply right-click the shared printer in the Explorer window and click Connect. A new printer is created in your Printers folder.



Hands-On Project 8-7: Mapping a Drive Letter

Time Required: 10 minutes

Objective: Map a drive letter by using different methods.

Required Tools/Equipment: Your classroom computer with Windows installed

Description: In this project, you map a drive letter to a network share that you create on your computer to avoid network and permission problems.

1. Log on to your computer as an administrator.
2. Click **Start, Computer**. Double-click the D drive if you have one; otherwise, double-click the C drive.

3. Create a new folder named **MyShare**. Right-click **MyShare**, point to **Share with**, and click **Specific people**. You can add users who can access the share in this dialog box, but because only the current user will access it, click **Share**. You're notified that the folder is shared. Click **Done**.
4. Click **Start**, type `\\localhost`, and press **Enter**. The `\\localhost` refers to your own computer, so a window opens showing available shares, including **MyShare**. Normally, you wouldn't map a drive to a folder on your own computer, and you would replace "localhost" with the name of a server hosting the share. You're using localhost just for practice. Right-click **MyShare** and click **Map network drive**.
5. You can choose the drive letter to map to this share. Click the **Drive** list arrow and click **X:**. Click to clear the **Reconnect at logon** check box. If you leave this option selected, the drive is mapped to the share each time you log on. Notice that you can also choose to connect to the share with different credentials (username and password). Click **Finish**.
6. An Explorer window opens, showing the share's contents. Close all windows. Click **Start, Computer**. You see the drive letter and share name listed under Network Location. Right-click **MyShare** and click **Disconnect** to delete the drive mapping.
7. Open a command prompt window. To map a drive letter from the command line, type `net use x: \\localhost\MyShare` and press **Enter**. To display current connections to shared resources, type `net use` and press **Enter**.
8. Click in the Computer window. The X drive letter is listed under Network Location again.
9. At the command prompt, type `net use /?` and press **Enter** to see a list of options for the `net use` command. You can use the `/persistent` option to make a drive mapping reconnect each time you log on. You can also connect with a different set of credentials. Type `net use x: /delete` to delete the drive mapping, and close the command prompt window.
10. To create a batch file for mapping a drive, open a new document in Notepad, and type the following two lines:
`net use X: /delete`
`net use X: \\localhost\Myshare`
11. The first command deletes any existing drive mappings for the X drive. Click **File, Save As**. In the left pane of the Save As dialog box, click **Desktop**. Click the **Save as type** list arrow and click **All Files**. In the File name text box, type `mapX.bat` and click **Save**. Close Notepad.
12. On your desktop, double-click `mapX`. Open the Computer window to verify that the X drive mapping has been created. Right-click **Myshare** and click **Disconnect**.
13. Batch files can come in handy if you need to connect to another computer periodically but don't want a permanent drive mapping. They're especially useful if you often need to enter a long command because they save you the time of having to remember and enter the command each time you need it. Close all open windows, and leave Windows running for the next project.



Hands-On Project 8-8: Creating and Connecting to a Shared Printer

Time Required: 10 minutes

Objective: Create a shared printer and then connect to a shared printer.

Required Tools/Equipment: Your classroom computer with Windows installed; a computer with a shared printer that all student computers can connect to, or students can connect to each other's shared printers.

Description: In this project, you create a shared printer and then connect to it.

1. If necessary, log on to your computer as an administrator.
2. Click **Start, Devices and Printers**. Click **Add a printer** in the Devices and Printers window.
3. Click **Add a local printer**. In the Choose a printer port window, leave the default option **Use an existing port** selected and click **Next**. In the Install a printer driver window, normally you select the printer's manufacturer and model, but because there's no physical printer, just accept the default selection and click **Next**.
4. In the Printer Sharing window, make sure **Share this printer so that others on your network can find and use it** is selected. For the share name, type **MyPrinter** and click **Next**. If you were actually installing a printer, you would click "Print a test page" in the next window. Just click **Finish**.
5. To connect to a shared printer, ask your instructor whether a printer share is set up. If not, use the shared printer one of your classmates created. Click **Start**, type **\\computer**, and press **Enter** (replacing *computer* with the name of the computer sharing the printer).
6. When the window opens, right-click the shared printer and click **Connect**. Open the Devices and Printers window to verify that the printer was created.
7. Close all open windows.

E-mail Client E-mail is the lifeblood of communication for most businesses and the people who work in them. Its complexity and importance combine to make it one of an IT department's biggest headaches. Most users, however, simply see e-mail clients as one of several communication tools they use every day and don't think much about how it works.

There's more to e-mail than just typing a message, attaching a file, and sending it to a colleague. E-mail is based on its own set of protocols, just as Web browsing and file sharing are. The most common e-mail protocols are as follows:

- *Post Office Protocol version 3 (POP3)*—E-mail clients use this protocol to download incoming messages from an e-mail server to their local desktops. POP3 clients must manage messages locally (not on the server, as they can with IMAP).
- *Simple Mail Transport Protocol (SMTP)*—This protocol is the standard protocol for sending Internet and other TCP/IP-based e-mail. POP3 is used to retrieve e-mail, and SMTP is used to send e-mail.
- *Internet Message Access Protocol (IMAP)*—This standard has advanced message controls, including the capability to manage messages locally yet store them on a server and fault-tolerance features.

Sending an e-mail message involves a series of steps. After a message has been written and the user clicks the Send button, the e-mail client software contacts an SMTP server. The SMTP server's address is part of the e-mail client's configuration. The SMTP server receives

the message, looks up the domain of the destination address, and contacts an SMTP server at the destination domain. The destination SMTP server sends the message to the POP3 server containing the recipient's mailbox. The POP3 server deposits the message in the recipient's mailbox, where it sits until the mailbox owner instructs the e-mail client software to retrieve messages.

When you start your e-mail client or click the Get Mail (or equivalent) button in your client, the client uses POP3 to contact the POP3 server containing your mailbox. The POP3 server forwards waiting messages to the client software and usually deletes them from the server. If you're using IMAP instead of POP3, only the message headers are sent, which include sender information and the subject. Only when you click the message header is the body of the e-mail sent. Messages aren't deleted from the server until you delete them with the client software. With IMAP, you have the advantage of being able to open and read e-mail on one computer and download the same messages on another computer. Because IMAP doesn't delete messages on the server automatically, they can be downloaded and opened from multiple locations. In addition, users' mailboxes can be backed up on the server so that users don't have to back up e-mail on client computers. Simulation 18 shows how e-mail works with SMTP and POP3.



Simulation 18: How e-mail works



Some ISPs support IMAP and some don't because of the extra space undeleted messages use on their servers. Also, POP3 has an option to leave downloaded messages on the server until they're deleted from the client computer, but most ISPs don't support this feature.

Although the use of e-mail client and server software is still the norm in medium and large businesses, many small businesses and home users access their e-mail by using a Web browser interface from sites such as Gmail.com and Yahoo.com. In this case, the same processes occur when transmitting and receiving e-mail, except the Web server you connect to for accessing your mail performs these tasks instead of a locally installed e-mail client.

The Role of a Server Operating System

Early NOSs installed on servers in PC networks were dedicated to providing network services to client computers and couldn't even run user applications and client network software. Today, the OS installed on a desktop computer is largely the same as that installed on a server, with the differences being the number and type of network services available and how server resources are used. For example, Windows Server 2008 is configured with Client for Microsoft Networks, and DHCP and DNS clients run user applications without complaint. However, you can install DHCP and DNS server components on Windows Server 2008 as well as the Active Directory directory service, all of which are unavailable in Windows 7. In Linux distributions, some installation programs let you choose a desktop or server configuration, but some distributions, such as Red Hat Enterprise, are designed as server OSs.

Memory, CPU, and disk usage on client OSs are optimized to run user applications and client network software. On server OSs, use of these resources is typically optimized to run network



services in the background to speed up responses to client requests. In addition, server OSs have more security and fault-tolerance features. The following is a list (but by no means an exhaustive list) of the features and functions most server OSs provide in a typical network:

- Centralized user account and computer management
- Centralized storage
- Infrastructure services, such as name resolution and address assignment
- Server and network fault tolerance
- Additional server features

Centralized User Account and Computer Management

Among the most compelling reasons to design a network, even a small one, as a server-based network is centralized management of network resources, which includes the following functions:

- User authentication and authorization
- Account management
- Security policy management

User Authentication and Authorization **Authentication** is the process of identifying who has access to the network. The most common form of authentication is a logon with a username and password. Other forms include digital certificates, smart cards, and biometric scanners.

Authorization is the process of granting or denying an authenticated user's access to network resources. Both authentication and authorization require users (and sometimes devices) to have a user account that stores properties about the user, such as a logon name and password. A user account is also used to grant permissions for the user to access network resources.

Account Management Most OSs, including those designed as client OSs, now incorporate account management for the purposes of authentication and authorization, but account management is centralized in the server OS. To better understand centralized account management and, therefore, centralized authentication and authorization, consider a network in which account management is decentralized, as in a Windows workgroup network. As discussed in Chapter 1, each computer in this type of network maintains its own list of user accounts and controls access to its own resources. In a network of 10 computers, if each computer shares resources that are accessed by users on other computers, each user account must be created 10 times, once on each computer. The password for each user must also be maintained on each computer. If a user's password is changed on one computer, he or she has to remember a different password to access that computer or have the password changed on all 10 computers. You can see how keeping up with this system could become tiresome quickly.

The server version of Windows OSs includes a centralized account management, authentication, and authorization system called Active Directory. Active Directory is a directory service that allows users to log on to the network once with their username and password and access resources they're authorized for regardless of which computer stores the resource. When Active Directory is installed on a server, the server becomes a domain controller, and users and computers with accounts in Active Directory are referred to as domain members.

Figure 8-14 shows the Active Directory Users and Computers management console. In the left pane are folders used to organize accounts and resources for easier management. In the right pane are user and computer accounts, distinguished by different icons.

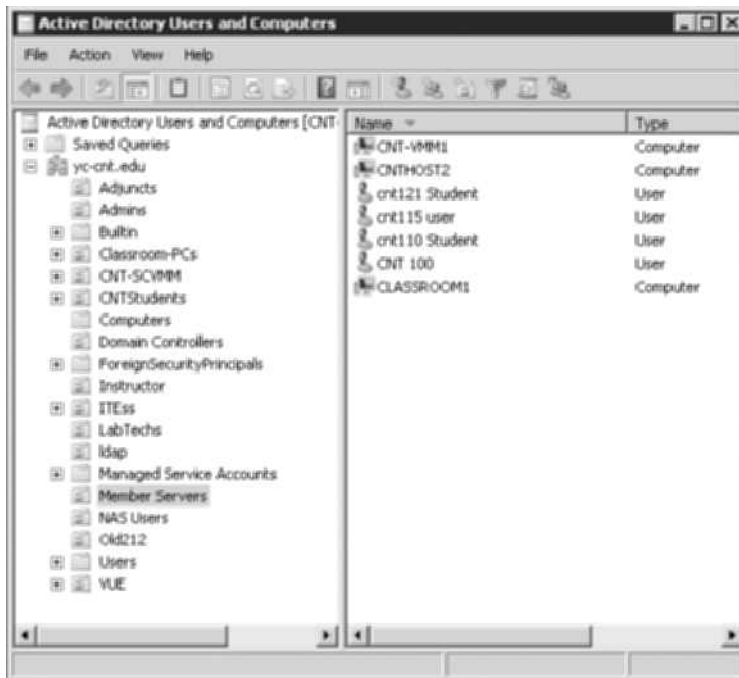


Figure 8-14 The Active Directory Users and Computers management console

Courtesy of Course Technology/Cengage Learning

When a computer running a desktop or server version of Windows becomes a domain member, an account is created for the computer in Active Directory. A computer becomes a domain member by changing its membership type from Workgroup to Domain in the Computer Name/Domain Changes dialog box accessed via the System Properties dialog box (see Figure 8-15). You work with accounts in Chapter 9.

Security Policy Management Aside from authentication and authorization, accounts in Active Directory are used to distribute and enforce policies for network use and security. These policies, called group policies in a Windows domain environment, can be applied to all domain members. Policies can range from user interface policies controlling what icons appear on the desktop and Start menu to security policies controlling password restrictions and what applications a user can run on a computer. They're just a few examples of the power of group policies; hundreds of different policy settings are available.

As discussed in Chapter 1, Linux OSs have a basic directory service for centralized logon called Network Information Service (NIS), but Lightweight Directory Access Protocol (LDAP), which Active Directory is based on, is gaining widespread acceptance in the Linux community. LDAP has the advantage of supporting both Windows and Linux user authentication and authorization.



Figure 8-15 Making a computer a domain member

Courtesy of Course Technology/Cengage Learning

Chapters 9 and 10 cover some details of how accounts and policies are created and managed. For now, you just need to know that these tasks are central to a server OS.

Centralized Storage

With large multimedia files being such a large portion of the data stored and processed on networks, network administrators are in a constant quest to better manage and maintain storage resources. Network storage includes file sharing, in which users store documents on network servers that other users can access. It also includes storing corporate e-mail, user files, application databases, and data backups, among many other resources.

Traditional servers use locally attached disk drives to store the installed OS and applications as well as user files. However, the amount of data stored in even small to medium networks is measured in dozens of terabytes, quite a burden for servers juggling numerous other network tasks. Although locally attached storage is still in common use, many network administrators are turning to specialized devices to help manage their storage requirements, including the following:

- Network-attached storage devices
- Storage area networks
- Cloud-based storage

Network-Attached Storage A **network-attached storage (NAS)** device is a dedicated server device designed solely for providing shared storage for network users. An NAS could be a regular server with NAS software installed or it could be a **network appliance**, a device equipped with specialized software that performs a limited task, such as file sharing. Network appliances are often packaged without video interfaces, so you don't configure them with an attached keyboard and monitor. They have a built-in Web server that you connect to with a Web browser to configure and manage the device. Many NASs integrate with Active Directory or an LDAP-based system for user authentication and authorization.

Storage Area Network A SAN, discussed in Chapter 1, is a high-speed, high-cost network storage solution that largely replaces locally attached drives on servers. SAN technology allows multiple servers to access an enormous amount of shared storage that appears as locally attached drives from the server and user’s perspective. Servers can even boot their OSs from a SAN instead of booting from local disks. This level of centralized storage offers better reliability and fault tolerance than traditional storage methods. Additionally, because storage is shared among several servers, power requirements for maintaining these systems are lower than those needed to maintain several servers with their own locally attached storage.

Cloud-Based Storage When a company’s storage needs have outgrown its storage capabilities, whether because of physical capacity limits or the lack of personnel to maintain in-house storage, the company can turn to the cloud. The cloud, in this context, simply means a third-party company that hosts computing solutions on behalf of other companies. In this case, the computing solution is network storage. With **cloud storage**, some or all of an organization’s data is stored on servers located offsite and maintained by a storage hosting company. The customer can manage storage by assigning permissions for user access and allocating storage for network applications and so forth without having to physically maintain the servers. If more storage is needed, the customer simply pays the storage hosting company for the additional space. The advantage of this approach is that the details of managing and backing up storage on local servers are offloaded to a third party, which enables a company to focus its monetary and personnel resources on business rather than IT tasks. Cloud-based storage isn’t for everyone. The data a company maintains might be too sensitive to trust to a third party, the data access speed might not be sufficient, or a host of other reasons. Cloud storage is a new model in network storage, and as you learn in Chapter 12, cloud computing in general is here to stay.



Infrastructure Services

Computers require infrastructure services for basic network functionality, including dynamic IP address assignment (DHCP) and name resolution services (DNS). Chapter 5 covered these protocols, and their client portions were discussed earlier in “The Role of a Client Operating System.” The following sections describe their server aspects.

DHCP Server A DHCP server is composed of the following elements:

- *IP address scope*—An **IP address scope** is a range of IP addresses the server leases to clients that request an IP address. In Windows, a scope is specified with starting and ending IP addresses, a subnet mask, and the address lease time, which can range from one minute to unlimited (meaning the address lease never expires). After the scope is created, an administrator can further configure it by using the following:
- *Scope options*—IP settings such as the default gateway, DNS servers, a domain name, and other address options are included in scope options. When a client requests an IP address, the client receives an address and a subnet mask from the scope and any options defined for the scope.



DHCP servers can maintain multiple scopes if they service more than one subnet.

- **Reservations**—A **reservation** is an IP address tied to a particular MAC address. When a client requests an IP address from the DHCP server, if the client’s MAC address matches an address specified by a reservation, the reserved IP address is leased to the client instead of getting it from the scope. In addition, reservations can have their own options that differ from regular scope options.
- **Exclusions**—An **exclusion** is one or more IP addresses excluded from the IP address scope; for example, if the scope ranges from 192.168.1.1 to 192.168.1.100, you can exclude addresses 192.168.1.1 through 192.168.1.10 if these addresses have been assigned statically.
- **DHCP server service**—This service runs in the background and listens on UDP port 69, the port reserved for client-to-server DHCP communication. The DHCP service responds to DHCP client requests for new IP addresses and IP address release and renewal requests.

After an address is leased, a record of the lease is stored in a database containing the IP address, the name and MAC address of the computer leasing the address, and the lease expiration time. Administrators can view the database’s contents to determine which computers are leasing which addresses. Figure 8-16 shows the Windows Server 2008 DHCP management console.

Client IP Address	Name	Unique ID	Lease Expiration
172.31.1.112	CNT212Server1.yc.cnt.edu	00188b060f7a	7/26/2010 7:56:18 AM
172.31.1.200	W2K3-Server1	00068b0c3fb	Reservation (inactive)
172.31.1.201	linuxServer	0040f457e466	Reservation (inactive)
172.31.1.202	disco.yc.cnt.edu	000c2990caab	Reservation (inactive)
172.31.1.203	CNTInas.yc.cnt.edu	0014fd100b4f	Reservation (active)
172.31.1.205	CNT-SERVER1	000c29c1e47f	Reservation (inactive)
172.31.1.206	172.31.1.206	001a4b4608c	Reservation (inactive)
172.31.1.210	yc.cnt.org	000c29969772	Reservation (active)
172.31.1.211	NFS2P2FAB.yc.cnt.edu	0001e65f2fad	Reservation (active)
172.31.1.212	YlueServer	000c2938f958	Reservation (inactive)
172.31.1.218	cnt-vmlogix	000c299a3f97	Reservation (inactive)
172.31.1.220	ubuntu.yc.cnt.edu	001a4b5208a7	Reservation (active)
172.31.1.225	CNT-VM041	001b21396100	Reservation (inactive)
172.31.1.250	CiscoRouter2600	000c054222c0	Reservation (inactive)
172.31.1.251	Mikrotik-VPN server	000c4254f6b0	Reservation (inactive)
172.31.209.1	cnt209-01.yc.cnt.edu	0024e633ed36	Reservation (active)
172.31.209.2	cnt209-02.yc.cnt.edu	0024e633efe3	Reservation (active)
172.31.209.3	cnt209-03.yc.cnt.edu	0024e633ed39	Reservation (active)
172.31.209.4	cnt209-04.yc.cnt.edu	0024e633ebac	Reservation (active)
172.31.209.5	cnt209-05.yc.cnt.edu	0024e633edeb	Reservation (active)
172.31.209.6	cnt209-06.yc.cnt.edu	0024e633ef53	Reservation (active)
172.31.209.7	cnt209-07.yc.cnt.edu	0024e633d4e9	Reservation (active)

Figure 8-16 The Windows Server 2008 DHCP management console

Courtesy of Course Technology/Cengage Learning

DNS Server DNS is a central component of every network for both Internet name resolution and local resource name resolution. The Linux environment has long used DNS for

name resolution; on Windows networks, DNS became the standard name resolution protocol starting with Windows 2000 Server. Before that time, Windows networks used Windows Internet Naming Server (WINS), a Windows-specific protocol for resolving Windows computer names. WINS is still supported in Windows Server 2008 but only as a legacy service for backward-compatibility with Windows 9x and older applications requiring WINS.

DNS servers are composed of the following elements:

- *DNS zones*—A **DNS zone** is a database of primarily hostname and IP address pairs that are related by membership in an Internet or a Windows domain. Each zone carries the name of the domain whose records it stores. Zone records are created manually by an administrator or dynamically by the host device. When a DNS client contacts a DNS server to resolve a name to an IP address, the domain name specified in the request is matched to the zones the DNS server manages. If a zone name matches the request, the zone is searched for a host record matching the hostname in the request. If the domain doesn't match a zone on the local DNS server, the server looks for a match in its cache. If there's still no match, the DNS server contacts other DNS servers by using its root hints (explained later in this list).
- *Cache*—When the local DNS server contacts another DNS server to satisfy a client's DNS query, the results are saved or cached so that if the same query occurs again, the local DNS server can respond without having to contact another server. Cached records expire after a period of time to prevent stale records.
- *Root hints*—When a DNS query can't be resolved from local zone records or cached records, a DNS server consults the root hints file, which contains a list of IP addresses of Internet root servers. Root servers maintain records for the Internet top-level domain (TLD) servers. TLD servers maintain records for DNS servers that manage second-level domains. These servers maintain different levels of domain information that form the basis of the hierarchical nature of the DNS system. Figure 8-17 shows a DNS query that involves root servers.

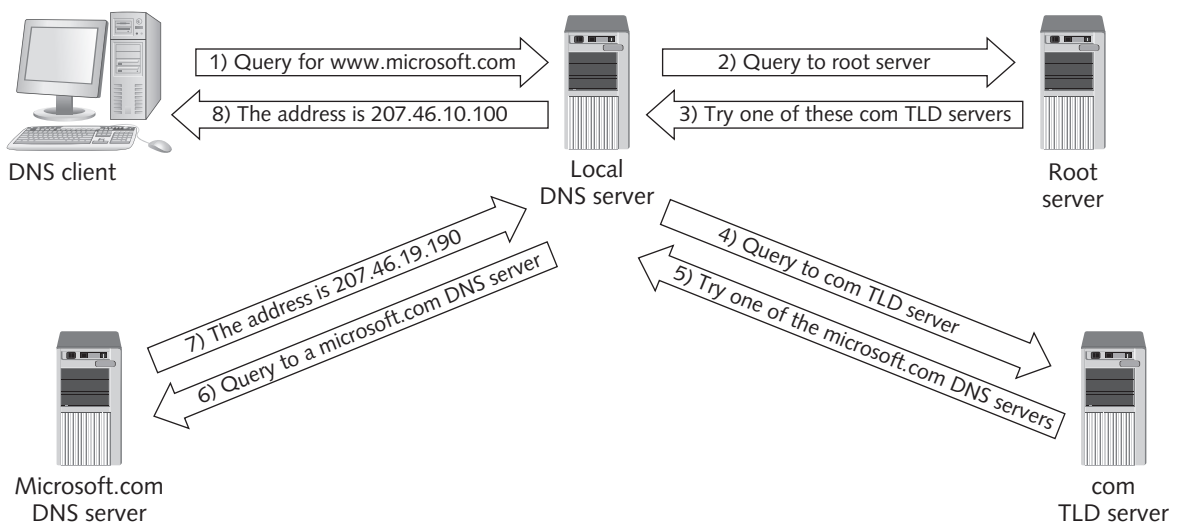


Figure 8-17 A DNS query making its way through the DNS hierarchy

Courtesy of Course Technology/Cengage Learning

- *DNS server service*—This service runs in the background and listens for DNS queries on UDP port 53.

DNS servers can also be configured with no zones at all—a configuration called a caching-only server. A home or small business network that doesn't maintain its own domain can still install and use a DNS server. After DNS is installed on the server, clients can be configured to use the server for DNS queries. Initially, the server has to query root servers for most requests, but because the results are cached, it can resolve queries for frequently visited Web sites. A caching-only DNS server has the advantage of reducing traffic to the ISP's DNS servers, and your local DNS server can often respond to queries faster, especially if your ISP's DNS servers are busy or down. Figure 8-18 shows the Windows DNS Manager console.

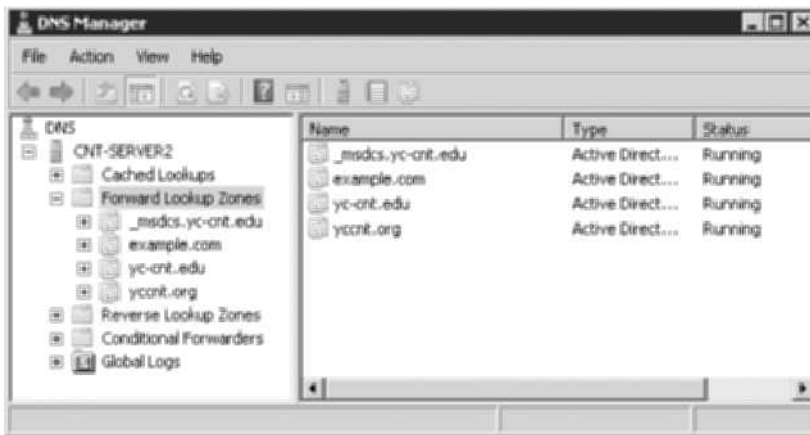


Figure 8-18 The DNS Manager console

Courtesy of Course Technology/Cengage Learning



Many Linux systems use a DNS system called Berkeley Internet Name Daemon (BIND), which has been around since the 1980s and is the original widely used DNS system.

Server and Network Fault Tolerance

By now, you know that a server is defined by the type of software installed on it. For example, Windows Server 2008 and most Linux distributions can be installed on an inexpensive laptop just as easily as they can on a \$20,000 server, but laptops are rarely adequate for an enterprise-level network. A network's servers can be critical to business operations, so keeping them running at peak performance is essential for user productivity and business transactions. For this reason, certain fault-tolerance features are built into server OSs, and only servers designed to use these features can access them. Some fault-tolerance features on a server OS that aren't usually available on desktop OS versions include the following:

- *Support for hot-swappable devices*—A **hot-swappable device** can be removed, replaced, or added to a server while it's running. Many low-end servers support hot-swappable

disk drives, but only high-end servers are likely to support hot-swappable memory and CPUs. All editions of Windows Server support hot-swappable disks, but you have to spring for the more expensive Enterprise or Datacenter edition to support hot-swappable memory and CPUs. Red Hat and other Linux distributions support hot-swappable devices (called “hotplug” in the Linux world), too.

- *Server clustering*—A **server cluster** is two or more servers configured to operate as a single unit. The most common types of server clusters are failover clusters and load-balancing clusters. A **failover cluster** is used to provide fault tolerance so that if one server fails, the other immediately takes over its functions with no or little downtime. A **load-balancing cluster** provides high-performance computing and data access by spreading the workload among multiple computers. Physically, there are multiple servers, but logically, they work as one unit. Load-balancing clusters have the added advantage that if one server fails, the others still operate, which ensures fault tolerance.
- *Redundant/high-end disk systems*—Hard drives are a critical component of a computer and one of the few moving parts in a computer, making them more susceptible to failure than other components. Most high-end servers support enterprise-class Small Computer System Interface (SCSI) or Serial Attached SCSI (SAS) disks designed for a 24/7 duty cycle. Low-end servers and desktop computers support only Serial ATA (SATA) disks, which lack some of the performance properties of SCSI and SAS. However, even high-end disk drives can fail, so most servers incorporate disk controllers capable of a disk arrangement known as **redundant array of independent disks (RAID)**. With RAID, you can configure two or more disks in a fault-tolerant arrangement so that if one disk fails, the data is preserved and the server can continue to operate. Even some desktop computers support variations of RAID, but the variations with higher performance and fault tolerance are standard on servers. RAID is discussed more in Chapter 9.



Additional Server Features

As mentioned, OS vendors reserve many high-end applications and network services for the server version of the OS. Common applications and services usually found only on servers includes the following:

- *Remote access*—Today’s mobile workforce needs convenient access to the corporate network from anywhere in the world. Most server OSs support virtual private networks (VPNs) and, if necessary, the older dial-up method of remote access. VPNs are discussed more in Chapter 10.
- *Database server*—Many applications rely on a database to store and retrieve vast amounts of data. Server OSs support advanced database systems, such as MySQL, SQL Server, and Oracle.
- *Client/server applications*—Client/server applications, such as corporate e-mail systems (Microsoft Exchange and Lotus Notes, for example), must run on a server OS. Web-based applications, too, need a server OS to handle the computing and network workload of these applications.
- *Virtualization*—Virtualization is an integral part of most IT data centers. Virtualization software runs on desktop systems, but for virtualizing production servers, you need a server-based product, such as Windows Hyper-V or VMware vSphere. For

open-source fans, Citrix XenServer might be a good fit with your data center. Virtualization is such an important aspect of the computing environment that the next section focuses on this topic.

The list of applications and services usually reserved for servers and server OSs is continually growing as networks play an increasingly important role in personal and work activities. For now, turn your attention to OS virtualization, one of the hottest topics in computing.

Operating System Virtualization

OS virtualization has become a mainstream technology in both small and large networks. **Virtualization** is a process that creates a software environment to emulate a computer's hardware and BIOS, allowing multiple OSs to run on the same physical computer at the same time. This environment can be installed on most current OSs, from Windows XP to Linux to MAC OS. In this case, a picture is worth a thousand words, so examine Figure 8-19. It shows a Windows 7 desktop running a Windows Server 2008 virtual machine, using VMware Workstation virtualization software.

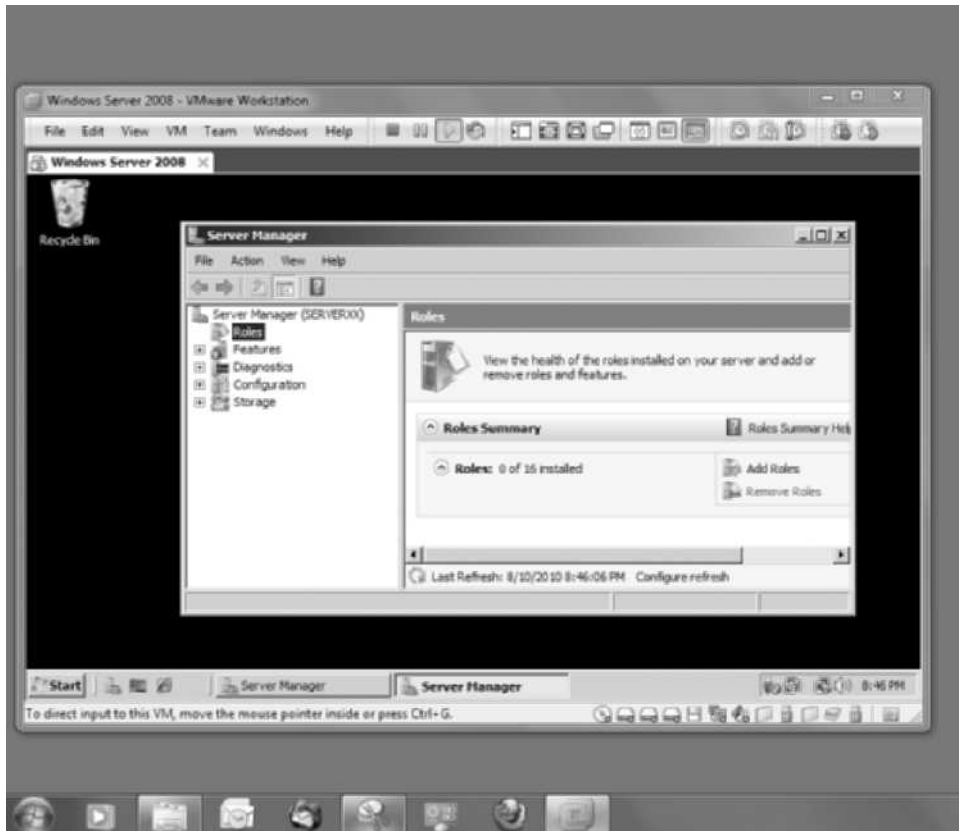


Figure 8-19 Windows Server 2008 running as a virtual machine on Windows 7

Courtesy of Course Technology/Cengage Learning

Like all technologies, virtualization has a collection of terms that define its operation and components:

- A **virtual machine (VM)** is the virtual environment that emulates a physical computer's hardware and BIOS. A **guest OS** is the operating system installed on a VM.
- A **host computer** is the physical computer on which VM software is installed and VMs run.
- Virtualization software is the software for creating and managing VMs and creating the virtual environment in which a guest OS is installed.
- The **hypervisor** is the virtualization software component that creates and monitors the virtual hardware environment, which allows multiple VMs to share physical hardware resources. (In some software, this component is called Virtual Machine Monitor [VMM].) The hypervisor on a host computer acts in some ways like an OS kernel, but instead of scheduling processes for access to the CPU and other devices, it schedules VMs.
- A type 1 hypervisor implements OS virtualization by running directly on the host computer's hardware and controls and monitors guest OSs. It also controls access to the host's hardware and provides device drivers for guest OSs. Also called **bare-metal virtualization**, it's used primarily for server virtualization in data centers.
- A type 2 hypervisor implements OS virtualization by being installed in a general-purpose host OS, such as Windows 7 or Linux, and the host OS accesses host hardware on behalf of the guest OS. Also called **hosted virtualization**, it's used primarily for desktop virtualization solutions.
- A virtual disk consists of files residing on the host computer that represent a virtual machine's hard drive.
- A virtual network is a network configuration created by virtualization software and used by virtual machines for network communication.
- A **snapshot** is a partial copy of a VM made at a particular moment; it contains changes made since the VM was created and is used to restore the VM to its state when the snapshot was taken.

One of the best ways to understand a technology is to understand the reasons it's used. The reasons to use virtualization are many and varied and are best discussed by splitting the topic into the two main types of virtualization: hosted and bare-metal.

Hosted Virtualization

As mentioned, hosted virtualization uses a type 2 hypervisor, which is installed in a standard desktop or server OS. It has the advantage of supporting a wider variety of guest OSs than bare-metal virtualization does, mostly because the guest OS uses the host OS to access host hardware, so there are few incompatibility problems between the guest OS and hardware. For example, you can run a distribution of Linux in a virtual machine on a host computer, even if you can't install Linux directly because of driver incompatibilities.

Another advantage of hosted virtualization is that it's easy and straightforward to use and is largely independent of companies such as Microsoft, VMware, and the open-source community. Open source means software in which the source programming code is available free to



the public. The executable program is often, but not always, distributed free, too. With hosted virtualization, you install the virtualization software on your computer and begin creating virtual machines. There are few hardware requirements, and most products run on Windows versions starting with Windows XP as well as most Linux distributions. All that's required are enough memory to support the host and guest OSs, adequate CPU power, and enough free disk space to store the virtual disk. A laptop running Windows 7 with 2 GB RAM, a 2.0 GHz CPU, and 20 GB free hard drive space can handily run Linux and Windows Server 2008 virtual machine at the same time. Performance might not be stellar, but the virtual machines should work well enough for experimenting or training (one of the main reasons for using hosted virtualization).

Hosted Virtualization Applications Hosted virtualization is so flexible and easy to use that its uses are varied and continuing to grow as people find different applications for it. Some common applications include the following:

- *OS training*—Whether in the classroom or at home, learning multiple OSs has often been a problem of not having enough computers or a lack of compatibility between the OS and available computers. With virtualization, a computer can have a host OS installed, such as Windows 7, and have virtual machines for numerous Linux distributions, Windows Vista, Windows Server 2008, even Novell NetWare. If you want to learn about the past, you can install Windows 3.11, DOS, or OS/2. In addition, you can run multiple VMs at the same time by using a virtual network, which enables you to work with both client and server OSs in situations that would normally take two or more physical computers.
- *Software training*—Students and employees can be trained on new software packages by giving them VMs with preinstalled software.
- *Application isolation*—Not all software plays well together, so if an application conflicts with other installed software, it can be installed in its own VM, effectively isolating it from the host machine's installed software.
- *Network isolation*—Installing some networking services, such as DHCP, can wreak havoc with an existing network. Virtual networks can be isolated from the rest of the network, however, so you can experiment with these services without causing the IT department to pay you a visit.
- *Software development*—Software developers often need to design software that works on multiple OSs and OS versions. Testing on VMs makes this process easier, compared with using a physical computer for each OS to be tested.
- *What-if scenarios*—If you want to try out a software package or see whether a configuration option you read about will actually improve performance on your computer, you might not want to risk destabilizing your physical computer. You can install software and make configuration changes safely on a VM before making the commitment on your host computer.
- *Use of legacy applications*—If you have a favorite application that won't run on a newer OS, you don't have to forgo the latest technology because of one application. You can install the old OS in a VM and run your legacy application on it.
- *Physical-to-virtual conversion*—Your six-year-old machine is getting slow and unreliable, so you bought a new desktop computer. However, you have several

applications on your old computer and no longer have the installation media. You can convert your old computer to a virtual machine, and then maintain all the software and run it on your new desktop computer as a VM. You'll probably even see a speed boost.

As you can see, virtualization can bring plenty of benefits to your computing experience. You have many choices of products in this category, and the good news is that most are free. The following section describes some products for hosted virtualization.

Hosted Virtualization Products There are a number of players in the hosted virtualization arena. The following lists the most well-known hosted virtualization products:

- *VMware Workstation*—VMware, the virtualization pioneer in the PC world, released VMware Workstation in 1999. It's the only product in this list that isn't free, but it offers the most features, including multiple snapshots, built-in physical-to-virtual conversion, and extensive guest OS support.
- *VMware Player*—This free download from VMware has a streamlined user interface and fewer advanced features than Workstation, but maintains excellent guest OS support.
- *Microsoft Virtual PC*—This free download from Microsoft enables you to create and run virtual machines. XP mode is a new feature available only in Windows 7 Professional or Ultimate that integrates Windows XP virtual machine applications with the Windows 7 desktop.
- *VirtualBox*—Originally developed by Innotek, it's now developed by Oracle Corporation. Two versions are available: a proprietary version that's free for home users and can be purchased for enterprise use and a free open-source version with a reduced feature set. VirtualBox runs on Linux, Mac OS X, or Windows hosts, and the proprietary version has features similar to VMware Workstation.



TIP

For more information on virtualization products, the platforms they run on, and supported guest OSs, review the article at http://en.wikipedia.org/wiki/Comparison_of_platform_virtual_machines.

These products have their strengths and weaknesses; the best approach is to work with different products to see which best serves your needs. The following sections discuss using these products.

Using VMware Workstation VMware Workstation isn't free, but you can download a trial version at no cost and try it for 60 days. Not-for-profit educational institutions can join the VMware Academic program to give students and faculty free downloads of VMware Workstation and other VMware products.

After VMware Workstation is installed, a wizard takes you through the steps of creating a virtual machine. You can choose the size for the virtual disk and set other hardware options or just accept the defaults.





One convenience of installing a guest OS in a VM is being able to boot to the installation program with a DVD image file rather than a DVD disk. This way, if you download the DVD image file, burning a DVD to do the OS installation is unnecessary. In addition, the image file can be stored on a server and be used by multiple users for VM installations.

An advanced feature of VMware Workstation is flexible networking options. You can configure the NIC on your VM to use one of the three virtual network options or you can create your own custom virtual network. VMware Workstation supports VMs with multiple NICs, and each NIC can be connected to a different virtual network. The three preconfigured options are as follows (see Figure 8-20):

- *Bridged*—This option connects the VM directly to the physical network, and the VM acts like any other computer on the physical network.
- *NAT*—With this default option, the host computer's IP address is shared with the VM by using Network Address Translation (NAT). This option is more secure than bridged because the VM isn't directly accessible. However, it's not a viable option for a VM providing server functions to the host network.

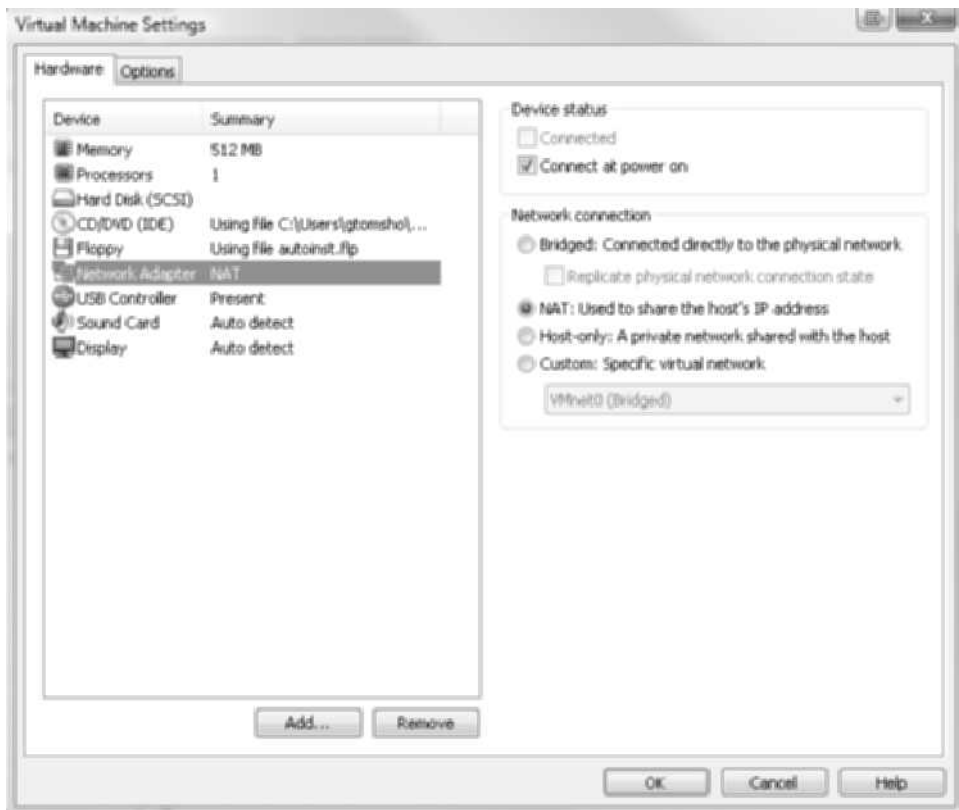


Figure 8-20 VMware virtual network options

Courtesy of Course Technology/Cengage Learning

- *Host-only*—This option isolates the VM from the host network and allows network communication only between VMs running on the host and the host computer. It's the most secure configuration and has the lowest risk of the VM causing problems with the host network. This configuration works well when you have multiple VMs that must communicate with one another but don't need to access computers or devices outside the host.

After the virtual machine is installed, you use it as you would any computer, except there are no physical on/off buttons. Figure 8-21 shows Windows Server 2008 running in a VM. Most controls for the VM are accessed with the icons below the menu bar.

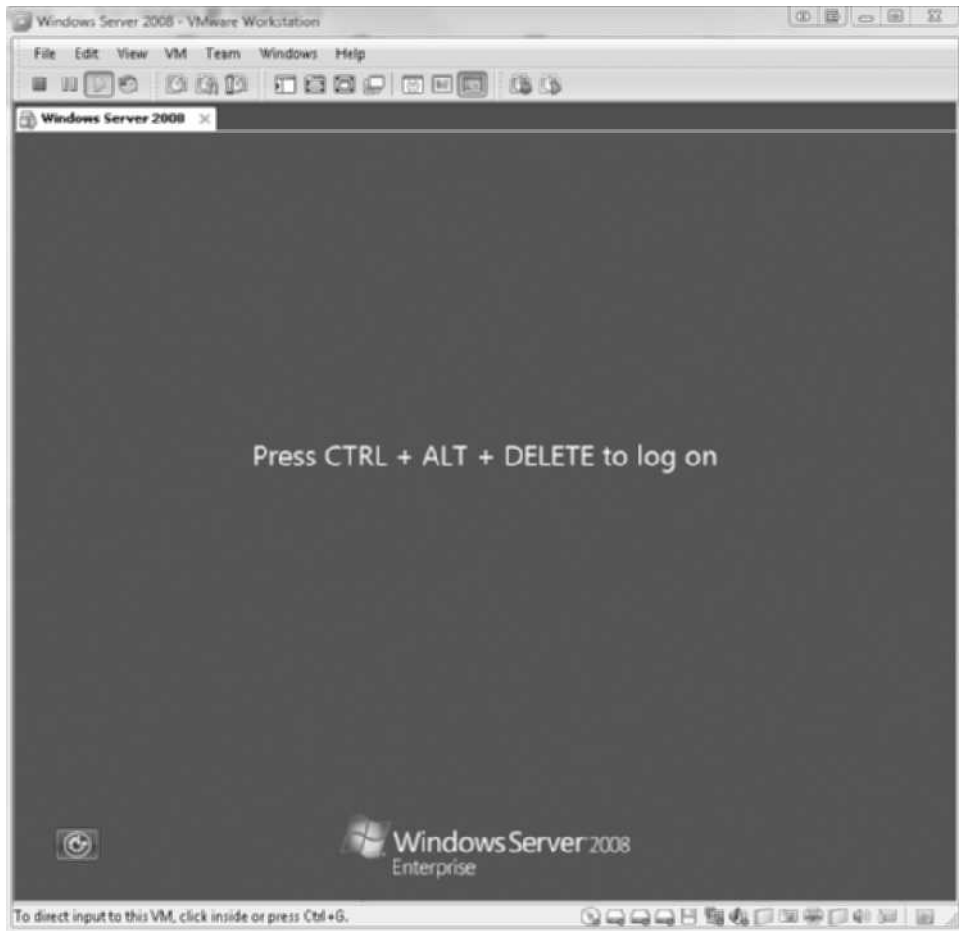


Figure 8-21 Windows Server 2008 in VMware Workstation

Courtesy of Course Technology/Cengage Learning

Another advanced feature is the concept of a virtual team, which is two or more VMs grouped together. For example, if you're testing a network setup involving several VMs, you can use a virtual team to open and start all the VMs at once instead of opening and starting them separately. A convenient interface lets you switch from one VM to another in the team easily. You can also define a virtual network and connect the NICs of all VMs in the team to the virtual network.

A set of tools and drivers called VMware Tools should be installed in the guest OS for the best performance and ease of use. VMware Tools adds optimized network, video, and disk drivers and guest-host integration tools that allow dragging and dropping files and cut and paste between the guest OS and host.

Other advanced features targeted to developers are available, which is why VMware Workstation is generally considered the flagship hosted virtualization product. However, if you don't need all the bells and whistles and simpler is better, try VMware Player.

Using VMware Player VMware Player is a stripped-down version of VMware Workstation but still offers the basics of desktop virtualization in a streamlined and easy-to-use interface. You can download it free from the VMware Web site, and it's also included with the VMware Workstation package. The opening window of VMware Player gives you an idea of its clean interface (see Figure 8-22).



Figure 8-22 The VMware Player Welcome window

Courtesy of Course Technology/Cengage Learning

Creating a new VM in VMware Player is a wizard-based affair, nearly identical to the one in VMware Workstation. Notice in Figure 8-22 the option to download a virtual appliance. Virtual appliances are ready-to-use VMs from OS and software vendors that contain a guest OS

with preconfigured applications or network services. In some cases, the virtual appliance is just a preinstalled guest OS. A virtual appliance is an easy way to use and evaluate a product or configuration without having to deal with installing it yourself. Virtual appliances can be run by VMware Player or Workstation and sometimes by VMware's bare-metal virtualization products.

VMware Player offers many of the same features as VMware Workstation, with the exception of snapshots, virtual teams, and customized virtual networks (although the three preconfigured network options are available). It's a good choice for new virtualization users and for classroom and training centers where the interface's simplicity is an advantage.

Using Microsoft Virtual PC The VMware products just discussed can be installed on both Windows and Linux OSs (another desktop product called VMware Fusions runs in Mac OS), but Virtual PC is a Windows-only product. The latest version, called Windows Virtual PC, runs only on Windows 7 Professional, Enterprise, or Ultimate; for older Windows OSs, you need to install Virtual PC 2007. This section focuses on Windows Virtual PC because it adds some much-needed enhancements, such as support for USB devices, that other virtualization products have long supported.



Virtual PC 2007 is discussed in Appendix E along with several other virtualization products. Appendix E includes reference information on several virtualization products and gives step-by-step instructions for using some of them. Be aware that products are revised frequently, so screenshots and instructions might vary.

The Virtual PC interface has two consoles. When you first start it, you see a Windows Explorer interface with a menu bar customized for Virtual PC (see Figure 8-23). From here, you can start a VM by double-clicking its configuration file, or you can create a new VM by clicking the Create virtual machine link at the top. This console also shows the VM's current status, such as Running or Powered down.

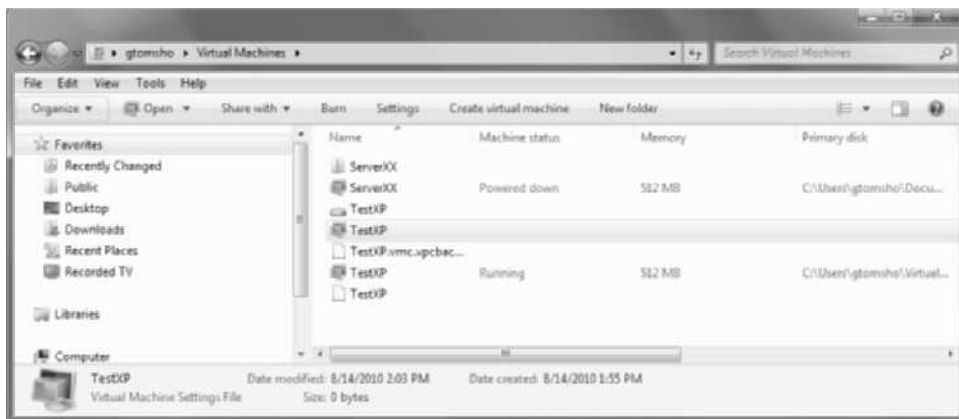


Figure 8-23 Virtual PC has a Windows Explorer interface for working with VMs

Courtesy of Course Technology/Cengage Learning

The user interface in Virtual PC is streamlined (much like VMware Player's) because Virtual PC has few advanced features (see Figure 8-24). Virtual PC does support a type of snapshot called an undo disk. You must enable the undo disk feature when creating the VM for this feature to be available.



Figure 8-24 A Windows XP VM running in Microsoft Virtual PC

Courtesy of Course Technology/Cengage Learning

An add-on feature for Virtual PC is XP mode, which you download separately. It removes some complexities of working with a VM, such as having two desktops and Start menus. With XP mode, a Windows XP VM is installed automatically for you. You can run Windows XP applications in a VM on a Windows 7 host without compatibility problems. The applications appear as shortcuts on your Windows 7 Start menu, which removes the sometimes confusing aspects of running a VM with its own desktop and Start menu. VMware Player and Workstation have a similar feature called Unity.

Like VMware, Virtual PC has guest-host integration tools that should be installed after the guest OS is installed. Integration services must be installed for USB support on the guest OS. Virtual PC guest OS support is officially limited to Windows OSs, but some users have had success installing Linux VMs.

Using VirtualBox VirtualBox can be installed on Windows, Mac OS X, Linux, and Solaris hosts and supports a wide range of Windows, Linux, and other guest OSs, making

it the most versatile of the products discussed. Like the other products, virtual machines are created with a wizard that walks you through selecting the guest OS and the VM's hard disk and RAM configuration; however, you can change all these settings after the VM is created. The VirtualBox user interface consists of a console where you can create VMs and view the status of all VMs. VirtualBox supports unlimited snapshots, so you can save a VM's state as you work with it and restore its state from any of the snapshots you make. You can even jump forward and backward in snapshots, meaning that if you have three snapshots, you could revert to the first snapshot and later go back to the third snapshot. Figure 8-25 shows the VirtualBox console with one running VM and a list of available snapshots.

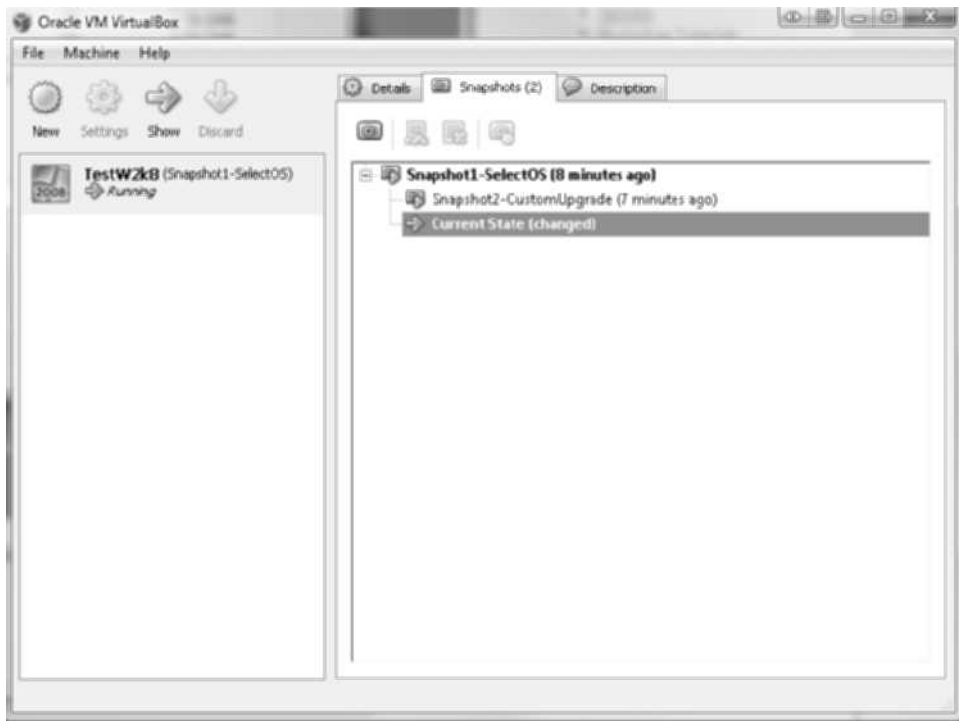


Figure 8-25 The VirtualBox console showing available snapshots for a VM

Courtesy of Course Technology/Cengage Learning

A VirtualBox feature called seamless mode is similar to Virtual PC's XP mode and VMware's Unity. With this feature, a second taskbar and Start button are created on your host desktop, allowing you to access your guest OS's applications and features without the distraction of a second desktop.

Virtualization Software Summary All the virtualization products discussed provide a type 2 hypervisor (hosted virtualization). Table 8-1 summarizes some major features and differences in these products. Flash Movie 1: Using VMware Workstation (available on this book's CD) shows creating a virtual machine in VMware Workstation and using this product's features.



Flash Movie 1: Using VMware Workstation

Table 8-1 Comparing features of hosted virtualization software

	VMware Workstation	VMware Player	Windows Virtual PC	VirtualBox
Price	\$189 or free with Academic Program membership	Free	Free	Free
Host OS support	Windows, Linux, Mac OS X (with VMware Fusion)	Windows, Linux	Windows	Windows, Linux, Mac OS X, Solaris
Guest OS support	Windows, several Linux distributions, NetWare, Solaris, DOS	Same as Workstation	Windows XP and later	Windows, several Linux distributions, Solaris, Mac OS X Server, DOS, OS/2, others
Snapshots	Unlimited	None	One (with Disk Undo enabled)	Unlimited
Virtual network options	Bridged, NAT, host-only, custom	Bridged, NAT, host-only	Bridged, NAT, internal (guest-to-guest only)	Bridged, NAT, host-only, internal
Host integration tools	VMware Tools, Unity	VMware Tools, Unity	Integration Services, XP mode	Guest additions, seamless mode
Other features	Virtual teams, screen capture and screen movie capture, physical-to-VM conversion, developer tools			Command-line management interface, built-in remote desktop, developer programming interface, open-source edition

A benefit of these virtualization products is that you can install all of them and run them at the same time on a single host computer, so you can download and install each one and evaluate it for yourself.

Bare-Metal Virtualization

Bare-metal virtualization products (type 1 hypervisors) are targeted mainly for production virtualization in data centers. These products are installed directly on hardware and have more stringent host machine requirements than hosted products do. Because they're targeted for IT departments, they have more features for managing VMs and have a performance advantage over hosted virtualization products. Their installation and use tend to require more sophisticated, knowledgeable users, too. Before learning about specific products, take a look at some applications for bare-metal virtualization products in the next section.

Bare-Metal Virtualization Applications Bare-metal virtualization products come with a price tag for the virtualization software, the hardware to run it on, or both. So

when considering whether to use virtualization in an IT data center, most IT managers look for a return on their investment in real money or in productivity gains. The following applications show that bare-metal virtualization can deliver both:

- *Consolidate servers*—Server consolidation is probably the original reason for using bare-metal virtualization and is done for the following reasons and benefits:
 - *Retire old or unreliable hardware:* Converting physical machines to VMs and running them on the latest hardware means you can get rid of old hardware, thereby gaining a reliability advantage and avoiding the tedious task of reinstalling and reconfiguring a server OS on new hardware. You might also improve performance.
 - *Make optimal use of multicore, high-performance servers:* Some server roles, such as Active Directory, should be the only major network service running on a server. With multicore server CPUs, you're likely to waste a lot of the server's power if you install a single-role OS. Instead, run two, three, or more VMs on the server, making optimal use of the available performance.
 - *Maintain application separation:* Some applications and services run best when they're the only major application installed on an OS. You avoid OS resource conflicts and gain stability and reliability.
 - *Reclaim rack or floor space:* By consolidating a dozen physical servers into three or four host servers, you're no longer tripping over a plethora of towers or wondering whether your rack can handle one more server. You can even clear enough room for an easy chair and a reading lamp so that you can catch up on the latest technical journals in comfort!
 - *Reduce cooling and power requirements:* In most cases, by reducing the number of servers (even with higher performance machines), you save money on cooling and powering a data center, especially when you reduce hundreds of servers down to dozens of hosts.
- *Test installations and upgrades*—Before you install a major software package or upgrade on your server, create a copy of the VM (referred to as “cloning” in some products), and go through a test run to iron out any potential problems or conflicts. If something still goes wrong on the production VM, you can revert to a snapshot.
- *Test a preconfigured application*—Not sure whether the application the vendor is trying to sell you is right for your company? Some vendors offer virtual appliances you can use to evaluate the application without the trouble of installing it.
- *Test what-if scenarios*—You can create a virtual network and run clones of your production VMs to test ideas for improving your network's performance, functionality, and reliability. However, this type of testing on live production systems is never a good idea.
- *Live migration*—Virtual machines can be migrated to new hardware while they're running for performance or reliability improvements with practically no downtime.
- *Dynamic provisioning*—Advanced VM management systems can deploy VMs and storage dynamically to meet application requirements. This advanced feature has uses in clustered computing and cloud computing (discussed in Chapter 12).





VMs that run distributed server applications, such as Active Directory, in which multiple servers synchronize a common database with one another, shouldn't be backed up or moved by copying the virtual hard disk, as it might result in database inconsistencies. Use only backup and migration tools approved for the application.

Bare-Metal Virtualization Products VMware dominated the type 1 hypervisor category for years, but now you have a number of products to choose from. The following are the most common bare-metal virtualization products:

- *Microsoft Hyper-V*—Hyper-V was introduced with Windows Server 2008 and can be installed as a server role, in which case the hypervisor is installed as a layer of software between Windows Server 2008 and the server hardware. Windows Server 2008 acts as a parent or management OS for VMs installed with Hyper-V. Hyper-V can also be installed as a stand-alone product directly on the server, with only a command-line interface available for rudimentary management tasks; it's managed remotely by another Windows Server 2008 computer. Hyper-V is included with Windows Server 2008 Standard, Enterprise, and Datacenter editions at no additional cost, or you can download the stand-alone Hyper-V Server free from the Microsoft Web site. Hyper-V R2, which comes with Windows Server 2008 R2, supports advanced features such as host server clustering and live migration. Hyper-V requires a 64-bit CPU with virtualization extensions on the host system. Virtualization extensions offload some virtualization work to the CPU and are available on Intel-VT CPUs and AMD-V CPUs.

A big advantage of using Hyper-V is that Microsoft provides virtual instances of the OS with no additional licensing fees. For example, Windows Server 2008 Standard Edition allows you to run one virtual instance (or one VM) of the OS at no additional cost. Enterprise Edition allows up to four virtual instances, and Datacenter Edition allows an unlimited number of virtual instances. Hyper-V has guest OS support for Windows server OSs from Windows 2000, SUSE and Red Hat Enterprise Linux distributions, Windows client OSs (Windows XP and later), and more. Figure 8-26 shows Hyper-V Manager.

- *Citrix XenServer*—This open-source hypervisor uses Linux as a management OS on the host. It's available free or as a commercial edition that adds enterprise-level features, such as fault tolerance, performance management, and host power management. A number of modified Linux versions and Solaris can run as the management OS, and like Hyper-V, a XenServer host computer requires a 64-bit CPU with virtualization extensions to run Windows guest OSs. Guest OS support includes most Windows OSs starting with Windows XP and SUSE, Red Hat, and CentOS Linux distributions. To manage your host and VMs, you download and install XenCenter on a Windows computer.
- *VMware vSphere*—vSphere includes VMware ESX Server, which is installed directly on the physical server without a management OS. After ESX Server is installed, a basic command-line console based on Linux is available for simple configuration tasks, such as IP address configuration. Most configuration tasks are done from a remote client OS, using vSphere Client that's downloaded and installed on a Windows or Linux OS (see Figure 8-27). You can also create, manage, and access VMs via a Web-based interface in Internet Explorer or Firefox. A unique feature is being able to install

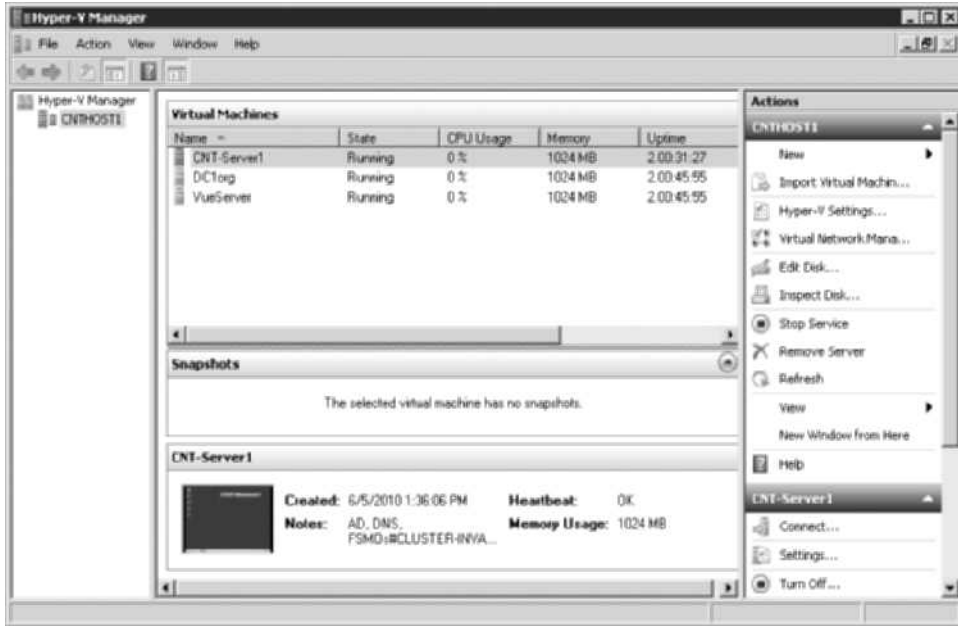


Figure 8-26 Microsoft Hyper-V Manager

Courtesy of Course Technology/Cengage Learning




Figure 8-27 VMware's vSphere Client software

Courtesy of Course Technology/Cengage Learning



ESX Server as a virtual machine in VMware Workstation or VMware Player, which gives you the opportunity to evaluate ESX Server without committing a physical server to its installation. You can even create VMs and install guest OSs on this ESX Server VM, which gives you a virtual machine running on a host that's a virtual machine. However, this option is intended only for evaluation, not for production purposes. ESX Server has the broadest guest OS support, including Windows versions back to Windows 3.1, more than eight Linux distributions, Novell NetWare, Solaris, and others.

All these products have extensive management tools for managing up to hundreds of hosts and a wide array of storage resources. These tools are available for a fee from virtualization software vendors. For example, Microsoft has System Center Virtual Machine Manager (SCVMM) for managing Hyper-V and ESX Server hosts. XenServer offers versions with different levels of management, depending on which product you purchase, and VMware sells vCenter Server to manage a vSphere virtualization environment. All these products are designed to provide a secure, reliable, and highly available virtualization infrastructure.

The basic tasks of creating and accessing VMs on bare-metal virtualization software are similar to using desktop products: a wizard that walks you through the procedures. The real differences lie in host and resource management and the capability to give IT managers the tools needed to virtualize a data center. This section serves as an introduction to the available products so that you have a starting point for doing your own research in the expanding field of virtualization.

Installing an OS

Installing an OS, whether it's a desktop or server version, has become the proverbial no-brainer. Essentially, OSs install themselves, and all you have to do is click Next and OK a few times and perhaps enter a license key and accept the license agreement. Even most Linux distributions, which in the past could stymie novices with a frustrating array of choices and options, are mostly hands-off installations now.

The real work of installing an OS, particularly on a network server, involves preinstallation and postinstallation tasks. The prerequisites for installing any OS are a copy of the installation medium and a computer that meets the installation requirements, including enough free (preferably unallocated) disk space, a CPU that meets minimum performance requirements, and enough RAM. The following sections explain the preinstallation planning process, the installation, and common postinstallation tasks for Windows Server 2008 and a common distribution of Linux.

Planning for and Installing Windows Server 2008

The role a server will play on the network is a key consideration in planning Windows Server 2008 installations. A server used only for file and printer sharing that supports a dozen users has different minimum hardware requirements than a server running Active Directory, a Web server, and a database and supporting a few hundred users. Windows Server 2008 is available in several editions with different costs and capabilities, so you need to determine which

edition best fits your needs. After Windows is installed, you need to perform some postinstallation tasks immediately before installing additional features or applications.

Selecting Server Hardware for Windows Server 2008 Table 1-4 in Chapter 1 lists the minimum hardware requirements for installing Windows Server 2008, but unless the server is to be used only for training or demonstration, the minimum requirements are rarely enough for a production server. So one of the main factors to consider for a server OS installation is the server's hardware features. The following list describes a few features you might need to decide on before purchasing a server:

- *CPU architecture*—The minimum requirement is a 1.4 GHz CPU. CPUs are available in speeds well over 3 GHz, and major CPU manufacturers typically have a workstation line and a server line of processors. Depending on the expected server workload, you must also consider how many physical processors and how many cores each processor should have. Although Windows Server 2008 can run on just about any CPU meeting the minimum requirements, a CPU designed for servers usually has other server-specific components on the motherboard, such as high-end disk controllers and memory slots.



Windows Server 2008 R2 is the first Microsoft OS that no longer supports a 32-bit CPU; you must use a 64-bit system for Windows Server 2008 R2.

- *Disk subsystem*—Before SATA drives, the only real choice of hard drives for servers was SCSI. Both specifications make performance improvements constantly, and now SAS is available. For entry-level or departmental servers, SATA is a good choice because it's inexpensive and offers excellent performance. For enterprise servers or servers accessed around the clock, SAS and the newest SCSI systems have better performance and reliability but are much more expensive than SATA. SCSI/SAS disks are generally designed for continuous use; SATA drives tend to be designed more for consumer use. Researching current technology and your network's needs before deciding is best. RAID configurations that provide fault tolerance are inexpensive and highly recommended, considering their usefulness in the event of a disk failure. Windows Server 2008 requires only 32 GB of free disk space, but you need additional space for data you store on the server. The OS should be installed on one disk (or RAID set), and at least one other disk (or RAID set) should be used for data and application storage.
- *Memory*—The minimum requirement is 512 MB RAM, and Windows Server 2008 runs quite capably with this amount, at least until you have more than a couple of users accessing the server. Server motherboards are typically equipped with more RAM slots than desktop systems are, and for good reason. After you start running database-driven Web applications, maintaining a few thousand users in Active Directory, or using virtualization on your servers, often you'll find that you need 8, 16, or 32 GB (or even more) RAM. Also, be aware that server memory usually costs more than desktop memory because it has features such as buffering and error correcting code (ECC) that make it more reliable.



This list covers just a few server hardware configurations you should consider before installing a server OS. The best advice is to forge a good relationship with a knowledgeable vendor you can consult when you need to make a purchase. This way, you can focus on managing your server, and your vendor can focus on keeping up with the latest hardware options.



TIP

To make sure your hardware selections are compatible with Windows Server 2008, check the Windows Server Catalog at www.windowsservercatalog.com.

Selecting the Right Windows Edition Windows Server 2008 comes in four main editions that target different types or sizes of customers. These editions can be summarized as follows:

- Standard Edition is suitable for most small to medium businesses that need a robust solution for file and printer sharing, centralized control over user accounts and network resources, and services found in most networks, such as Web services, DNS, and DHCP. Standard Edition supports up to four multicore 64-bit processors and up to 32 GB RAM, so you might think this edition is all you need. However, some advanced features, such as system clustering and support for more processors, are reserved for higher-end editions of Windows Server 2008.
- Enterprise Edition has all the features of Standard Edition but includes others that make this edition suitable for medium to large businesses that need high-availability network services. It supports up to eight processors and up to 2 TB RAM. Server clustering is the most notable feature that isn't available in Standard Edition. With this feature, network administrators can tie physical servers together logically to act as a single, high-performance, fault-tolerant machine. Enterprise Edition permits up to 16 cluster nodes. Another fault-tolerance feature is hot-add memory, which means RAM can be added to a server without shutting the system down, as long as the server hardware supports this feature.
- Datacenter Edition might be a good fit for organizations managing huge amounts of data, using virtualization on a large scale, consolidating servers, or running high-volume, transaction-heavy applications. It includes all the features of Enterprise Edition with support for 64 processors. In addition, it includes these fault-tolerance features: hot-replace memory, hot-add processors, and hot-replace processors. Datacenter Edition can't be purchased as individual licenses; it must be purchased through volume licensing agreements or from original equipment manufacturers (OEMs), preinstalled on server hardware.
- Windows Web Server 2008 is designed to operate as a single-purpose Web server running Internet Information Services (IIS) 7.0. It has hardware support similar to Standard Edition but has no virtualization support and can't be installed as a domain controller. It lacks many features of other editions, such as remote access, Active Directory, and Terminal Services. Windows Web Server 2008 is a cost-effective solution, however, when you want a full-featured Web server but don't need the advanced features of the other editions.

Windows Server 2008 Preinstallation Decisions When installing a new server in a network, you must make some decisions shortly after finishing the installation. Many of these configuration decisions should be made before you actually begin the installation so that you can dive right into postinstallation tasks. Some are fairly straightforward, but others take some thought and consultation. Here's a list of some decisions you need to make:

- What should you name the server? This decision is more important than it sounds. Every computer needs a name so that it can be identified on the network. A server name must be unique on the network and should include some description, such as its location or primary function. Server names should also be simple and easy to remember because users often access servers by name.
- Which network protocols and addresses should you use? By default, Windows installs both TCP/IPv4 and TCP/IPv6 in Windows Server 2008. You can't uninstall them, but you can disable them in a network connection's Properties dialog box. Disabling a protocol is recommended if you're not using it. TCP/IPv4 is still the predominant LAN protocol and probably will be for years. Previous Windows versions had the option of installing other protocols and services, such as IPX/SPX (NWLink) and client/server components for NetWare. Windows Server 2008 has no additional protocol or client options, so if they're important, you need to find a third-party solution or use Windows Server 2003 or earlier.
- How should I assign an IP address to the server? By default, Windows Server 2008 is configured to use DHCP, but a server should have a static IP address. Some server roles actually require assigning a static address. If you haven't devised your addressing scheme, now is the time to do that. Generally, servers use one of the first or last addresses in the address range, such as 192.168.1.10 or 192.168.1.200. Whatever you decide, be consistent so that when more servers are added, you can assign addresses easily.
- Setting the correct time zone isn't really a decision but a task you must complete because having the wrong time zone can cause all manner of problems, particularly in a domain environment. Certain functions in a domain network, such as user authentication, depend on client and server computers having their clocks well synchronized.
- Should I use the workgroup or domain model? The Windows domain model has a number of advantages in usability, manageability, and security. If you've invested in a Windows server OS, it makes sense to get the most out of it by using the domain model and installing Active Directory. With a small network of fewer than 10 users, however, the workgroup model is a viable option, particularly if the main administrator isn't familiar with Active Directory. With either model, you need a workgroup or domain name, unless you're using the workgroup model and keep the default name "Workgroup." If you're using the domain model, you need to decide whether the domain name will be registered on the Internet. If it isn't, many Active Directory administrators use the top-level domain name "local," such as mycompany.local.
- What services should you install? This decision is one of the most important because it determines how the server will be used and what network services will be available to users. Windows Server 2008 refers to services such as Active Directory, DNS, and DHCP as "server roles." With the domain model, you must install Active Directory on at least one server. Active Directory requires DNS, so the DNS Server role is installed



automatically. Other basic roles to consider on a first server include DHCP (for IP address configuration) and File Services and Print Services, which include tools for sharing and managing file storage and printer resources. Many other roles and features can be installed to meet your network and business needs.

After you have a plan, it's time to move on to the actual installation of Windows Server 2008. Instead of reviewing installation steps here, you can run Flash Movie 2: Installing Windows Server 2008 (available on this book's CD) to see these steps in action.



Flash Movie 2: Installing Windows Server 2008

Windows Server 2008 Postinstallation Tasks Now that Windows Server 2008 is installed, it's time to attend to some postinstallation tasks. Some were discussed earlier, such as naming the server and configuring protocols and addresses. Here's a summary of the tasks you should perform immediately after installation:

- Activate Windows Server 2008.
- Set the correct date, time, and time zone.
- Assign a static IP address.
- Assign a computer name.
- Configure automatic updates.
- Download and install available updates.
- Add and configure roles and features.

Except for activating Windows Server 2008, all these tasks are in the Initial Configuration Tasks applet that opens each time you log on as Administrator (unless you click the “Do not show this window at logon” check box at the lower left). Windows Server 2008 requires activation within 60 days after installation. After 60 days, you can't log on until you do so. Windows Server 2008 activates automatically after several days, or you can activate it manually by clicking “Activate Windows now” in the System Properties dialog box. As you can see, most of the work of installing Windows Server 2008 is in the planning and postinstallation tasks. The same is true of most Linux installations, covered next.

Planning for and Installing Linux

Planning for a Linux server installation isn't much different from a Windows Server 2008 installation. Minimum hardware requirements must be met, and more important, hardware requirements for the role the server will play in your network must be met. Linux has come a long way in hardware compatibility but still doesn't have the broad support for different hardware that Windows does.



TIP

To research hardware compatibility for Linux distributions, go to www.linux-drivers.org.

One of the biggest decisions to make before you install Linux is which distribution to use. There are so many distributions, each with its own target audience, that making a recommendation without knowing the intended environment is impossible. A Web site called DistroWatch.com lists dozens of distributions along with descriptions and links to get more information. Most Linux distributions are open source and governed by the GNU General Public License (GPL), which allows users to run the program for any purpose, make changes to the program, and redistribute the program to others under the same GPL license terms.

After deciding on a Linux distribution, the next step is downloading a disk image of the installation medium and burning it to a CD or DVD. Many Linux distributions are offered as a Live CD that you can use to boot your system (physical or virtual) from the CD/DVD and run the OS without having to install it on a hard drive. Running a Live CD isn't a replacement for installing the OS on a disk, but it's a good way to evaluate a distribution. In addition, many specialized Linux installations come on Live CD and contain disk and system repair utilities to help you fix a Linux or Windows installation.

**TIP**

You can find a list of Live CDs at www.livcdlist.com.

The preinstallation and postinstallation tasks for a Linux OS aren't very different from those for Windows Server 2008, except there's no need to activate Linux and most tasks, such as IP address assignment and time zone selection, are done during the Linux installation. For demonstration purposes, the procedure for installing CentOS version 5.4, a popular Linux distribution used for server deployment, is shown in Flash Movie 3: Installing CentOS 5.4 (available on the book's CD).

**SIMULATION**

Flash Movie 3: Installing CentOS 5.4

Although installing Linux isn't difficult, it requires more input and decision-making during installation, whereas almost all configuration decisions in Windows are made after installation. Linux is a popular server OS, particularly for running Web applications and applications that use server-based databases. Compared with Windows, it offers all the basic infrastructure services, such as DHCP and DNS, but lacks a comprehensive directory service, such as Active Directory. Also, although more Linux services can be managed in a GUI, Linux still tends to make heavy use of the command line, which can be a drawback for administrators who are more at home with a GUI. Most large network environments use a combination of Windows and Linux servers, placing them in roles where they excel.



Chapter Summary

- A computer's OS provides services that enable users and devices to interact with the computer and manage the computer's resources. These services include a file system, process and service management, and the kernel.
- File systems provide a method for storing, organizing, and managing access to files on a storage device, such as a hard drive. In addition, they provide an indexing system for fast file retrieval and permissions for securing access to files.
- A process is a program that's loaded into memory and run by the CPU. It can be an application a user interacts with or a program with no user interface that communicates with and provides services to other processes. The latter type of process is called a service.
- The kernel schedules processes to run, making sure high-priority processes are taken care of first; manages memory to ensure that two applications don't attempt to use the same memory space; and makes sure I/O devices are accessed by only one process at a time, in addition to other tasks.
- An NOS provides all the features of a non-networked OS plus services that provide a convenient method for sharing and accessing network resources. Computers running an NOS usually play the role of a client or a server. Clients usually access network resources, and servers provide shared resources.
- Client computers typically run a number of client software components including file and printer sharing, DNS, DHCP, and e-mail. Servers provide the server side of the client/server relationship with file and printer sharing servers, DNS servers, DHCP servers, and e-mail servers, among others. In addition, servers provide directory services and policy management. Additional features in most server OSs include fault tolerance, remote access, database access, and virtualization.
- Virtualization can be divided into two categories: hosted virtualization and bare-metal virtualization. Hosted virtualization products are installed on a desktop OS and include VMware Workstation, Virtual PC, and VirtualBox. Bare-metal virtualization is used in data centers, is installed on servers, and includes products such as Microsoft Hyper-V, VMware vSphere, and Citrix XenServer.
- The real work of installing an OS consists of preinstallation and postinstallation tasks. The prerequisites for installing any OS are a copy of the installation medium and a computer that meets the hardware requirements, including enough free (preferably unallocated) disk space, a CPU meeting minimum performance requirements, and enough RAM.
- Some features to look for in a server system include CPU architecture, disk subsystem, and amount of memory. You must also be sure to select the correct edition of the OS you're going to install.
- Preinstallation decisions include the server name, the protocols to use, the networking model (domain or workgroup) you should use, and the services should be installed. Postinstallation tasks include activating the OS if necessary, setting the correct date and time, configuring IP settings, configuring the computer name, installing updates, and installing roles and features.

Key Terms

authentication The process of identifying who has access to the network. The most common form of authentication is a logon with a username and password.

authorization The process of granting or denying an authenticated user's access to network resources.

bare-metal virtualization The hypervisor implements OS virtualization by running directly on the host computer's hardware and controls and monitors guest OSs. *See also* virtualization.

batch file A text file containing a list of commands you ordinarily type at the command prompt.

cloud storage A data storage method in which some or all of an organization's data is stored on servers located offsite and maintained by a storage hosting company.

context switching Occurs when the OS suspends one process and activates another process.

cooperative multitasking In this form of multitasking, the OS can't stop a process; when a process gets control of the CPU, it maintains control until it satisfies its computing needs and informs the OS that another process can be activated.

DNS zone A database of primarily hostname and IP address pairs that are related by membership in an Internet or a Windows domain.

exclusion A configuration option for the IP address scope; excludes specified IP addresses from the DHCP IP address scope. *See also* IP address scope.

failover cluster A server cluster configuration used for fault tolerance so that if one server fails, the other takes over its functions immediately, with no or little downtime.

file system The method by which an OS stores, organizes, and manages access to files on a storage device, such as a hard drive.

host computer The physical computer on which virtual machine software is installed and virtual machines run.

hosted virtualization The hypervisor implements OS virtualization by being installed in a general-purpose host OS, such as Windows 7 or Linux, and the host OS accesses host hardware on behalf of the guest OS. *See also* virtualization.

hot-swappable device A computer device that can be removed, replaced, or added to a server while it's running.

hypervisor The component of virtualization software that creates and monitors the virtual hardware environment, which allows multiple VMs to share physical hardware resources.

IP address scope A component of a DHCP server, it's a range of IP addresses the server leases to clients requesting an IP address.

load-balancing cluster A server cluster configuration that provides high-performance computing and data access by spreading the workload among multiple computers.

multiprocessing A feature of some OSs that allow two or more threads to be run concurrently by separate CPUs or CPU cores. *See also* thread.

multithreaded application An application that has two or more threads that can be scheduled separately for execution by the CPU. *See also* thread.



network appliance A device equipped with specialized software that performs a limited task, such as file sharing. Network appliances are often packaged without video interfaces, so you don't configure them with an attached keyboard and monitor.

network-attached storage (NAS) A dedicated server device designed solely for providing shared storage for network users.

preemptive multitasking A form of multitasking in which the OS controls which process gets access to the CPU and for how long.

process A program that's loaded into memory and run by the CPU. It can be an application a user interacts with or a program with no user interface that communicates with and provides services to other processes.

redirector An OS client component that intercepts resource requests and determines whether the resource is local or remote.

redundant array of independent disks (RAID) A storage configuration of two or more disks, usually in a fault-tolerant arrangement so that if one disk fails, data is preserved and the server can continue to operate.

reservation A configuration option for an IP address scope that ties an IP address to a MAC address. When a client requests an IP address from the DHCP server, if the client's MAC address matches an address specified by a reservation, the reserved IP address is leased to the client instead of getting it from the scope. *See also* IP address scope.

server cluster Two or more servers configured to operate as a single unit. The most common types of server clusters are failover clusters and load-balancing clusters.

service A process that runs in the background and provides services to other processes; for example, DNS client and server components are services.

snapshot A partial copy of a virtual machine made at a particular moment, used to restore the virtual machine to its state when the snapshot was taken. *See also* virtual machine (VM).

thread The smallest unit of software that can be scheduled to run.

time slicing The process by which a CPU's computing cycles are divided between more than one process.

virtual machine (VM) A software environment that emulates a physical computer's hardware and BIOS.

virtualization A process that creates a software environment to emulate a computer's hardware and BIOS, allowing multiple OSs to run on the same physical computer at the same time.

Review Questions

1. Which of the following is an objective of a file system? (Choose all that apply.)
 - a. Organize space on a drive.
 - b. Organize files hierarchically.
 - c. Schedule access to applications.
 - d. Secure access to files.
2. A cluster is composed of which of the following?
 - a. One or more 512-bit blocks

- b. Two or more 2K-byte sectors
 - c. One or more 512-byte sectors
 - d. One or more 2K-byte blocks
3. Large cluster sizes can result in which of the following on a disk drive?
- a. Faster performance for large files, more wasted space for small files
 - b. Faster performance for small files, less wasted space for small files
 - c. More fragmentation, faster performance for large files
 - d. Less fragmentation, more wasted space for large files
4. Which of these file systems includes file and folder permissions? (Choose all that apply.)
- a. FAT32
 - b. Ext3
 - c. NTFS
 - d. Ext2
5. What feature of a file system makes it possible to find a file based on keywords in it?
6. Which of the following is best described as a program loaded into memory that has no user interface but communicates with other programs?
- a. Process
 - b. Task
 - c. Service
 - d. Application
7. The DNS function is built into most applications. True or False?
8. Which best describes context switching?
- a. Dividing computing cycles equally among processes
 - b. The OS suspending the running process and activating another process
 - c. Cooperative multitasking
 - d. Changing from one OS to another in a virtual environment
9. The most common form of multitasking in current OSs is cooperative multitasking. True or False?
10. The OS component that schedules processes to run is the _____.
- a. File system
 - b. User interface
 - c. Memory manager
 - d. Kernel
11. Which best describes a thread?
- a. A process you can view in Task Manager



- b. The smallest schedulable unit of software
 - c. A process in a preemptive multitasking OS
 - d. A multiprocessing computer
12. Which answer shows the correct order of DHCP packets generated when a computer requests a new address lease?
- a. DHCPDiscover, DHCPOffer, DHCPRequest, DHCPACK
 - b. DHCPDiscover, DHCPRequest, DHCPOffer, DHCPACK
 - c. DHCPRequest, DHCPDiscover, DHCPOffer, DHCPACK
 - d. DHCPRequest, DHCPOffer, DHCPDiscover, DHCPACK
13. If you want a computer to query DNS by appending more than one domain name to the computer name, which of the following should you configure?
- a. Two or more DNS server addresses
 - b. Additional DNS port numbers
 - c. Additional DNS suffixes
 - d. Two or more primary DNS suffixes
14. Which port does your Web browser use to communicate securely with the Web server by using SSL?
- a. 80
 - b. 25
 - c. 110
 - d. 443
15. Why should you set the “Register this connection’s addresses in DNS” option?
- a. So that you can query multiple domains when looking up a computer name
 - b. So that the computer’s name and address are added to the DNS database automatically
 - c. So that the NIC receives a DNS address from DHCP
 - d. To disable Dynamic DNS for that computer
16. Which is the correct syntax for mapping drive letter W to a shared folder named Accounting on the Finance server?
- a. net use W: \\Finance\Accounting
 - b. net share W: \\Accounting\Finance
 - c. net use W: \\Accounting\Finance
 - d. net share W: \\Finance\Accounting
17. A text file containing a list of commands is called which of the following?
- a. Logon process file
 - b. Service file
 - c. Task file
 - d. Batch file

18. The default protocol Windows uses to share folders is which of the following?
 - a. NFS
 - b. SMB
 - c. WPA
 - d. FTP
19. Which of the following refers to a Windows server with Active Directory installed?
 - a. Member server
 - b. NIS server
 - c. Domain controller
 - d. LDAP controller
20. Which of the following best describes an NAS?
 - a. A dedicated device designed to provide shared storage for network users
 - b. A high-speed network storage solution that can replace locally attached drives on servers
 - c. A storage solution in which some or all data is stored on offsite servers
 - d. A SATA or SCSI drive connected to a server
21. Which of the following is an element of a DNS server? (Choose all that apply.)
 - a. Zone
 - b. Root hints
 - c. Scope
 - d. Reservations
22. If you want one server to take over the processing of a server that has failed, what should you configure?
 - a. RAID
 - b. Failover cluster
 - c. Redundant array of servers
 - d. Load-balancing disk system
23. Software that creates and monitors the virtual hardware environment is called what?
 - a. Host computer
 - b. Hypervisor
 - c. Snapshot
 - d. Guest OS
24. Bare-metal virtualization is best for desktop virtualization. True or False?
25. If you want your virtual machine to have direct access to the physical network, which virtual network option should you configure?
 - a. Bridged



- b. NAT
- c. Host-only
- d. Internal

Challenge Labs



Challenge Lab 8-1: Capturing DHCP Packets

Time Required: 40 minutes

Objective: Capture the four packets generated by a DHCP client lease request.

Required Tools/Equipment: Your classroom computer with an IP address assigned via DHCP and Wireshark installed

Description: In this challenge lab, you use Wireshark to capture DHCP packets. Make sure you create a valid capture filter based on the Transport-layer protocol and port number DHCP uses so that only DHCP packets are captured. To generate DHCP packets, you need to force your computer to release and then renew its DHCP-assigned address. For instructions on doing this, review the Hands-On Projects in Chapter 5. Next, force your computer to generate packets used to renew an existing DHCP address. Answer the following questions:

- What filter did you use to capture only DHCP packets?

- What commands did you use to create the DHCP request packets:

- List the packets created when your computer made the DHCP request; you can find packet descriptions in the Info column of Wireshark.

- What type of packets were created in the DHCP request: unicast, multicast, or broadcast?

- Capture the packets involved in the DHCP renewal of the address lease. What packets were created (based on the description in the Info column of Wireshark), and were they unicast, multicast, or broadcast?



Challenge Lab 8-2: Installing and Using VMware Player with Ubuntu Linux

Time Required: 1 hour or longer, depending on download and installation time

Objective: Install VMware Player and the latest Ubuntu Linux distribution.

Required Tools/Equipment: Your classroom computer with Windows installed; you need a physical computer to install virtualization software, so a VM can't be used for this lab. The computer must have access to the Internet, or the instructor can download VMware Player and the Ubuntu Linux distribution and make them available to students. (Windows 7 is used to describe settings, but other Windows versions can be used.)

Description: Download and install VMware Player (www.vmware.com/products/player/), and then download and install the latest Linux Ubuntu distribution (www.ubuntu.com). Install VMware Player first, and then create a VM for installing Ubuntu Linux. Start the installation wizard, and in the Guest Operating System Installation window, select the "Installer disc image file (iso)" option and browse to the Ubuntu .iso file. After the installation is finished, answer the following questions:

- What advantages do you see in using an .iso file instead of a physical DVD to install an OS?

- List three examples of using hosted virtualization in a business environment:

- What disadvantages do you see with using hosted virtualization?



Case Projects



Case Project 8-1

From this book, you have learned a little about the different file systems in Windows and Linux. To get a better handle on the differences between them, write a short memo describing three properties of each of these file systems: FAT16, FAT32, NTFS, and EXT2. Relate the properties to the objectives of all file systems discussed in the chapter.

Case Project 8-2

You need to set up a network that meets the following requirements:

- Automatic IP address configuration

- Name resolution
- Centralized account management
- Ability to store files in a centralized location easily

Write a memo explaining what services must be installed on the network to satisfy each requirement.

Case Project 8-3

Your boss wants to purchase a new graphics design application to be distributed to approximately 40 users in the company. The problem is that although the company says it has broad OS support, he wants to be sure it will run on the five different OSs running on the company's user stations. He wants you to verify compatibility by using evaluation copies of the software without disrupting users or their computers. You have the installation disks for all five OSs your company uses, but you don't have a lot of computers available to install them on. What's your plan?

Case Project 8-4

You've been called in to recommend a server for a company that's opening a new office. You're meeting with the operations manager to get preliminary information about what the company needs and will make a recommendation for this new server's hardware and OS. List the top five questions you should ask the operations manager so that you can make the best recommendation.

Server Management and Administration

After reading this chapter and completing the exercises, you will be able to:

- Create and work with user and group accounts
- Create and manage permissions on storage volumes
- Work with shared files and printers
- Monitor a system's performance and reliability
- Describe fault tolerant and backup solutions

You've learned the basics of networks and network operating systems; now it's time to turn your attention to using an OS to perform typical management and administrative tasks. This chapter discusses user and group management, storage and file system management, and working with shared files and printers. In addition, as a server administrator, you need to know how to monitor system performance and prevent loss of data with fault-tolerance methods and backups.

Managing User and Group Accounts

Working with user accounts is one of an administrator's key tasks. User accounts are the link between real people and network resources, so user account management requires both technical expertise and people skills. When users can't log on to the network or access the resources they need, IT staff members get the phone calls. Your understanding of how user accounts work and how to configure them along with group accounts can reduce the frequency of these phone calls. User accounts have two main functions in a network:

- *Provide a method for users to authenticate themselves to the network*—Using a username and password is the most common way for users to log on to a network to gain access to network resources. User accounts can also contain restrictions about when and where a user can log on. Administrators use user accounts to assign permissions to network resources and define the types of actions a user can perform (referred to as **rights** in Windows), such as creating file shares or installing software.
- *Provide detailed information about a user*—User accounts can hold information such as a user's phone number, office location, department, and so forth for use in a company directory or for use by the IT department to identify users for support purposes.

Group accounts are used to organize users so that assignment of resource permissions and rights can be managed more easily than working with dozens or hundreds of individual user accounts. For example, an administrator can make a group account for each department in the company and add the users who work in each department as members of the corresponding group. Then, when a shared folder containing documents used by a certain department is created, the administrator just needs to assign permission to the group, which gives all members of that group the necessary permission. If a user changes departments, the administrator moves the account from one group to another, thereby changing the resources to which the user has permissions.

Account and Password Conventions

In a small network with only a few users and network resources, establishing a naming convention for accounts might be more trouble than it's worth. When you're working with dozens of servers and hundreds or thousands of users, however, a scheme for naming user and group accounts as well as network devices is crucial. For user accounts, some considerations for a naming convention include the following:

- Is there a minimum and maximum number of characters user account names should have?
- Should the username be based on the user's real name, or, if security is of utmost importance, should usernames be more cryptic and, therefore, difficult to guess?
- Some OSs distinguish between uppercase and lowercase letters. Should usernames contain both as well as special characters, such as periods and underscores?

There's no right or wrong answer to these questions, but after you devise a policy, you should stick to it so that when it's time to create a new user account, your naming conventions make the process straightforward.

As part of creating user accounts, passwords must also be created. The considerations for password naming conventions include the following:

- *Minimum length*—In environments where a user account is based on a user's real name, all that's needed to access the account is guessing the password. Longer passwords are harder to guess and, therefore, more secure.
- *Complexity requirements*—Using uppercase and lowercase characters along with special characters (such as @, \$, %, and so forth) makes passwords considerably more difficult to guess, even with password-guessing software.
- *User or administrator created*—In most cases, users create their own passwords after an administrator gives them an initial password. However, to ensure that passwords are complex enough, a randomly created password is sometimes used.
- *Password change frequency*—Many networks require frequent password changes to enhance security. However, if changes are required too frequently, users are more apt to write down their passwords, which is a major security risk.

There are other considerations for working with passwords, some of which are particular to the OS on which the user account is created. You learn more about password-handling options later in “Working with Accounts in Windows.”

Group account names also warrant careful planning. The group name should reflect the group membership (such as a department name) or the resource to which the group is assigned permissions or rights. An example of a group name that reflects a resource permission assignment is `NAS_4thFloor`, indicating that group members have access to the NAS server on the 4th floor. In some cases, a group name might reflect the role group members have in the company, such as supervisors, administrators, executives, and so forth.

The most important aspect of naming conventions is that after you have established them, stick to them and allow only users who are well versed in the conventions to create accounts.

Working with Accounts in Windows

The details of account creation differ, depending on whether accounts are created on a Windows client OS or in Windows Server with Active Directory. You can specify many more user account properties in Active Directory, but basic account creation is similar to Windows 7.

This section discusses user and group accounts in Windows Server 2008 with Active Directory, and you work with accounts in Windows 7 in the hands-on projects. A key difference to remember when working with accounts in Active Directory is that they're used to log on to the Windows domain and can be used to access resources on all computers that are domain members. An account created in a Windows client OS is used only to log on to that particular computer and access resources only on that computer.

When Windows is first installed, two users are created: Administrator and Guest. On a Windows Server 2008 domain controller, the Guest account is disabled, and in Windows 7,



both Administrator and Guest are disabled. In Windows 7, you create a new user with administrator privileges during installation. The Guest account is rarely used and poses a security risk, which is why it's disabled. You can enable the account if you like, but best practices dictate creating new accounts for guest users of your network. The Administrator account has full access to a computer, and in a Windows domain, the domain Administrator account's access is extended to all computers that are domain members. You must give careful consideration to who can log on as Administrator and who's a member of the Administrators group.

Creating User Accounts in Windows Domains Windows domain users are created in Active Directory Users and Computers (see Figure 9-1). As you can see, several folders are available for organizing users, groups, and other domain elements. You can also create your own folders, called organizational units (OUs), to match your company's organizational scheme. For example, you can create a folder for each department or create folders representing corporate office locations. In Figure 9-1, the open folder named Users contains the Administrator and Guest accounts and many of the default groups created when Active Directory is installed. You can create additional users in this folder, but it's better to add OUs and create users and groups in the OU structure you specify.

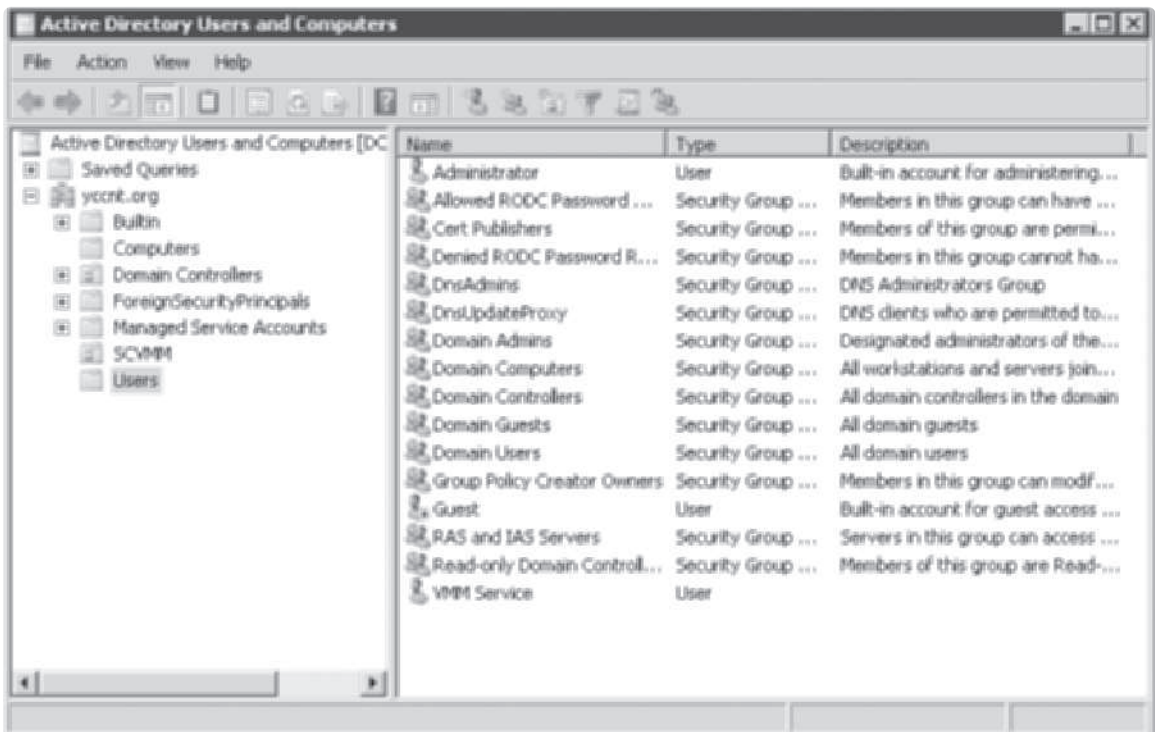


Figure 9-1 Active Directory Users and Computers

Courtesy of Course Technology/Cengage Learning

To create a user, open the folder where you want to create the account. Right-click the folder, point to New, and click User, or you can click the user icon on the Active Directory Users and Computers toolbar. The New Object - User dialog box opens (see Figure 9-2). Everything you create in Active Directory is considered an object.

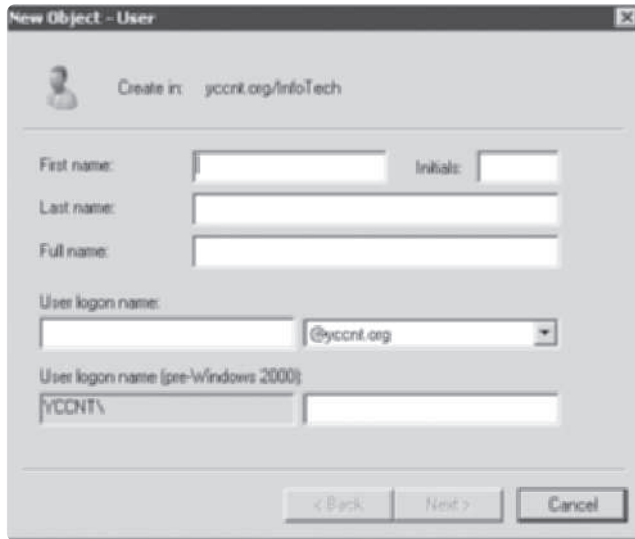


Figure 9-2 Creating a user in Active Directory

Courtesy of Course Technology/Cengage Learning

You don't have to fill in all the fields, but you must enter something in the Full name and User logon name text boxes. The user logon name isn't case sensitive, so if the logon name is JSmith, the user can log on with jsmith or JSMITH or any combination of upper-case and lowercase letters. The drop-down list next to the User logon name text box shows the default domain the user logs on to. In a network with multiple domains, the user might need to log on to the network with the syntax *LogonName@domain*, which is referred to as the user principal name (UPN). In most cases, a user needs only the logon name to log on. After entering the full name and user logon name, you click Next to get to the window shown in Figure 9-3, where you enter the password and confirm it. The password is case sensitive. As Figure 9-3 shows, the password isn't shown as you type it for security reasons. You can also choose the following options for the user's initial logon and password:

- *User must change password at next logon*—The user is prompted to change the password at the next logon. Administrators sometimes create accounts with a default password based on the user's name or phone number that must be changed at the next logon. This option can also be set when users forget their passwords and the support staff changes passwords for them.
- *User cannot change password*—In environments where the administrator wants to maintain control of passwords, this option can be set to prevent users from changing their own passwords. It's also used when multiple users have a common generic account for logging on (such as "salesperson").



- *Password never expires*—Users can be required to change their password periodically. If this option is set, the user isn't subject to the password change requirement.
- *Account is disabled*—If a user account is created several days before it's going to be used, the account can be disabled at first and then enabled when the user joins the company. In addition, if a user leaves the company or will be gone for an extended period, the account can be disabled. Often when a user leaves the company, the account is disabled rather than deleted so that the user hired as the replacement can use the same account after you rename it and change the password. In this way, the new user has all the same permissions and rights as the previous user.




Figure 9-3 Setting the password and additional account options

Courtesy of Course Technology/Cengage Learning

After a user account is created, double-click it to open its properties. Compared with user accounts in client OSs, user accounts in Active Directory have far more tabs in property dialog boxes. Figure 9-4 shows properties for a user in Active Directory on the left and for an account in Windows 7 on the right. Notice that the two property dialog boxes have the Member Of tab in common, where you can see which groups a user belongs to and add or remove the user from groups.



TIP

When a user is added to or removed from a group, the setting takes effect the next time the user logs on; if a user is already logged on, he or she must log off and log back on.

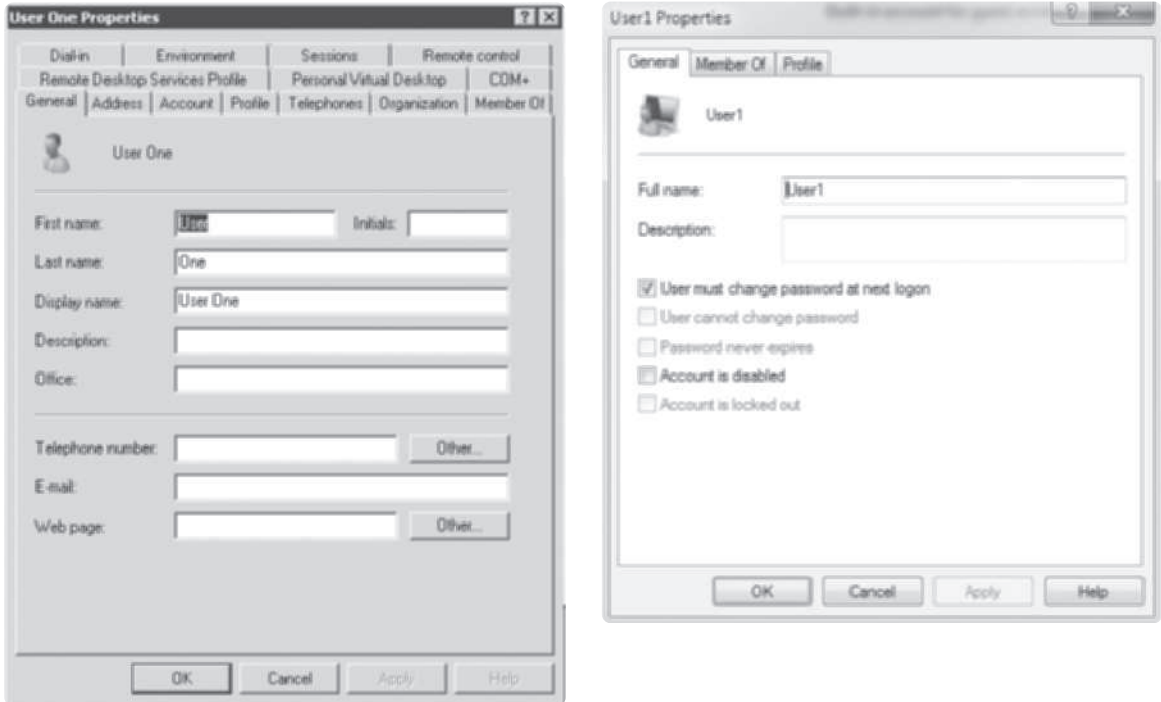


Figure 9-4 User properties in Active Directory (left) and in Windows 7 (right)

Courtesy of Course Technology/Cengage Learning



Hands-On Project 9-1: Creating Users in a Windows Client OS

Time Required: 20 minutes

Objective: Create a user account in a Windows client OS.

Required Tools/Equipment: Your classroom computer with a Windows client OS installed. Windows 7 is used in this project, but the steps are similar for Windows Vista and XP.

Description: In this project, you create a new user in the Computer Management console.

1. Log on to your computer as an administrator.
2. There are two tools for creating new accounts in a Windows client OS. One is User Accounts in Control Panel, which is mostly for home users. The other is Local Users and Groups in the Computer Management console. Local Users and Groups gives administrators more control over user properties and has more in common with Active Directory Users and Computers, so it's used in this project. First, you need to make access to common administrative tools easier. (If you're already logged on with an administrator account, you can skip this step.) Right-click **Start** and click **Properties**.



Click the **Start Menu** tab, if necessary, and click **Customize**. In the Customize Start Menu dialog box, scroll all the way down until you see System Administrative Tools. Click **Display on the All Programs menu and the Start menu**, and then click **OK** twice. The Administrative Tools folder is now on your Start menu.

3. Click **Start**, point to **Administrative Tools**, and click **Computer Management**. Click **Local Users and Groups**. In the left pane, you see two folders: Users and Groups. Double-click **Users** to display a list of users on your computer in the middle pane (see Figure 9-5). Notice the Guest user in the figure, shown with a black arrow in a white circle to indicate that the account is disabled.

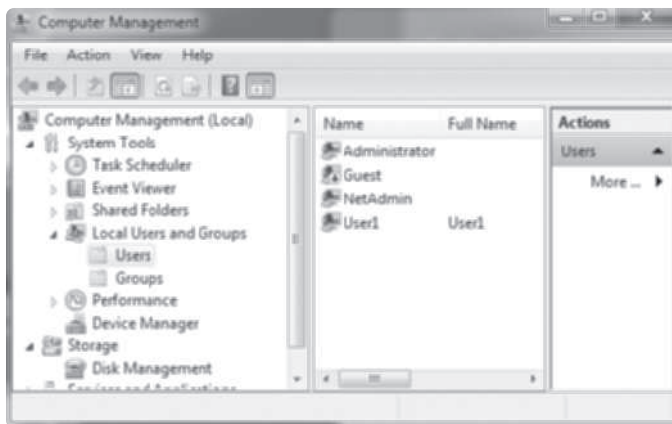


Figure 9-5 Viewing Local Users and Groups in Computer Management

Courtesy of Course Technology/Cengage Learning

4. Right-click empty space in the middle pane and click **New User**. In the New User dialog box, type **NewGuest1** in the User name text box. In the Full name text box, type **New Guest User 1**, and in the Description text box, type **A new guest user account**.
5. Type **guestpass** in the Password text box and again in the Confirm password text box.
6. Leave the **User must change password at next logon box** check box selected (see Figure 9-6), and click **Create**. The New User dialog box clears so that you can create another user. Click **Close**.
7. In Local Users and Groups, double-click **NewGuest1** to view its properties. Click the **Member Of** tab. By default, all new users are put in a group called Users; this is also the case when a new user is created in Active Directory. Click **Cancel**.
8. Double-click the **Groups** folder in the left pane to see a list of groups Windows creates. Double-click the **Users** group. You'll see a list of users who are members of the Users group, including NewGuest1. You'll probably also see some special groups named Authenticated Users and INTERACTIVE, which are internal groups used by Windows. Click **Cancel**.

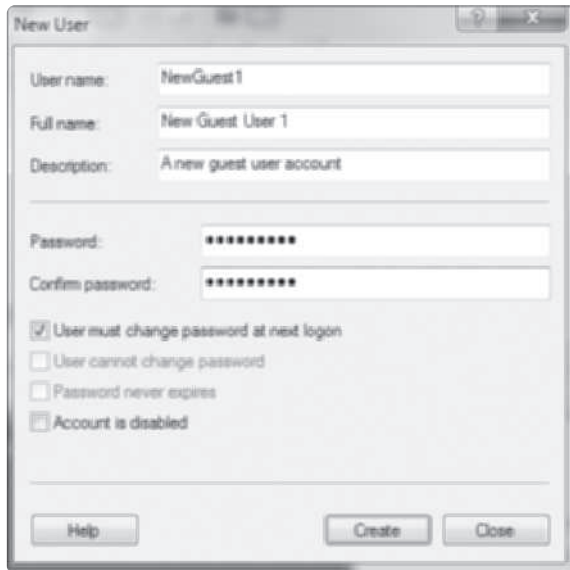


Figure 9-6 Creating a user in Windows 7

Courtesy of Course Technology/Cengage Learning

9. Close Computer Management and log off Windows. When the logon window opens, New Guest User 1 is shown as a user. Click **New Guest User 1**. Type **guestpass** in the Password text box and press **Enter** or click the arrow to log on. In the message box stating that your password must be changed, click **OK**. Type **guestpass1** in the New password and Confirm password text boxes. Press **Enter** or click the arrow.
10. In the message box stating that the password has been changed, click **OK** to log on. When you see your desktop, log off again, but leave Windows running for the next project.

Creating Group Accounts in Windows Domains Group accounts are easy to create. All they require is a name, and after they're created, you can begin adding users as members. The process is similar to creating a user. In Active Directory, the New Object - Group dialog box looks like Figure 9-7. The Group name (pre-Windows 2000) text box is used for backward compatibility with older Windows OSs. The other options, Group scope and Group type, are used only in Windows domains. The Group scope has three options:

- *Domain local*—Can be used to assign permissions to resources only in the domain in which the group is created. Although domain local groups can contain users from any domain, they're used mainly to hold global groups and assign permissions to global group members.
- *Global*—The default option, global groups contain users from the domain in which they're created but can be assigned permissions to resources in other domains in a



multidomain network. Their main purpose is to group users together who require access to similar resources.

- *Universal*—Used in multidomain networks; users from any domain can be members and be assigned permission to resources in any domain.

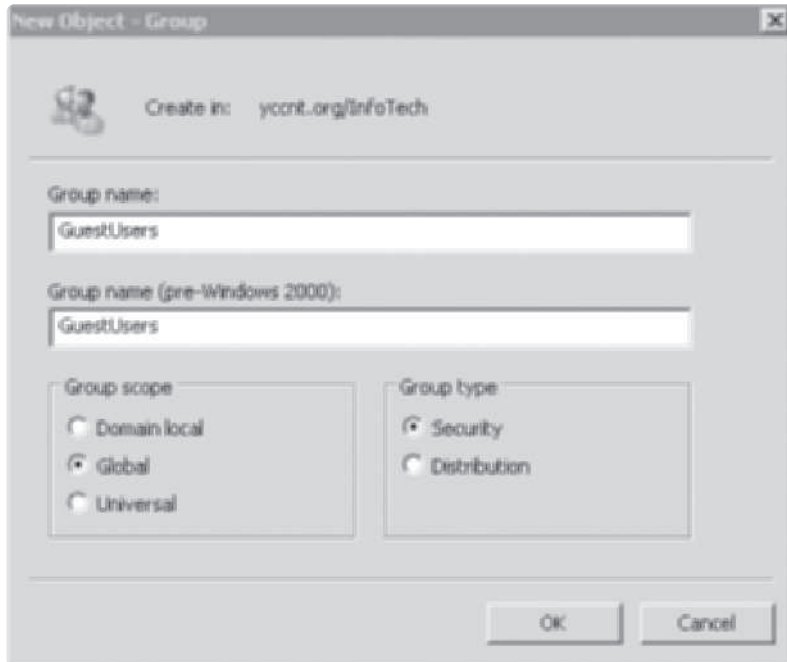


Figure 9-7 Creating a group in Active Directory

Courtesy of Course Technology/Cengage Learning



A detailed discussion on group scope is beyond the scope of this book. For a complete discussion, see *MCTS Guide to Microsoft Windows Server 2008 Active Directory Configuration* (Course Technology, 2010, 1423902351).

The group type option is set to Security by default. Distribution groups are used only for tasks such as sending all group members an e-mail when you run an Active Directory–integrated e-mail program, such as Microsoft Exchange.

Windows Default Groups Aside from groups you create to organize users and assign permissions, Windows defines a number of default groups. **Default groups** have preassigned rights that apply to all group members. Table 9-1 shows the most important default domain local groups in Windows Server 2008 running Active Directory and the rights assigned to these groups.

Table 9-1 Some Windows Server 2008 default domain local groups

Group	Rights
Administrators	Has complete control over the computer and domain
Account Operators	Can administer user and group accounts for the local domain
Backup Operators	Can back up and restore files that users normally can't access
Guests	Is allowed guest access to domain resources; same access as the Users group
Print Operators	Can add, delete, and manage domain printers
Server Operators	Can administer domain servers
Users	Has default access rights that ordinary user accounts have

In addition, Windows Server 2008 has numerous default global groups, including Domain Admins, Domain Users, and Domain Guests. Essentially the same as domain local groups with similar names, these groups apply to entire domains rather than a single machine.

Special Identity Groups Special identity groups, some described in Table 9-2, don't appear as objects in Active Directory Users and Computers or in Local Users and Groups, but they can be assigned permissions and rights. Membership in these groups is controlled dynamically by Windows, can't be viewed or changed manually, and depends on how an account accesses the OS. For example, membership in the Authenticated Users group is assigned to a user account automatically when the user logs on to a computer or domain.



Table 9-2 Some Windows special identity groups

Special identity group	Description
Authenticated Users	Members are any user account (except Guest) that logs on to a computer or domain with a valid username and password.
Creator Owner	A user becomes a member automatically for a resource he or she created (such as a folder).
Everyone	Refers to all users who access the system. Similar to the Authenticated Users group but includes the Guest user.
Interactive	Members are users logged on to a computer locally or through Remote Desktop.
Network	Members are users logged on to a computer through a network connection.
System	Refers to the Windows OS.
Self	Refers to the object on which permissions are being set.



Hands-On Project 9-2: Working with Groups in a Windows Client OS

Time Required: 10 minutes

Objective: Create a group and add a user to the group.

Required Tools/Equipment: Your classroom computer with a Windows client OS installed. This project uses Windows 7, but the steps are similar for Windows Vista and XP.

Description: In this project, you create a new group in the Computer Management console and then add a user to the group.

1. Log on to your computer as an administrator, if necessary.
2. Click **Start**, point to **Administrative Tools**, and click **Computer Management**. Click **Local Users and Groups**, and then double-click **Groups** in the left pane.
3. Right-click empty space in the middle pane and click **New Group**. In the New Group dialog box, type **GuestUsers** in the Group name text box. In the Description text box, type **A group for guest users of this computer**.
4. Click **Add**. Examine the Select Users dialog box shown in Figure 9-8. It's similar to what you see when adding a user to a group in Active Directory.

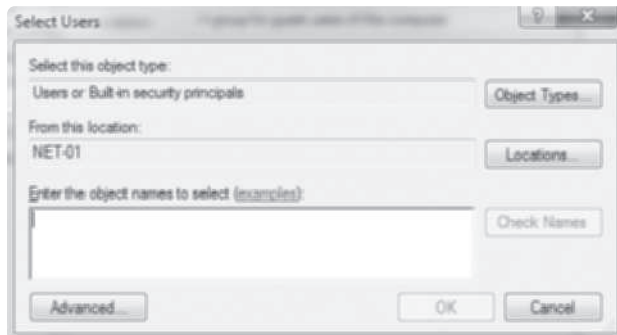


Figure 9-8 Selecting users to add to a group

Courtesy of Course Technology/Cengage Learning

5. To limit the types of objects Windows shows if you click the **Advanced** button to search for objects to add to a group, you use the **Object Types** button. Click **Locations**. You have only one option unless your computer is a member of a domain. If so, you can select objects from the domain; otherwise, you can choose only objects created on your computer. Click **Cancel**.
6. You can type the user's name in the text box, but to select from a list, click **Advanced**. Click **Find Now** to list available users and groups you can add as group members. Click **NewGuest1** and click **OK**. Notice in the **Select Users** dialog box that the user is specified as **NET-01\NewGuest1**. **NET-01** is the name of the computer or domain where the user was created; in this case, it's the computer name. Click **OK**.

7. NewGuest1 is then listed as a member of the group. Click **Create** to finish creating the group, and then click **Close**.
8. NewGuest1 is now a member of both the GuestUsers and Users groups. Remember that the result of changing group membership takes effect the next time the user logs on. If you wanted to remove NewGuest1 from the default Users group, you would double-click the Users group, right-click NewGuest1, and click Remove. However, doing so removes NewGuest1 from the list of users in the Windows 7 logon window, so for now, leave this account as a member of both groups. Close Computer Management, and log off Windows for the next project.

User Profiles A user profile is a collection of a user's personal files and settings that define his or her working environment. By default, a user profile is created when a user logs on to a computer for the first time and is stored in a folder that usually has the same name as the user's logon name. On a Windows Server 2008 or Windows 7 computer, a profile is created as a subfolder of the Users folder, which is on the same drive as the Windows folder, usually C. Figure 9-9 shows the profile folder hierarchy on a typical Windows 7 system.

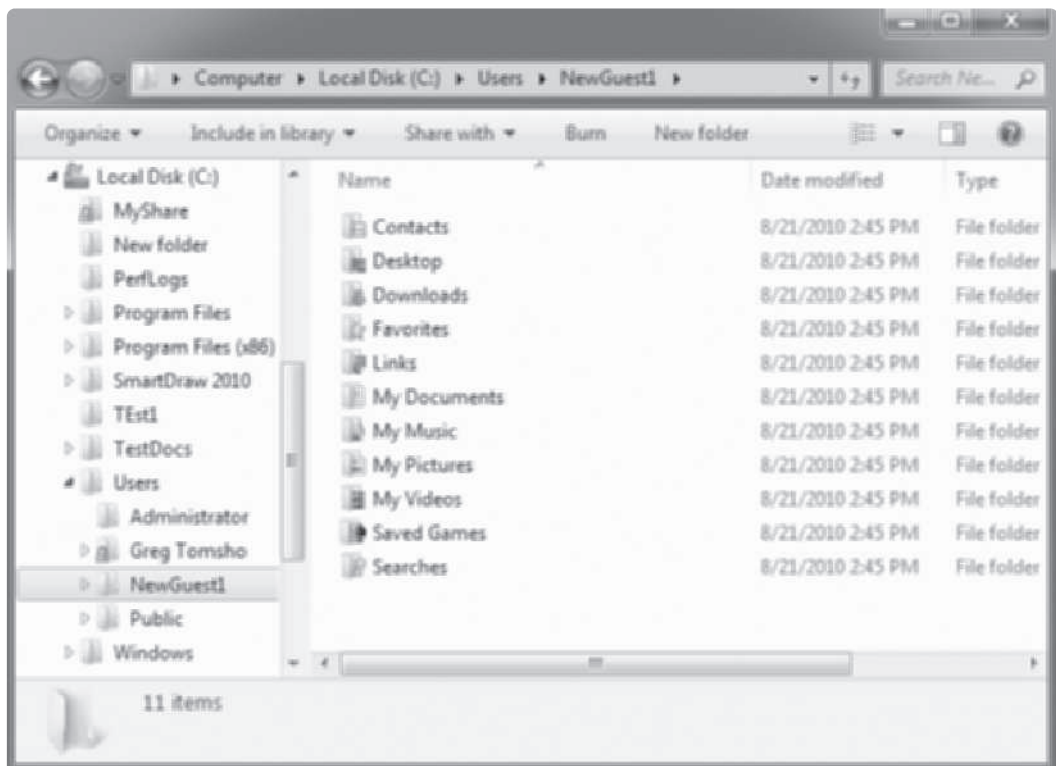


Figure 9-9 The files and folders composing a user profile

Courtesy of Course Technology/Cengage Learning

A profile contains personal data folders a user maintains as well as files and folders containing user and application settings. Some files and folders in the profile are hidden or system files that can't be viewed with the default Windows Explorer settings. To view all files in the profile, you must enable the option to view hidden and system files in Windows Explorer. The following list describes some key files and folders in a user's profile. Because some folder names have changed in Windows 7 and Server 2008, the Windows XP name is in parentheses. If the folder doesn't exist in a Windows XP profile, N/A is noted in parentheses.

- *AppData (N/A)*—A hidden folder that's the default location for user application data
- *Desktop*—Contains desktop items, such as shortcuts and files
- *Documents (My Documents)*—The default location where applications store saved documents
- *Downloads (N/A)*—The default location for files downloaded via a Web browser
- *Favorites*—Bookmarks in Internet Explorer
- *Music (My Music)*—The default location for saved music files
- *Pictures (My Pictures)*—The default location for saved picture files
- *Ntuser.dat*—A hidden system file containing user preferences for Windows and application settings; merged with the Registry when a user logs on to Windows

A user profile stored on the same system where the user logs on is called a **local profile**. A local profile is created from a hidden profile called Default the first time a user logs on to a system; to see this profile, you must enable the option for viewing hidden and system files in Windows Explorer. When users log off, their profile settings are saved in their local profiles so that the next time they log on, all their settings are preserved. However, if a user logs on to a different computer, the profile is created again from the Default profile. If administrators want to make users' profiles available on any computer they log on to, they can set up roaming profiles, discussed next.



To view hidden and system files in Windows Explorer, click Organize, Folder and search options, and then click the View tab. Click Show hidden files, folders, or drives, and click to clear the Hide protected operating system files check box.

A **roaming profile** follows the user no matter which computer he or she logs on to. It's stored on a network share so that when a user logs on to any computer in the network, the profile is copied from the network share to the profile folder on the local computer. This local copy of the roaming profile is referred to as the profile's "cached copy." Any changes the user makes to the profile are replicated from the locally cached copy to the profile on the network share when the user logs off.

The location of a roaming user's profile is specified in the Profile tab of a user's properties (see Figure 9-10). This tab is identical in both Active Directory and a client OS. The profile path points to a network share by using the UNC path (discussed in Chapter 8).




Figure 9-10 User profile settings

Courtesy of Course Technology/Cengage Learning

Roaming profiles are rarely used in workgroup networks, as you would have to add the user account to every computer in the workgroup, but it's a feature Active Directory administrators use often. A third type of profile, called a mandatory profile, discards a user's profile changes at logoff so that the profile is always the same when the user logs on. Mandatory profiles are sometimes used on shared computers and for guest accounts.

Working with Accounts in Linux

User and group accounts in Linux are used for the same purposes as in Windows: user authentication and authorization. Linux OSs also have a default user who has full control over all aspects of the system. In Linux, this user is named root. As in Windows, creating additional users to log on to and use the system is recommended so that the root user account is used only when you're performing tasks that require root privileges. In fact, some Linux distributions require creating a user during installation because logging on as root isn't allowed. You can access root privileges only by entering a special command.

Because most Linux administration takes place at the command line, this method for creating new users is discussed first. You boot Linux to a command prompt without the GUI or boot to the GUI and open a terminal window. In its simplest form, user creation is a matter of using the `adduser newuser` command (replacing `newuser` with the logon name for the user account you're creating). You're prompted to create a new password and



enter the user's full name and other information. Both the logon name and password are case sensitive in Linux.

On most Linux systems, you can't run `adduser` and similar commands unless you're logged on as the root user or (preferably) preface the command with `sudo`, as in `sudo adduser newuser`. The `sudo` command, which stands for "superuser do," executes the command with root privileges. If you know you're going to use many commands requiring root privileges, you can change to the root user temporarily with the `su` command (which means "switch user"). This command attempts to switch to the root user when no user is specified, and you must enter the root user's password when prompted.

With some commands, if they're entered with a username, as in `passwd testuser`, the command is executed only for this user account. If they're entered without a username, they're executed only for the current user. For example, users or administrators can change a user's password with the `passwd` command. User information can be changed with the `usermod` command, and you delete users with the `deluser` command.

The `adduser` command has many options. For example, you can specify another home directory, assign group memberships, and so forth. You practice using this command in Hands-On Project 9-3.



TIP

You can view extensive help on most Linux commands by typing `man command` (replacing `command` with the command you want information on). The `man` command means manual, and the help pages it displays are referred to as "man pages."

All users must belong to at least one group in Linux. When a user is created, a group with the same name as the user is also created, and the new user is made a member of this group. However, you can create groups and add users to them, just as you can in Windows.

Groups are created with the aptly named `addgroup` command. To add users as members of a group you create, you can specify this option when the user is created or use the `adduser username groupname` command or the `usermod` command.

To view the list of users, display the `/etc/passwd` file's contents with the `cat /etc/passwd` command, and to view the list of groups, display the `/etc/groups` file's contents with the `cat /etc/groups` command. The `cat` command lists a text file's contents onscreen.

For those who prefer a GUI to manage users and groups, most Linux distributions have convenient graphical interfaces for doing so. In Ubuntu Linux, the Users and Groups control panel is available (see Figure 9-11).

One reason many administrators prefer the command-line method for creating users is because they can import user information from a text file and add many users at one time by using the `newusers` command, which accepts as input a text file listing users to create.

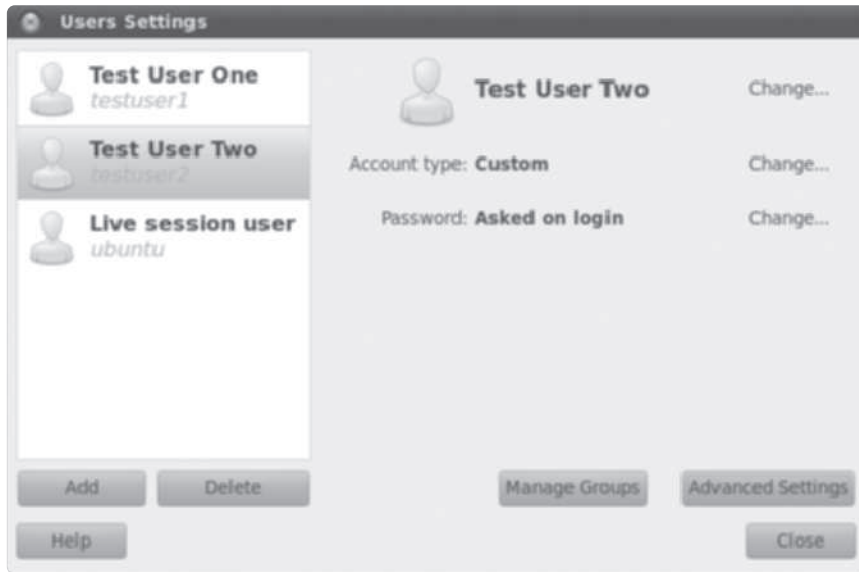


Figure 9-11 The Users and Groups control panel in Ubuntu Linux

Courtesy of Course Technology/Cengage Learning



Hands-On Project 9-3: Working with Users and Groups in Linux

Time Required: 15 minutes

Objective: Create users and groups with Linux command-line tools.

Required Tools/Equipment: Your classroom computer with Linux installed or a Linux Live CD. This project uses Ubuntu Linux 10.4, but the steps are similar in most Linux distributions.

Description: In this project, you create users with the `adduser` command and groups with the `addgroup` command. Next, you add users as members of these groups with the `adduser` and `usermod` commands.

1. Log on to your Linux computer and open a terminal window. These steps assume you don't log on as root. If you do, you don't need to preface commands with `sudo`. In Ubuntu Linux, click **Applications**, point to **Accessories**, and click **Terminal**.
2. At the terminal prompt, type `man adduser` and press **Enter** to get an overview of what the man pages for the `adduser` command contain. Press the **Page Up** and **Page Down** keys to scroll through the man pages. Press `q` when you're finished.
3. To view current users on the Linux system, type `cat /etc/passwd` and press **Enter**. Another way to view a text file you can page up and down through is using the `less` option. Type `less /etc/passwd` and press **Enter**. Use the arrow keys or **Page Up** and **Page Down** keys to scroll through the file. Many of the user accounts you see in this file are system accounts and aren't used to log on to the OS. Press `q` to quit.

4. Display the list of groups by typing `less /etc/group` and pressing **Enter**. When you're finished, press `q`.
5. To create a user, type `adduser testuser1` and press **Enter**. If you aren't logged on as root, you get a message stating that only the root user can add users and groups. Type `sudo adduser testuser1` and press **Enter**. If you're prompted for your root password, enter it. Periodically, when you use the `sudo` command, you're prompted for this password.
6. Next, you're prompted to enter a password for creating a user. Type `Password01` and press **Enter** and type it again. (Your keystrokes aren't displayed.)



If you don't enter the same password when asked to retype it, you get a message stating that the passwords don't match, and you're prompted to try again.

7. Next, you're prompted to enter the user's full name. (The full name and the other information you're prompted to enter are optional.) Type `Test User One` and press **Enter**, and then press **Enter** four times for the remaining prompts. When asked whether the information is correct, press `y` and press **Enter**. The user is then created.
8. Create another user with the logon name `testuser2`.
9. Type `less /etc/passwd`, press **Enter**, and page to the bottom of the file, where you see the users you created. Press `q` and then display the group file to see that groups named `testuser1` and `testuser2` were also created. (*Hint:* Remember that you can use the arrow keys to scroll through recently used commands.)
10. Type `sudo addgroup testgroup1` and press **Enter**. To add `testuser1` to `testgroup1`, type `sudo adduser testuser1 testgroup1` and press **Enter**. To add a user to a group with the `usermod` command, type `sudo usermod -a -G testgroup1 testuser2` and press **Enter**. Type `cat /etc/group` and press **Enter** to list all groups. You should see your new group at the end of the file along with a list of its members.
11. You can view a user's group memberships with the `groups` command. Type `groups testuser1` and press **Enter**. `Testuser1` is listed as a member of the `testuser1` and `testgroup1` groups.
12. Close the terminal window, and leave Linux running for the next project.



Hands-On Project 9-4: Using a GUI to Work with Linux Accounts

Time Required: 15 minutes

Objective: Create users and groups in the Linux Users and Groups control panel.

Required Tools/Equipment: Your classroom computer with Linux installed or a Linux Live CD. This project uses Ubuntu Linux 10.4, but the steps are similar in most Linux distributions.

Description: In this project, you work with Linux accounts in the Users and Groups control panel.

1. Log on to your Linux computer, if necessary. Click **System**, point to **Administration**, and click **Users and Groups**.

2. The Users Settings dialog box (shown previously in Figure 9-11) shows all current users (excluding root and system users). Click **Test User One** to select it, and click **Advanced Settings** to open the Change Advanced User Settings dialog box.
3. There are three tabs. Click the **Advanced** tab to see information such as the account's home directory, shell, main group, and user ID (see Figure 9-12). The home directory is similar to the user's profile folder in Windows. The shell is the default command-line interface used when a terminal window is opened. The main group (also called a primary group) is used when setting permissions on files the user creates. The user ID is a numeric value that identifies the account internally.

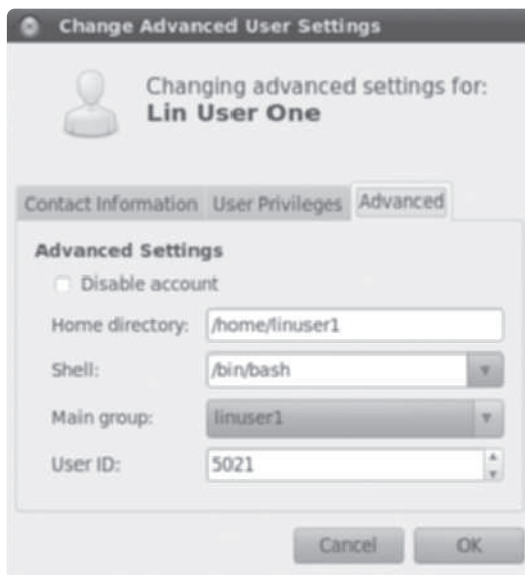


Figure 9-12 The Change Advanced User Settings dialog box in Ubuntu Linux

Courtesy of Course Technology/Cengage Learning

4. Click the **User Privileges** tab, where you can set the types of tasks the user account is allowed to perform. In Windows, these tasks are called rights.
5. Click the **Contact Information** tab. The information shown here is what you were prompted to enter when you used the `adduser` command and is optional. Click **Cancel**.
6. In the Users Settings dialog box, click **Manage Groups**. You see a list of groups, which is the same list displayed when you used the `less /etc/group` command, except the groups are sorted alphabetically here. You can create or delete groups and view group membership in this dialog box. Double-click `testgroup1` to see its members. You can add other users to the group by clicking the check boxes next to their usernames. Click **Cancel** and then **Close**.
7. In the Users Settings dialog box, click **Add**. Type **Test User Three** in the Name text box and `testuser3` in the Short Name text box. You can click the check box to encrypt the user's home directory, if you want. Click **OK**.



8. In the Change User Password dialog box, click the **Generate random password** option button. A random password is created. Click **OK**.
9. In the Users Settings dialog box, click **Manage Groups**. Double-click **testgroup1**, and then click the **Test User Three** check box to add this user to the testgroup1 group. Click **OK**, and then click **Close** twice.
10. Shut down the Linux computer. Note that if you're using a Live CD to run Linux, all your changes are discarded when you shut down.

Storage and File System Management

Managing storage on networks is becoming more of a challenge. Users are storing larger amounts and more varied types of data, and network administrators have to make sure enough storage space is available as well as manage who has access to it. In some cases, the challenge is preventing users from storing inappropriate types of data on company servers, such as music files and pictures unrelated to work. This section describes some tools for managing locally attached storage on a server and then discusses how an administrator controls access to files and folders on the file system.

Locally attached storage is a device, such as a hard disk, that's connected to a storage controller on the server. Usually, the storage is physically housed inside the server case and connected to a SATA or SCSI controller. However, it can also be external, attached via USB or external SATA (eSATA) connectors. In either case, the server sees the storage the same way—as a disk containing one or more volumes or, if the disk is empty, unallocated space.

Volumes and Partitions

A **volume** is part or all of the space on one or more disks that contains or is ready to contain a file system. In Windows, volumes with file systems are usually assigned a drive letter. In Linux, volumes are mounted in the file system and accessed as though they were just another folder. Starting with Windows Vista, Windows volumes can also be mounted in the file system instead of assigning a drive letter.

The term “partition” is sometimes used interchangeably with “volume,” but these terms don't always describe the same thing. To understand the difference, look at how Windows views a hard disk. Disks are numbered starting with Disk 0, Disk 1, and so forth. A Windows disk can be categorized as a basic disk or a dynamic disk. By default, all newly installed disks are considered basic disks. A **basic disk** can be divided into one to four partitions as follows, with a maximum of four partitions consisting of the following:

- *One to four primary partitions*—A **primary partition** can be formatted with a file system and assigned a drive letter or mounted in an empty folder on an existing drive letter. It's also a volume.
- *One extended partition*—An **extended partition** can't be formatted with a file system or assigned a drive letter. It's divided into one or more logical drives, each of which can be formatted and assigned a drive letter. A logical drive is considered a volume,

but an extended partition is not. You can create an extended partition only if there are fewer than four primary partitions.

Only a primary partition can be the **active partition**, which is a partition that can hold boot files (called the “boot loader”) the BIOS loads before it can start the OS. An extended partition/logical drive can’t be booted, but it can store OS files.

The active primary partition storing the Windows boot loader is referred to as the **system partition**. The partition or logical drive holding the Windows OS files is called the **boot partition**. In most cases, these two are the same because both the boot loader files and Windows OS files are often located on the C drive.

A **dynamic disk** can be divided into one or more volumes; the term “partition” isn’t used in this context. You can create up to 1000 volumes per dynamic disk (although no more than 32 are recommended). A dynamic disk offers features that a basic disk doesn’t, namely RAID (discussed later in “Protecting Data with Fault Tolerance”) and disk spanning, which is creating a volume that occupies space on two or more disks.

Linux systems refer to disks by using their device driver name plus a letter, starting with “a.” For example, the first SATA or SCSI disk on a Linux system is named `/dev/sda`, the second disk is `/dev/sdb`, and so forth. Partitions or volumes are referred to by using the device name and a number. The first volume on the first disk in Linux is named `/dev/sda1`, the second volume is `/dev/sdb2`, and so on.

Whether you’re working with a partition or a volume in Windows or Linux, what makes disk storage usable is the file system. The next sections discuss the primary file systems in Windows and Linux: FAT, NTFS, and Ext3/Ext4.

The FAT File System

The File Allocation Table (FAT) file system has two variations: FAT16 and FAT32. FAT16 is usually referred to simply as “FAT.” It’s been around since the mid-1980s, which is one of its biggest strengths—it’s well known and well supported by most OSs. FAT32 arrived on the scene with the release of Windows 95 OSR2 in 1996.

The main difference between FAT16 and FAT32 is the size of the disk partition that can be formatted. FAT16 is limited to 2 GB partitions in most implementations (although Windows NT permits partitions up to 4 GB). FAT32 allows partitions up to 2 TB, but in Windows 2000 and later, Microsoft limits them to 32 GB because the file system becomes noticeably slower and inefficient with larger partition sizes. This 32 GB limitation applies only to creating partitions; Windows can read FAT32 partitions of any size. FAT16 supports a maximum file size of 2 GB, and FAT32 supports files up to 4 GB.



The number in FAT file system names refers to the number of bits available to address disk clusters. FAT16 can address up to 2^{16} (65536) disk clusters, and FAT32 can address up to 2^{32} (4,294,967,296) disk clusters. The number of disk clusters a file system can address is directly proportional to the largest partition size it supports.



As you can see, FAT has severe limitations in today's computing environment. The file size limitation alone prevents storing a standard DVD image file on a FAT system. The limitations are even more apparent when you consider reliability and security requirements of current OSs. FAT doesn't support file and folder permissions for users and groups, so any user logging on to a computer with a FAT disk has full control over every file on that disk. In addition, FAT lacks support for encryption, file compression, disk quotas, and reliability features, such as transaction recovery and journaling, all of which NTFS supports.

You might think that FAT isn't good for much, especially compared with the more robust NTFS, but FAT still has its place. It's the only file system option when using older Windows OSs, such as Windows 9x. In addition, FAT is simple and has little overhead, so it's still the file system of choice on removable media, such as flash drives. For hard drives, however, particularly on Windows servers, NTFS is unquestionably the way to go.

The NTFS File System

NTFS is a full-featured file system that Microsoft introduced with Windows NT in 1993. Since that time, its features have been expanded to help administrators gain control of expanding storage requirements. NTFS has supported file and folder permissions almost since its inception, which was a considerable advantage over FAT. Many other compelling features are available in NTFS that aren't available with FAT:

- *Disk quotas*—Enable administrators to limit the amount of disk space that users' files can occupy on a disk volume or in a folder.
- *Volume mount points*—Make it possible to associate the root of a disk volume with a folder on an NTFS volume, thereby forgoing the need for a drive letter to access the volume.
- *Shadow copies*—Enable users to keep historical versions of files so that they can revert a file to an older version or restore an accidentally deleted file.
- *File compression*—Allows users to store documents in a compressed format without needing to run a compression/decompression program to store and retrieve the documents.
- *Encrypting File System (EFS)*—Makes encrypted files inaccessible to everyone except the user who encrypted the file, including users who have been granted permission to the file.

Disk Quotas With the number and types of files requiring more disk space on corporate servers, **disk quotas** are a welcome tool to help administrators get a handle on server storage. Typically, disk quotas are set on an NTFS volume and, by default, apply to all users except administrators. Quotas can put a hard limit on the amount of storage a user's files can occupy, thereby preventing the user from storing any more files after the limit has been reached. Quotas can also be configured to create a log entry when a user has exceeded the quota, so you can determine who's using a lot of space without actually preventing users from exceeding the limit. Quotas are configured in the Quota tab of an NTFS volume's Properties dialog box (see Figure 9-13).



Figure 9-13 The Quota tab

Courtesy of Course Technology/Cengage Learning

Volume Mount Points Volume mount points enable you to access a volume as a folder in another volume instead of using a drive letter. The volume that holds the folder serving as the mount point must be an NTFS volume, and the folder must be empty. In UNIX and Linux, mount points rather than drive letters have always been used to access disk volumes, so users of these OSs should be quite comfortable with mount points. Windows volumes can be assigned both a mount point and a drive letter, if needed.

Shadow Copies Like quotas, shadow copies are enabled on an entire volume. When this feature is enabled, users can access previous versions of files in shared folders and restore files that have been deleted or corrupted. You configure shadow copies in the Shadow Copies tab (called Previous Versions in Windows 7) of a volume's Properties dialog box (see Figure 9-14). Shadow copies are disabled by default.



In Windows 7, shadow copies (called Previous Versions in Windows 7) are enabled for the Windows boot volume (usually the C drive) by default, but you must enable it for other volumes.





Figure 9-14 The Shadow Copies tab

Courtesy of Course Technology/Cengage Learning

File Compression and Encryption File compression and encryption on an NTFS volume are implemented as file attributes, like the Read-only and Hidden attributes. One caveat: These attributes are mutually exclusive, so a file can't be both compressed and encrypted. You can set only one of these two attributes.

Files can be compressed and accessed without users needing to take any explicit action to uncompress them. When a compressed file is opened, the OS decompresses it automatically. On NTFS volumes, you can enable file compression on the entire volume, a folder and its contents, or a file.

File encryption on NTFS volumes is made possible by Encrypting File System (EFS) and works in a similar manner to file compression. You can set the encryption attribute on a file or folder but not on a volume. By default, encrypted folders and files can be identified by their filenames displayed in green.

Encrypted files can usually be opened only by the user who encrypted the file. However, this user can designate other users who are allowed to access the file. In addition, in a domain environment, the domain Administrator account is designated as a recovery agent. A designated recovery agent can decrypt a file if the user account that encrypted it can no longer access it. This can happen if an administrator resets a user's password, the user account is deleted, or the user leaves the company. To encrypt a file, click the Advanced button in the General tab of a file's Properties dialog box, and then click Encrypt contents to secure data (see Figure 9-15).

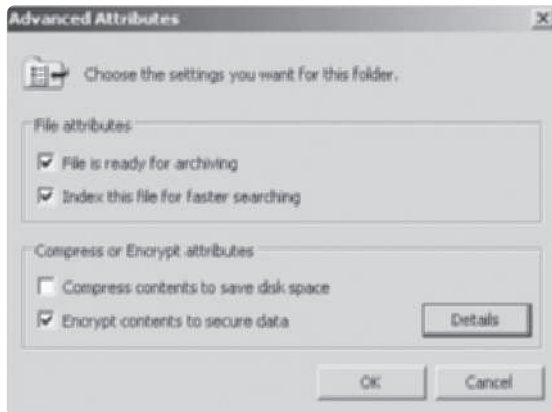


Figure 9-15 The Advanced Attributes dialog box

Courtesy of Course Technology/Cengage Learning



Most of the NTFS features discussed here are available in both client and server Windows OSs, with the exception of disk quotas.

NTFS Permissions It's important to know that there are two modes for accessing files on a networked computer: network (sometimes called remote) and interactive (sometimes called local). It follows, then, that there are two ways to secure files: share permissions and NTFS permissions. Share permissions are applied when a user attempts network access to shared files. NTFS permissions always apply, whether file access is attempted interactively or remotely through a share. That last statement might sound confusing, so take a closer look at how permissions work.

Permissions can be viewed as a gatekeeper to control who has access to folders and files. When you log on to a computer or domain, you're issued a ticket containing information such as your username and group memberships. If you attempt to access a file or folder, the gatekeeper examines your ticket and compares your username and group memberships with the file or folder's permissions list. If neither your username nor your groups are on the list, you're denied access. If you or your groups *are* on the list, you're issued an access ticket that combines all your allowed permissions. You can then access the resource as specified by your access ticket.

At least, that's how the process works when you're attempting interactive access to files. If you're attempting network access, there are two gatekeepers: one that checks your ticket against the share permissions list and, if you're granted access by share permissions, another that checks your ticket against the NTFS permissions list. The NTFS gatekeeper is required to examine your ticket only if you get past the share gatekeeper. If you're granted access by share permissions, you're issued an access ticket. Then if you're granted access by NTFS permissions, you're allowed to keep only the access ticket that gives you the *least* permission, or is the most restrictive, of the two.

For example, Mike is granted Read access by share permissions and Read and Write access by NTFS permissions. Mike gets to keep only the Read access ticket because it's the lesser of the two permissions. Another example: Neither Mike nor any of Mike's groups are on

the share permissions access list. There's no need to even examine NTFS permissions because Mike is denied access at the share permissions gate. As a final example, Mike is granted Full Control access by share permissions and Modify access by NTFS permissions. Mike's access ticket gives him Modify permission because it allows less access than Full Control.

The general security rule for assigning permissions to resources is to give users the least access necessary for their job. This rule is often referred to as the “least privileges principle.” Unfortunately, this axiom can be at odds with another general rule: Keep it simple. Sometimes determining the least amount of access a user requires can lead to complex permission schemes. The more complex a permission scheme is, the more likely it will need troubleshooting, and the more troubleshooting that's needed, the more likely an administrator will assign overly permissive permissions out of frustration.

NTFS permissions give administrators fine-grained access control over folders and files for both network users and interactive users. Unlike share permissions, which can be configured only on a shared folder, NTFS permissions can be configured on folders and files. By default, when permissions are configured on a folder, subfolders and files in that folder inherit the permissions. However, inherited permissions can be changed when needed, making it possible to have different permission settings on files in a folder.

To view or edit permissions on an NTFS folder or file, you simply access the Security tab of the object's Properties dialog box. NTFS folders have six standard permissions, and NTFS files have five. NTFS standard permissions for folders and files are as follows (see Figure 9-16):



Figure 9-16 NTFS permissions

Courtesy of Course Technology/Cengage Learning

- *Read*—Users can view file contents, copy files, open folders and subfolders, and view file attributes and permissions.
- *Read & execute*—Grants the same permissions as Read and includes the ability to run applications or scripts. When this permission is selected, List folder contents and Read are selected, too.
- *List folder contents*—This permission applies only to folders and grants the same permission as Read & execute. However, because it doesn't apply to files, Read & execute must also be set on the folder to allow users to open files in the folder.
- *Write*—Users can create and modify files and read file attributes and permissions. However, this permission doesn't allow users to read or delete files. In most cases, the Read or Read & execute permission should be given with the Write permission.
- *Modify*—Users can read, modify, delete, and create files. Users can't change permissions or take ownership. Selecting this permission automatically selects Read & execute, List folder contents, Read, and Write.
- *Full control*—Users can perform all actions given by the Modify permission with the addition of changing permissions and taking ownership.

Permissions and rights assignments should be made by using groups whenever possible. Users can be members of more than one group and, therefore, have all the rights and permissions assigned to all the groups of which they're members. In this sense, rights and permissions are cumulative. So a user has the permission to read a file if a group in which he or she is a member has been assigned Read permission for the file. However, if another group the user is a member of has Modify permission for the same file, the user also has the Modify permission. The exception is the Deny permission, which takes precedence over the Allow permission. For example, if a user belongs to a group that has the Allow Modify permission for a file and also belongs to a group that has been assigned Deny Modify, the user is denied access to the file.



Hands-On Project 9-5: Using Windows Disk Management

Time Required: 15 minutes

Objective: Create and delete volumes in Windows Disk Management.

Required Tools/Equipment: Your classroom computer with one unallocated disk

Description: In this project, you create and delete volumes and convert a basic disk to dynamic and then back to basic.

1. Log on to your computer as an administrator.
2. Click **Start**, point to **Administrative Tools**, and click **Computer Management**. In Computer Management, click **Disk Management**.
3. Figure 9-17 shows Disk Management in Windows 7; Disk 0 has three primary partitions. The partition labeled System Reserved is the Windows system partition and is marked "Active." Remember that the system partition is where boot files are located. It doesn't have a drive letter assigned, which is the default configuration in Windows 7.

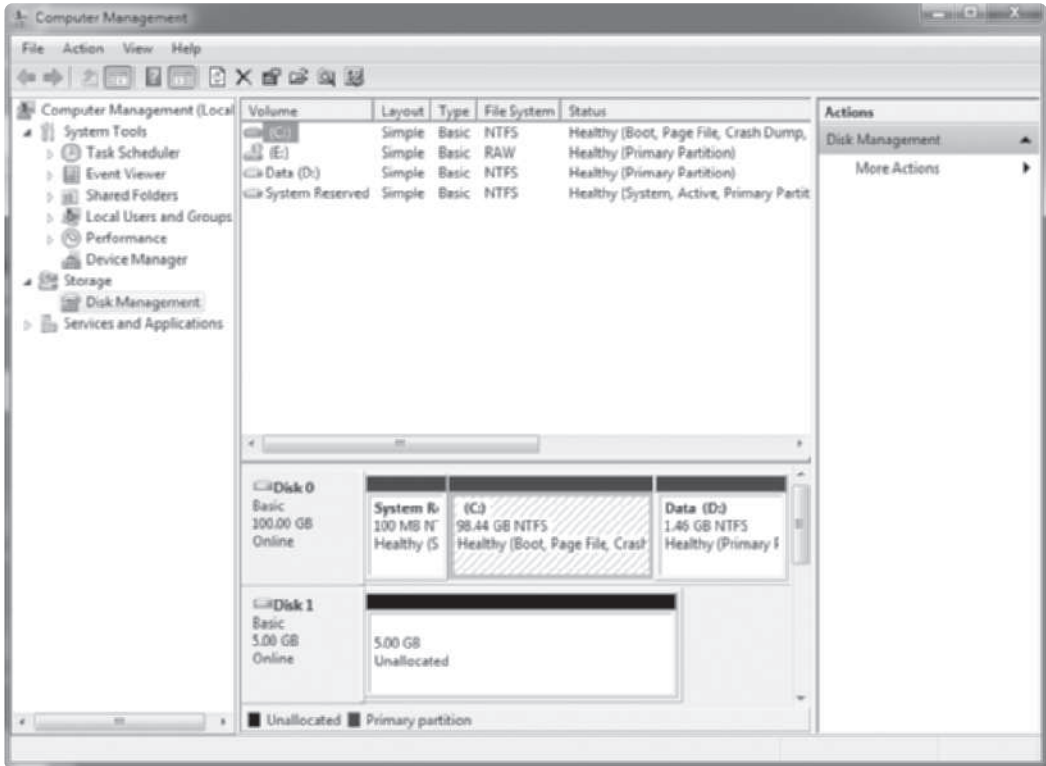


Figure 9-17 The Disk Management console

Courtesy of Course Technology/Cengage Learning

- The C drive is a primary partition and is the Windows boot partition, which indicates that a Windows OS is used on that partition. The D drive is also a primary partition. Disk 1 is unallocated, which means no volumes have been created on it. Right-click **Disk 1** in the box marked Unallocated. (If Disk 1 isn't unallocated, choose a disk that is.) Notice that the other volume types are grayed out. Disk 1 is a basic disk and supports only simple volumes. Click **New Simple Volume**.
- In the New Simple Volume Wizard welcome window, click **Next**. In the Simple volume size in MB text box, type **500**, and then click **Next**. In the Assign Drive Letter or Path window, click the drive letter list arrow and click **M**. You also have the option to mount the drive in an empty NTFS folder or make no assignment. Click **Next**.
 - In the Format Partition window (shown in Figure 9-18), you can choose the file system for the volume. Click the **File system** list arrow to see the choices, and click **NTFS**. Click the **Allocation unit size** list arrow to view the choices. The allocation unit size specifies the size of disk clusters. Click **Default** to select the default size. In Windows 7, the default cluster size is 4096 (4 KB) for volumes up to 2 TB. In the Volume label text box, type **Vol1**, and make sure the **Perform a quick format** check box is selected. Click **Next** and then **Finish**.




Figure 9-18 The Format Partition window

Courtesy of Course Technology/Cengage Learning

6. The volume is formatted, and then the drive letter is assigned and the status is set to Healthy (Primary Partition). Right-click the **Vol1** volume and view the actions you can perform. You can mark the partition as active, change the drive letter or mount path, format it, extend or shrink the volume, and delete it. Click **Extend Volume**. In the Extend Volume Wizard welcome window, click **Next**. Type **100** in the Select the amount of space in MB text box. Click **Next** and then **Finish**.
7. Next, convert the basic disk to a dynamic disk. Right-click **Disk 1** (or the disk you used to create the volume) and click **Convert to Dynamic Disk**. Click **OK** in the Convert to Dynamic Disk prompt. In the Disk to Convert prompt, click **Convert**. You get a message stating that you can't start an installed OS from the disk. Click **Yes**. Converting a disk from basic to dynamic retains the data on the disk.
8. To convert the disk back to basic, you must delete the volume first. Right-click **Vol1** and click **Delete Volume**. Click **Yes**. The disk reverts back to a basic disk.
9. Close Computer Management, and leave Windows running for the next project.



Hands-On Project 9-6: Using the NTFS File System

Time Required: 25 minutes

Objective: Work with file and folder attributes and permissions.

Required Tools/Equipment: Your classroom computer with a volume designated as the D drive and formatted with NTFS

Description: In this project, you assign permissions in NTFS and set file attributes.

1. Log on to your computer as an administrator, if necessary.
2. Open the Computer window and double-click the **D** drive.

3. Create a new folder at the root of the D drive named **TestFiles**.
4. Right-click **TestFiles** and click **Properties** to see the folder's properties. Click the **Security** tab. You see a list of groups at the top and a list of permissions at the bottom (see Figure 9-19). Click each group to see the permissions assigned to it. Notice that the Users group has Read & execute, List folder contents, and Read permissions. These three permissions allow all users of the computer to view files and folders in the TestFiles folder, but they can't change them.

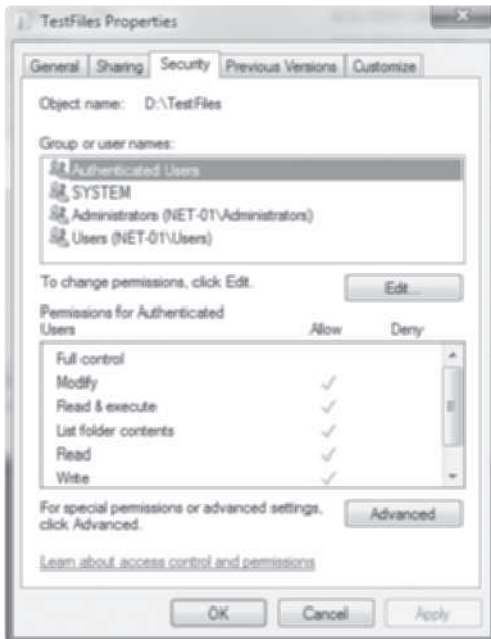


Figure 9-19 The Security tab showing NTFS permissions

Courtesy of Course Technology/Cengage Learning

5. Click **Edit**, and then click **Add** to add a user or group to the permissions list. In the Enter the object names to select text box, type **GuestUsers** and click **Check Names**. You created this group in Hands-On Project 9-2. Click **OK**.
6. Click **Guest Users** in the list of users, and then click the **Modify** permission in the Allow column of the permissions list. When you select the Modify permission, the Write permission is selected automatically. Click **OK**, and then click **OK** again to close the TestFiles Properties dialog box.
7. Double-click the **TestFiles** folder to open it, and create a new text document named **test1** in this folder. Right-click **test1** and click **Properties**. Click the **Security** tab. Notice that the list of groups and their permissions are the same as for the TestFiles folder. By default, all new files created in a folder automatically inherit the permissions of the folder in which they're created. To change inherited permissions, you must disable permission inheritance.

8. Click **Advanced**, and then click **Change Permissions**. If you want to disable permission inheritance, click to clear the **Include inheritable permissions for this object's parent** check box. For now, leave permission inheritance enabled. Click **OK** twice. Leave the Properties dialog box for test1 open.
9. Click the **General** tab, and click **Advanced** at the bottom.
10. You see the Compress and Encrypt attributes at the bottom of the Advanced Attributes dialog box. Click the **Compress contents to save disk space** check box. Now click the **Encrypt contents to secure data** check box. Notice that the Compress check box is cleared when you click the Encrypt check box because both can't be enabled at the same time. Click **OK** twice.
11. Next, you're prompted to encrypt the file and parent folder (which causes all files placed in the folder to be encrypted) or encrypt only the file. Click **Encrypt the file only**, and then click **OK**. The test1 filename is then displayed in green, indicating it's encrypted.
12. Now set the Compress attribute for test1. The filename is displayed in blue now.
13. Close all windows, and log off Windows for the next project.



Hands-On Project 9-7: Using the Linux Disk Utility

Time Required: 15 minutes

Objective: Create and delete volumes in Linux with Disk Utility.

Required Tools/Equipment: A Linux system with an unallocated disk

Description: In this project, you create and delete volumes in the Linux Disk Utility.

1. Log on to your Linux system.
2. Click **System**, point to **Administration**, and click **Disk Utility**. Click the first hard disk listed under Storage Devices (see Figure 9-20). The device is named `/dev/sda`. If your Linux installation uses a single disk, as in Figure 9-20, you see one large partition named `/dev/sda1` and a smaller extended partition named `/dev/sda2` that contains the swap space partition (`/dev/sda5`) Linux uses as virtual memory. Click the **Extended** partition to see the device name `/dev/sda2`, and then click the **Swap Space** partition to see its name.
3. Click the `/dev/sda1` partition. You'll see that the type is Ext4. Ext3 and Ext4 are the most common file systems on Linux volumes. Notice the mount point, which indicates the volume is mounted at `/` (the root of the file system).
4. Click the second hard disk listed under Storage Devices, and then click **Create Partition**. In the Size text box, type **500 MB**. Click the **Type** list arrow to view the types of file systems you can use. Linux supports several file systems, including NTFS and FAT. Click **Ext4**. Type **vol1** in the Name text box (making sure you use the numeral "1" for the last character), and then click **Create**.
5. If you're prompted for a password, enter your logon password and click **Authenticate**. The new volume is created, but it's not ready for use until you mount it in the file system. Click **Mount Volume**. By default, Ubuntu mounts the new volume in a folder named `\media\volume name` (with *volume name* representing the name you entered when you created the volume; see Figure 9-21).



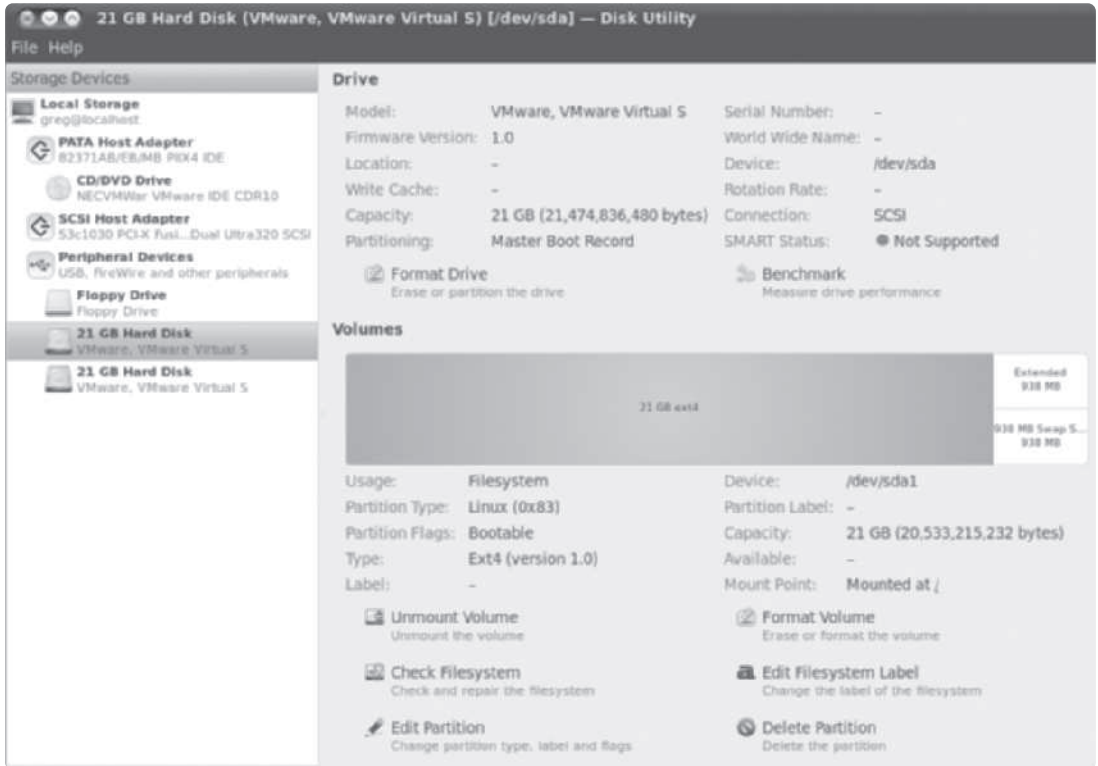


Figure 9-20 Linux Disk Utility
 Courtesy of Course Technology/Cengage Learning

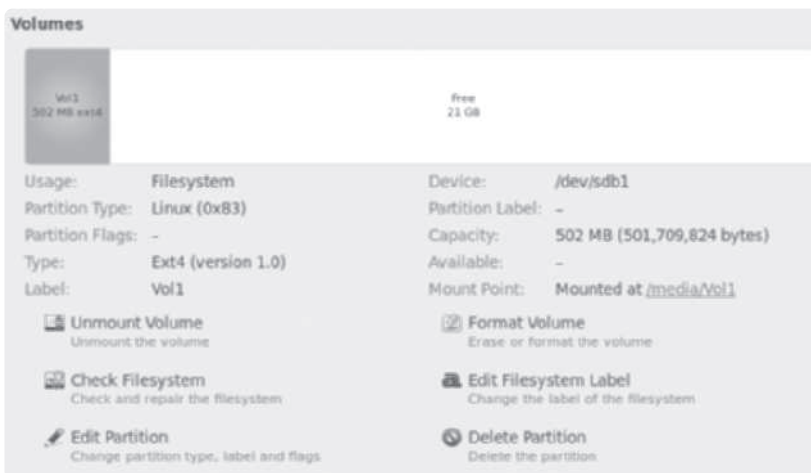


Figure 9-21 A new volume in Linux
 Courtesy of Course Technology/Cengage Learning

6. You can also mount a volume manually in a terminal window. Open a terminal window, type **sudo umount /media/vol1**, and press **Enter** to unmount the volume. Enter your password if prompted.
7. To create a folder where you'll mount the volume, type **sudo mkdir /vol1** and press **Enter**. Then mount the volume by typing **sudo mount /dev/sdb1 /vol1** and pressing **Enter**. To view the contents of this volume, type **ls /vol1** and press **Enter**. A system file called `lost+found` is displayed. Type **touch /vol1/newfile** and press **Enter** to create a file, and then type **ls /vol1** and press **Enter** to verify that the file was created.
8. To unmount `vol1`, type **sudo umount /vol1** and press **Enter**. Type **ls /vol1** and press **Enter** to verify that the file is no longer there because you unmounted the volume from this folder.
9. Shut down Linux.

The Linux File System

Linux supports a number of file systems, as you saw in Hands-On Project 9-8 when you selected the file system type. Among them are Ext3, Ext4, ReiserFS, and XFS. Ext3 (and, more recently, Ext4) is the default file system for most Linux distributions. Both file systems use a feature called journaling that ensures reliability by maintaining a record of changes to the file system before they're written to the disk. If a disk operation is interrupted because of a power failure or system crash, partial changes can be undone to prevent corrupting the file system.

Linux file systems support using permissions to control access to files and folders but much differently than in NTFS. In Linux, there are only three permissions—read, write, and execute—and three user types that can be assigned one or more of these permissions. The user types are as follows:

- *owner*—The owner of the file or folder, which is usually the user who created it
- *group*—The primary group to which the owner belongs
- *other*—All other users

Permissions are specified by using a single letter: *r* for read, *w* for write, and *x* for execute. For example, a file named `newfile` created by a user named `greg` who belongs to a group named `greg` is shown by using the `ls` command, as follows:

```
- rw- r-- r-- greg greg newfile
```

A few details are missing, but this line shows how permissions are displayed. The dash (`-`) in the first position indicates the file is a regular file. Folders or directories are indicated with a `d` in the first position. Permissions for each user type are displayed with three characters. The first three characters (`rw-` in this example) are the owner's permissions. The next three characters (`r--`) are the group permissions, and the last three (`r--`) are the permissions for all other users. The first name is the owner of the file (`greg`), and the next name is the Owner's primary group (`greg`). So to summarize, the permissions on `newfile` are as follows: Owner `greg` has read/write access, group `greg` has read access, and everybody else has read access.

The GUI in many Linux distributions shows permissions in a less cryptic form. Figure 9-22 shows permissions for the same file in the Linux GUI.



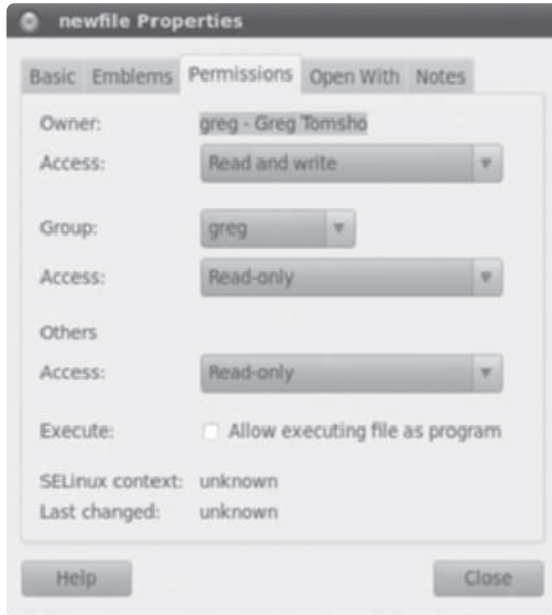


Figure 9-22 File permissions in the Linux GUI

Courtesy of Course Technology/Cengage Learning



Linux also supports many advanced file system features, such as disk quotas, encrypted files, and file compression.

Working with Shared Files and Printers

File and printer sharing is one reason businesses began to outfit computers with network interfaces and software. OSs and computers have evolved since the early days of stand-alone computers, when considerable effort was needed to make a lone PC part of a busy network. Most computers come with a NIC installed and an NOS that includes all the protocols for sharing files and printers.

The dominant file-sharing protocol is **Server Message Block (SMB)**, used by Windows and supported by Linux and MAC OS computers. SMB is the native Windows file-sharing protocol, and **Network File System (NFS)** is the native Linux file-sharing protocol. However, Linux supports SMB, and Windows can support NFS with the right software installed.



NFS support on Windows is available in Windows Server 2008 and the Ultimate and Enterprise editions of Vista and Windows 7. Third-party solutions are available for other Windows versions.

Printer sharing also uses SMB on Windows and Linux. The native Linux printer-sharing protocol is line printer daemon/line printer remote (LPD/LPR). LPD is the server side of a shared printer session, and LPR is the client component of the software.

Sharing Files and Printers in Windows

The first thing you need to know about file sharing in Windows is that users are subject to both share permissions and NTFS permissions when accessing files over the network. You learned about NTFS permissions earlier in the chapter. Thankfully, share permissions are somewhat simpler, and there are only three (see Figure 9-23):

- *Read*—Users can view contents of files, copy files, run applications and script files, open folders and subfolders, and view file attributes.
- *Change*—All permissions granted by Read, plus create files and folders, change contents and attributes of files and folders, and delete files and folders.
- *Full Control*—All permissions granted by Change, plus change file and folder permissions as well as take ownership of files and folders.

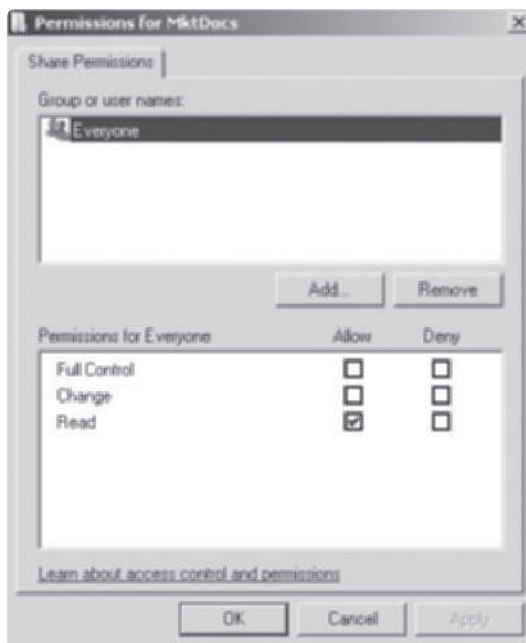


Figure 9-23 Viewing share permissions

Courtesy of Course Technology/Cengage Learning

Windows assigns default permissions depending on how a folder is shared. Generally, the default share permission is Read Everyone. On FAT volumes, share permissions are the only way to secure files accessed through the network.

Sharing files on the network, as you have seen in previous activities, isn't difficult in a Windows environment. Nonetheless, you should be familiar with some techniques and options before forging ahead with setting up a file-sharing server. You can use the following methods to configure folder sharing in Windows Server 2008. The procedures are similar in Windows 7 and Vista:

- *File Sharing Wizard*—To start this wizard, right-click a folder and click Share (or “Share with” in Windows 7). The File Sharing Wizard (see Figure 9-24) simplifies sharing for novices by using easier terms for permissions and by setting NTFS permissions to accommodate the selected share permissions. In Figure 9-24, the



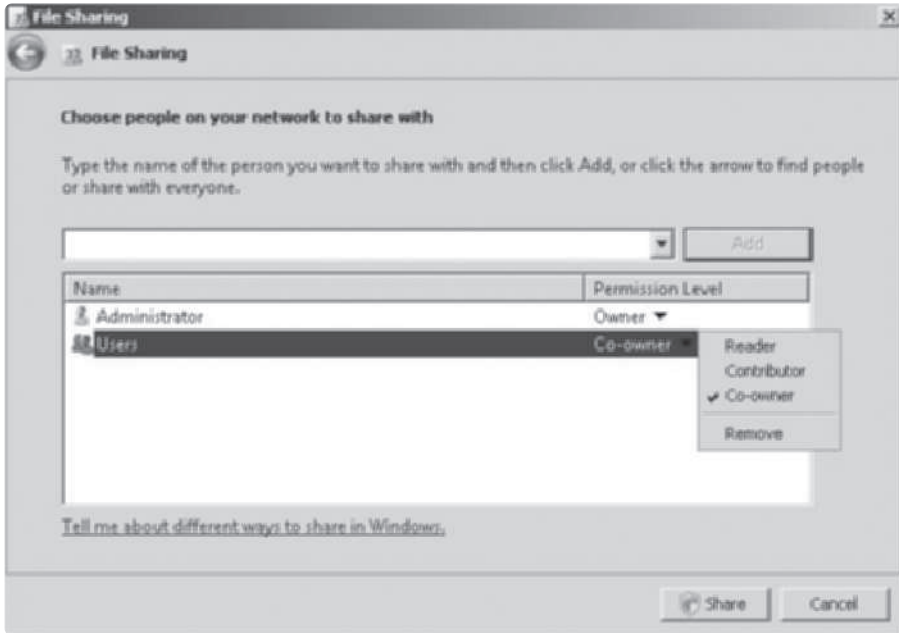


Figure 9-24 The File Sharing Wizard

Courtesy of Course Technology/Cengage Learning

permissions you see—Reader, Contributor, and Co-owner or Owner—correspond to the Read, Change, and Full Control share permissions, respectively.

- *Advanced Sharing dialog box*—To open this dialog box, click Advanced Sharing in the Sharing tab of a folder’s Properties dialog box. There are quite a few options in this dialog box (see Figure 9-25):

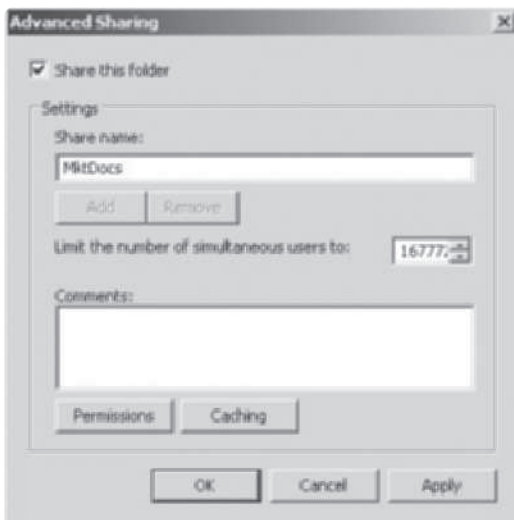


Figure 9-25 The Advanced Sharing dialog box

Courtesy of Course Technology/Cengage Learning

- **Share this folder:** Sharing can be enabled or disabled for the folder by clicking this check box.
- **Share name:** The share name is the name users see in the Network folder of Windows Explorer when browsing the server. To put it another way, it's the name you use to access the folder with the UNC path (`\\server\share name`). You can add or remove share names. A single folder can have multiple share names and different permissions, simultaneous users, and caching settings for each share name.
- **Limit the number of simultaneous users to:** In Windows Server 2008, the default limit is 16,777,216, which is, practically speaking, unlimited. In Windows client OSs, the maximum number of users who can access a share simultaneously is 10.
- **Comments:** You can enter a description of the share's contents and settings in this text box.
- **Permissions:** Click this button to open the Permissions dialog box shown previously in Figure 9-23.
- **Caching:** This option controls how offline files are configured. Offline files enable users to disconnect from the network and still have the shared files they were working with available on their computers.
- **Shared Folders snap-in**—You use this component of the Computer Management console to monitor, change, and create shares on the local computer or a remote computer. To create a new share, right-click the Shares node under the Shared Folders snap-in and click New Share. The Create a Shared Folder Wizard walks you through selecting the folder to share or creating a new folder to share, naming the share, configuring offline files, and setting permissions.
- **Share and Storage Management**—This snap-in is the most advanced method for creating shares. Like the Shared Folders snap-in, you use a wizard to create shares. You can select the folder to share or create a new share, configure NTFS permissions, and choose sharing protocols.



Hands-On Project 9-8: Sharing a Folder with the File Sharing Wizard

Time Required: 15 minutes

Objective: Create a test folder and then share it by using the File Sharing Wizard.

Required Tools/Equipment: Your classroom computer with a volume designated as the D drive and formatted with NTFS

Description: In this project, you try the File Sharing Wizard to see how it sets permissions automatically.

1. Log on to your computer as an administrator, if necessary.
2. Open the Computer window and double-click the D drive.
3. Create a new folder at the root of the D drive named **TestShare1**.
4. Open the TestShare1 folder's Properties dialog box, and click the **Security** tab. Make a note of the permissions assigned on this folder. Close the Properties dialog box.



5. Right-click **TestShare1**, point to **Share with**, and click **Specific people** to start the File Sharing Wizard.
6. Type **newguest1** in the text box, and then click **Add**. New Guest User 1 is added. Click the list arrow in the Permission Level column next to New Guest User 1, and make sure **Read** is selected.
7. Click **Share**. The UNC path for the share is displayed. Click **Done**.
8. Right-click **TestShare1** and click **Properties**. Click the **Sharing** tab, and then click **Advanced Sharing**.
9. Click **Permissions**. The Administrators and Everyone groups have Full Control permission to the share. The NTFS permissions, as you'll see in the next step, restrict New Guest User 1's permissions to only Read & execute, List folder contents, and Read, which effectively gives the user the ability to open and view the file. Click **Cancel** twice.
10. In the TestShare1 folder's Properties dialog box, click the **Security** tab. Click **New Guest User 1** and notice that the account's NTFS permissions are Read & execute, List folder contents, and Read.
11. Close all open windows.

Sharing Printers in Windows To understand how to work with and share printers in a Windows environment, first you need to understand the terminology for defining the components of a shared printer:

- *Print device*—The physical printer containing paper and ink or toner to which print jobs are sent. There are two basic types of print devices:
 - Local print device: A printer connected to an I/O port on a computer, with a parallel or USB cable, or through a TCP/IP port, which is used to access a printer attached directly to the network through the printer's NIC
 - Network print device: A printer attached to and shared by another computer
- *Printer*—The icon in the Printers folder that represents print devices. Windows programs print to a printer, which uses a printer driver to format the print job and send it to the print device or print server. A printer can be a local printer, which prints directly to a local or network print device, or a network printer, which prints to a print server.
- *Print server*—A Windows computer that's sharing a printer. It accepts print jobs from computers on the network and sends jobs to the printer to be printed on the print device.
- *Print queue*—A storage location for print jobs awaiting printing. In Windows, the print queue is implemented as a folder (by default, C:\Windows\System32\Spool\Printers) where files that make up each print job are stored until they're sent to the print device or print server.

A configured print server can perform a host of printing functions that aren't possible when users' computers print directly to a print device:

- *Access control*—Using permissions, administrators can control who can print to a printer and who can manage print jobs and printers.

- *Printer pooling*—A single printer represents two or more print devices. Users can print to a single printer, and the print server sends the job to the print device that's least busy.
- *Printer priority*—Two or more printers can represent a single print device. In this case, printers can be assigned different priorities so that jobs sent to the higher priority printer are sent to the print device first.
- *Print job management*—Administrators can pause, cancel, restart, reorder, and change preferences on print jobs waiting in the print queue.
- *Availability control*—Administrators can configure print servers so that print jobs are accepted only during certain hours of the day.

To configure a print server, you just need to share a printer. After a printer is installed, right-click it and click Sharing. The Sharing tab of a print server's Properties dialog box (see Figure 9-26) contains the following options:

- *Share this printer*—When this check box is selected, the print server is shared. By default, the Everyone special identity group is assigned Print permissions to shared printers.
- *Share name*—By default, it's the name of the print server in the Printers folder. You can enter a shorter share name or one that's easier to remember.



Figure 9-26 The Sharing tab for a print server

Courtesy of Course Technology/Cengage Learning



- *Render print jobs on client computers*—When this check box is selected (the default setting), client computers process the print job and send it to the print server in a format that’s ready to go directly to the print device. If this option isn’t selected, more processing occurs on the print server.
- *List in the directory*—When this check box is selected, the print server is displayed in Active Directory and can be found by Active Directory searches. By default, this option isn’t selected. This option is only available on computers that are domain members.
- *Additional Drivers*—When a client connects to a shared printer, the printer driver is downloaded to the client from the server automatically when possible. You can click this button to install different printer drivers on the server to support different Windows versions.

Sharing Files and Printers in Linux

Linux supports Windows file sharing by using SMB in a software package called Samba. Depending on the Linux distribution, you might have to install this component. In Ubuntu Linux, Windows file sharing isn’t installed by default. If you try to share a folder, you’re prompted to install the Windows networks sharing service. If you choose to install the service, the Samba package is installed.

After Samba is installed, you can right-click a folder and click Sharing options. To enable the share, click the Share this folder check box, as shown in Figure 9-27. You can allow others to create and delete files in the folder or enable guest user access.



Figure 9-27 Linux folder sharing

Courtesy of Course Technology/Cengage Learning

Printer sharing in Linux is straightforward after Samba has been installed. When you create a new printer in Linux, it’s shared automatically. You might need to enable printer publishing in the printer’s Server Settings dialog box (see Figure 9-28). In Linux, “printer publishing” is just another name for printer sharing.

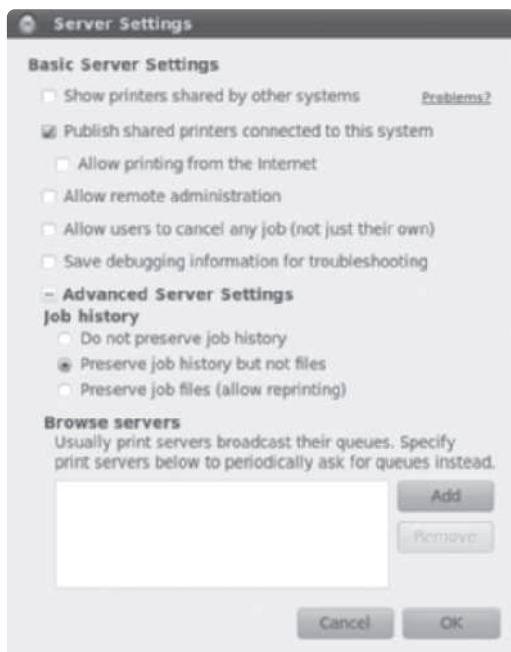


Figure 9-28 Linux print server settings

Courtesy of Course Technology/Cengage Learning



Monitoring System Reliability and Performance

Windows Server 2008 provides tools to manage and monitor server operation and resources, including the following:

- Task Manager
- Event Viewer
- Performance Monitor
- Windows System Resource Manager

All these tools are available in Windows Server 2008, and all but Windows System Resource Manager are available in Windows 7 and Vista. You have already used Task Manager, so this section focuses on Event Viewer, Reliability and Performance Monitor, and Windows System Resource Manager.

Event Viewer

Administrators use Event Viewer to examine event log entries generated by system services and applications. A typical event log can contain hundreds or thousands of events, but usually, administrators are interested in events that indicate a problem. Events are categorized by these levels:

- *Information*—Indicated by a blue “i” in a white circle, these events are normal operations, such as service stops and starts.
- *Warning*—Indicated by a black exclamation point inside a yellow triangle, warnings provide information about events that should be brought to the administrator’s attention. Warnings aren’t necessarily an indication of a problem but often indicate a condition that can lead to a more serious error.
- *Error*—Error events, indicated by a white exclamation point inside a red circle, are often generated when a process or service is unable to perform a task or stops unexpectedly. Error messages should be addressed immediately, as they indicate a configuration error or an operational problem.

You can examine several log files in Event Viewer (see Figure 9-29), including the Application, Security, Setup, and System logs. In addition, some applications and services have their own log files. You can click the Level column header to sort events and group them by level to spot the most serious events easily. When an event is selected, descriptive information about it is displayed in the bottom pane of the General tab. The Details tab shows additional technical data about the event. For many events, you can also click the Event Log Online Help link in the General tab to get more information.

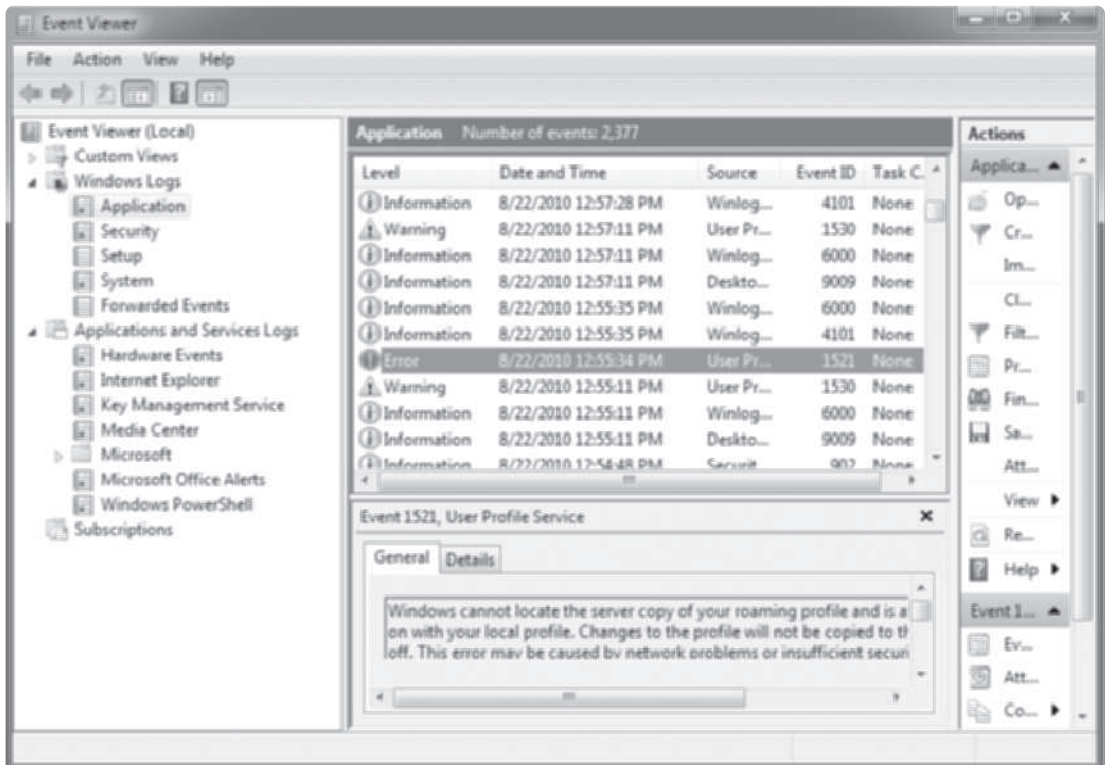


Figure 9-29 Event Viewer in Windows 7

Courtesy of Course Technology/Cengage Learning

Performance Monitor

Performance Monitor consists of a collection of tools for pinpointing which resources are being overloaded and how they're being overloaded. You open it from the Administrative Tools folder. Performance Monitor contains the following folders:

- *Monitoring Tools*—Contains the Performance Monitor tool
- *Data Collector Sets*—Contains user- and system-defined templates with sets of data points called data collectors
- *Reports*—Contains system- and user-defined performance and diagnostic reports



Performance Monitor is called Performance and Reliability Monitor in Windows Server 2008 R1 and Windows Vista. It was renamed Performance Monitor in Windows Server 2008 R2 and Windows 7, and the look was also changed. This book uses screenshots from Windows 7.

Performance Monitor, accessed under the Monitoring Tools folder, uses counters to track the performance of a variety of objects. Performance can be tracked in real time or scheduled for later review and analysis. A counter is a value representing some aspect of an object's performance. For example, disk drives have counters representing the percent of time the disk is used for read operations and the number of disk requests waiting to be serviced, among many others. There are counters for almost every hardware and OS component on a server, including, of course, directory services.

Performance Monitor can track counters with a line graph (the default), with a histogram (bar graph), or as raw data saved to a report. To use Performance Monitor in real-time mode, you simply add counters to the selected graph or report. You can add as many counters as you like, but as you can see in Figure 9-30, the display can get crowded.

Performance Monitor has two modes. You can display counters in real time, or you can open a saved performance log file and view data that has been captured over a period of time. To create a performance log, you create a new data collector set or start a saved data collector set. After the data collector set has finished running, you can view collected performance data in Performance Monitor.

Collecting Baseline Performance Data Viewing performance data in real time is helpful if you want to see the impact certain actions have on selected counters. For example, you might want to see the effect a large network file copy has on CPU and network utilization. Real-time monitoring of performance counters can also be useful for tracking the cause of a sluggish system. Unless you have a good idea what part of the system to examine, however, finding the cause of the problem can be a hit-and-miss proposition.

One reason that tracking causes of poor performance with real-time monitoring is difficult is that you have no point of reference for comparing data. This point of reference, called a performance baseline (or simply a **baseline**), is a record of performance data gathered when a system is performing well under normal operating conditions. Generally, baseline data is collected shortly after a system is put into service and then again each time changes are made, such as installing or removing a network service or application, or when many new



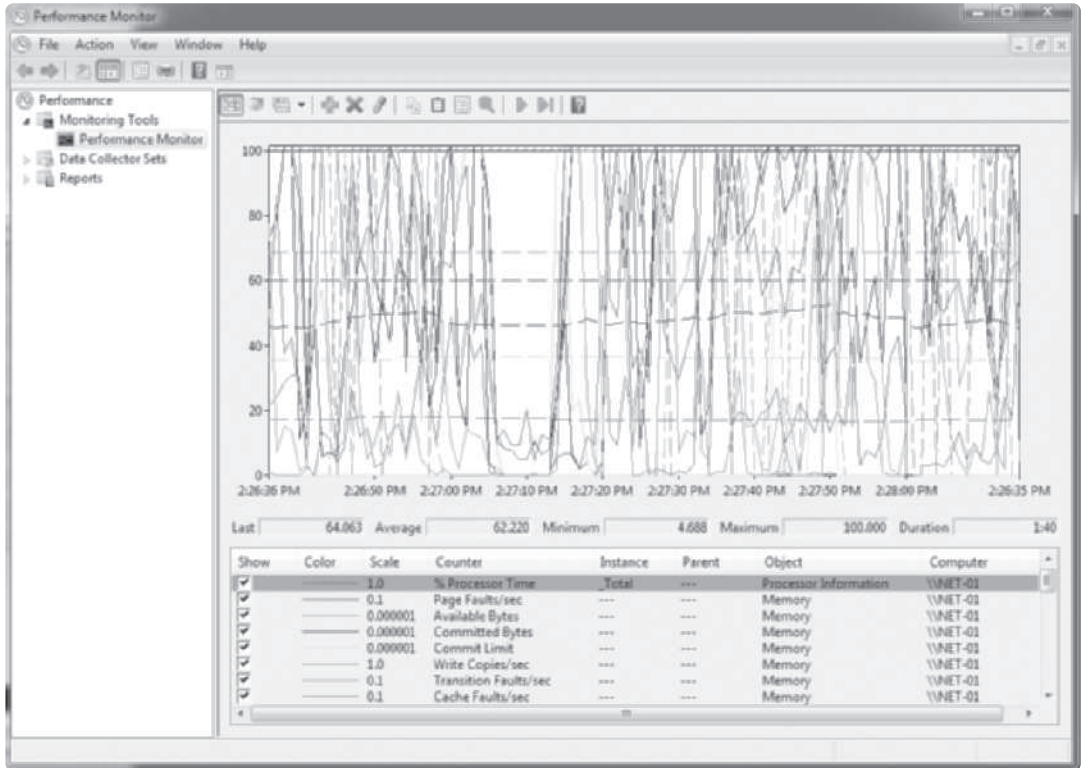


Figure 9-30 Performance Monitor with too many counters

Courtesy of Course Technology/Cengage Learning

users are using the system. The baseline data collected during normal operation conditions can then be compared with data collected during peak resource demands to give you insight into your system's capabilities and limitations.

To create a baseline of performance data, you create a **data collector** set that specifies the performance counters you want to collect, how often to collect them, and the time period. You can create multiple data collector sets that capture different aspects of system performance and measure performance during different time periods. For example, if you know a database application is used heavily between 10:00 a.m. and 3:00 p.m., you can collect CPU, disk, memory, and network performance data during that time period. You should also collect data for critical resources over an entire day so that you can spot usage trends.

Be aware that performance monitoring uses system resources. It takes memory to run Performance Monitor, CPU cycles to collect and display counter data, and disk resources to update log files. With Performance Monitor, however, you can select a remote computer as the target for monitoring. By monitoring remotely, you lessen the monitoring session's impact on the computer being monitored. You can also adjust the counter sampling interval to collect counter data less frequently than the default values. The more often counter data is collected, the more impact the monitoring session has on system resource use.



Hands-On Project 9-9: Viewing Real-Time Performance Data

Time Required: 10 minutes

Objective: Add counters to Performance Monitor to view real-time performance data.

Required Tools/Equipment: Your classroom computer

Description: You have installed a server recently and want to look at key performance indicators to be sure the server is handling its current load. First, you look at performance data in real time. In the next activity, you save data in a log to view later.

1. Log on to your computer as an administrator.
2. Click **Start**, point to **Administrative Tools**, and click **Performance Monitor**. The initial view in the right pane is Overview of Performance Monitor, which provides information on using Performance Monitor and a system summary. There's also a link to Resource Monitor.
3. Click **Open Resource Monitor** to see a tabular and graphical view of CPU, Disk, Network, and Memory use. Click the **CPU**, **Memory**, **Disk**, and **Network** tabs to view additional performance details for each resource. Resource Monitor can also be opened from Task Manager. Close Resource Monitor.
4. Scroll through the System Summary pane, which shows key performance indicators in tabular form. Click **Performance Monitor** in the left pane.
5. Click the **Add** toolbar icon (green plus sign) to open the Add Counters dialog box. You can specify the computer where you want to add counters in the Select counters from computer list box. For now, leave the setting as <Local computer>.
6. Click the **Show description** check box at the bottom so that you can see descriptions of counters you select.
7. Scroll through available counters to see what type of data can be monitored. Click **PhysicalDisk**. The Instances of selected object list box displays the physical disk objects you can select. You can monitor just one disk or multiple disks or add a counter representing the total counter values for all disks. Click **0 C: D:** (assuming disk 0 contains the C and D drives), click the **Add** button, and then click **OK**.
8. Notice that several counters have been added to Performance Monitor. In the bottom pane, click the **Avg. Disk Bytes/Transfer** counter. Emphasize it in the display by clicking the **Highlight** toolbar icon (looks like a yellow highlighter pen). If the counter isn't showing much activity, create some activity by opening and then closing Internet Explorer.
9. Right-click **Avg. Disk Bytes/Transfer** and click **Remove All Counters**. When prompted to confirm, click **OK**.
10. Click the **Add** toolbar icon. In the Add Counters dialog box, click to expand **PhysicalDisk**. To select a counter for PhysicalDisk, click **% Disk Time**. (If necessary, verify that the **Show description** check box is selected. You might need to check it whenever you open this dialog box.) Read the description of the counter, and then click **Add**.
11. Click to expand **Processor**, and click **% Interrupt Time**. Read the description, and then click **Add**. Click to expand **System**, and click **Processor Queue Length**. Read the



description, and then click **Add**. Queue counters indicate how many activities are waiting for work to be done. For most objects with queue counters (such as PhysicalDisk and Network Interface), a sustained queue value of more than a few items in the queue often indicates a bottleneck. Click **OK**.

- In the bottom pane of Performance Monitor, notice that the value in the Scale column of the Processor Queue Length counter is 10. This value means the graph is showing the counter's actual value multiplied by the scale value. In this case, if the graph shows a value of 30 for Processor Queue Length, the actual value is 3. To adjust the scale, right-click **Processor Queue Length** and click **Properties**. In the Data tab, you can select the color, width, style, and scale of the line graph for the counter. Click the **Scale** list arrow, click **1.0** in the list, and then click **OK**.



You might want to change the scale value for a counter so that the line on the graph is more distinct and shows variations in the counter value more clearly.

TIP

- Keep Reliability and Performance Monitor open for the next activity.



Hands-On Project 9-10: Creating a Data Collector Set

Time Required: 15 minutes

Objective: Create a custom data collector set.

Required Tools/Equipment: Your classroom computer

Description: You want to create a performance baseline for your server, so you decide to create a data collector set.

- Log on to your computer as an administrator and open Performance Monitor, if necessary.
- In the left pane, click to expand **Data Collector Sets**. Right-click **User Defined**, point to **New**, and click **Data Collector Set**. In the Name text box, type **SysPerformance1**. Verify that the default **Create from a template** is selected, and then click **Next**.
- In the Which template would you like to use? window, click each template in the Template Data Collector Set list box, and read its description. Click **System Performance** to select this template, and then click **Next**.
- In the next window, you can change the default path where data is saved. For now, leave the default location, and then click **Next**.
- In the Create the data collector set window, you can change the user account for running the data collector set. Leave the default setting of **<Default>** in the Run as text box, and then click **Finish**.
- Notice that the new data collector set has been created and its status is **Stopped**. Right-click **SysPerformance1** and click **Properties**.
- Click the **Schedule** tab, where you can create a schedule of when you want the data collector set to run. Click **Add**. You can choose a beginning date and an expiration date. If you choose an expiration date, the data collector set stops collecting data after that date. You can also specify a start time and the days of the week the data

- collector set should run. You're going to start this data collector set manually, so click **Cancel**.
8. Click the **Stop Condition** tab, where you specify the duration for running the data collector set. If no conditions are selected, the data collector set runs until it's stopped manually. Accept the default value of 1 minute in the Overall duration text box, and then click **OK**.
 9. In Performance Monitor, right-click **SysPerformance1** and click **Start**. A green arrow on the data collector set icon indicates it's running. Open and close Internet Explorer and Wireshark several times to generate some resource use.
 10. When the status returns to Stopped, right-click **SysPerformance1** and click **Latest Report**. You see a performance report similar to Figure 9-31. In the Diagnostic Results section, counters with suspect values are flagged as warnings. In this example, a high CPU load was detected.

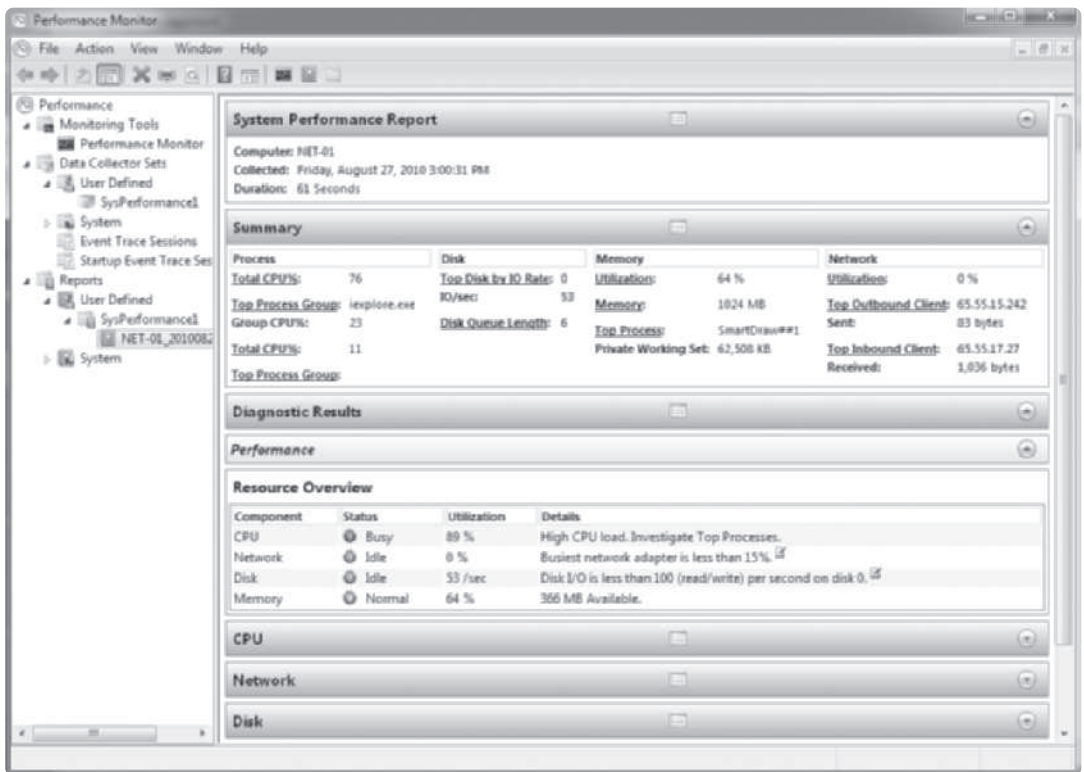


Figure 9-31 Performance report generated from a data collector set

Courtesy of Course Technology/Cengage Learning

11. Scroll down in the report to view sections for major hardware systems, such as CPU, Network, Disk, and Memory. Click to expand each section to view more detailed information. You can get quite a bit of information from this report.
12. Close all open windows.

Windows System Resource Manager

Windows System Resource Manager (WSRM) is a Windows Server 2008 feature installed in Server Manager that helps you manage processor and memory resources on heavily used systems. By managing resources, you can give high-priority services and applications a larger share of CPU time and memory to make sure they can perform critical tasks in a timely manner. You can also fine-tune resource use so that each process has a fairly equal share of resources to ensure that no one service dominates CPU and memory use. WSRM includes the following features for managing running services efficiently:

- Preconfigured and custom policies that allocate resources on a per-process or per-user basis
- Policies based on calendar rules to allow fine-tuning system resource use according to time of day
- Automatic policy application based on server events or changes in memory or CPU resources
- Resource monitoring data stored in a Windows internal database or an SQL database

A more detailed discussion of this advanced performance-tuning tool is beyond the scope of this book, but you can learn more in *MCTS Guide to Microsoft Windows Server 2008 Active Directory Configuration* (Course Technology, 2010, 1423902351).

Backup and Fault Tolerance

Performing regular backups is a necessary but often disliked task for many network administrators. Using tape to back up servers is still a common practice, but external disks connected through USB or eSATA and network storage are becoming more popular because of their speed and convenience. However, if your company has a policy that backup media must be stored off-site periodically, tapes are probably the best solution.

Regular backups provide a safety net to restore a system to working order in the event of a disk failure or file corruption. They also allow you to restore files that were accidentally deleted or older versions of modified files that you might need. Traditional backup requires using the backup program to restore files. If system files are lost or damaged to the extent that the system can no longer boot, you have to reinstall the OS before you can use the backup program to restore the system. Another popular type of backup is an image backup, in which a copy of an entire disk is created that can be restored without reinstalling the OS. With many image backups, however, you can't restore separate files, so image backups are usually done along with traditional file backup.

Fault tolerance provides methods for a system to continue running after a system failure has occurred. System failures can be power failures, disk failures, and entire computer

failures. Fault tolerance isn't a replacement for backups but complements a regular backup routine. For example, a fault-tolerant disk system allows a system to keep functioning after a disk fails, but you still need backups for situations involving deleted files or a corrupt file system.

Windows Backup

Windows Server Backup comes with Windows Server 2008 and has the following features:



These features refer to Windows Server Backup in Windows Server 2008 R2. The backup program in the first release of Windows Server 2008 is more limited in its capabilities.

- Backups can be run manually or scheduled to run automatically.
- You can create a system recovery backup that automatically includes all volumes containing critical system data, such as the volume with the Windows folder and the volume with the Active Directory database and log files.
- Manual backups can be stored on network drives, fixed and removable basic disk volumes, and CDs/DVDs. Tape drives aren't supported.
- Backups can be stored on a hard disk dedicated for backups, a nondedicated volume, or a shared network folder. Microsoft recommends using a dedicated disk for backup.
- You can use a Volume Shadow Copy Service (VSS) backup, which means even open files can be backed up.
- By default, Windows Server Backup is configured to back up the local computer, but you can also connect to another computer to back up files remotely.

Although Windows Server Backup is a fine tool for backing up servers, you should be aware of its limitations. It's not a substitute for an enterprise-class backup program, such as Symantec NetBackup and CommVault Galaxy Backup and Recovery; both offer advanced disaster recovery solutions. These programs are called for when you need a comprehensive backup and recovery solution for a large number of servers and servers distributed across multiple sites. Most of these products use a distributed backup strategy, in which backup agents are installed on servers and workstations throughout the enterprise and are controlled by a management console. Network and server administrators who manage large networks should familiarize themselves with these products, but a detailed discussion of them is beyond the scope of this book.

Windows 7 backup is called, simply enough, Backup and Restore (see Figure 9-32) and has straightforward features. You can use it to create a system image, create a system repair disc, and back up all files or separate files and folders. When you click the Set up backup link, Backup and Restore creates a scheduled backup that occurs weekly by default, but you can change the time, day, and frequency settings.



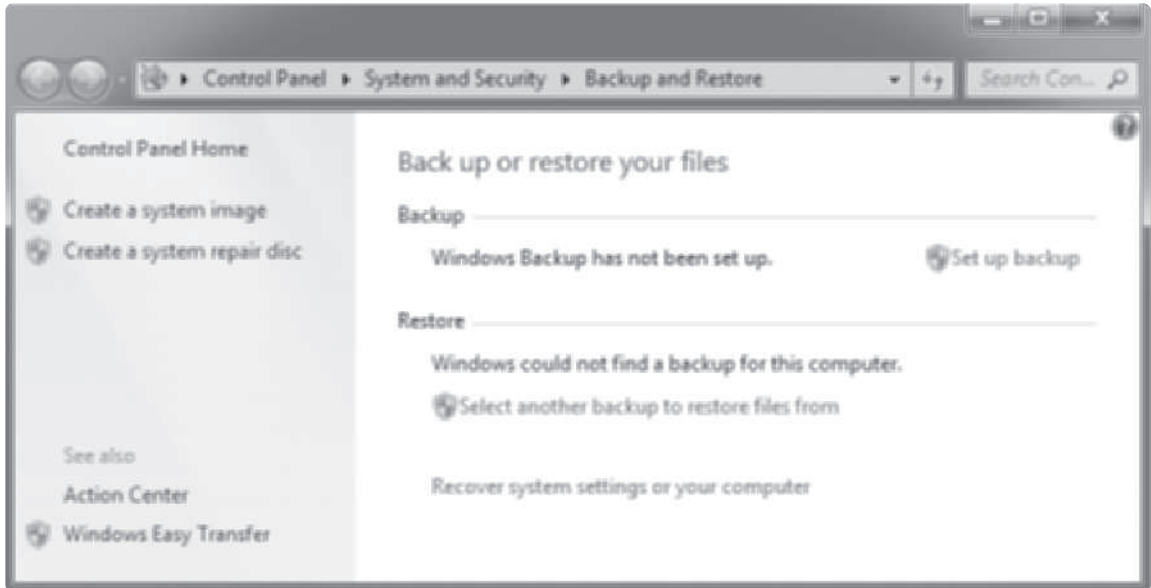


Figure 9-32 Backup and Restore in Windows 7

Courtesy of Course Technology/Cengage Learning



Hands-On Project 9-11: Using Windows 7 Backup and Restore

Time Required: 15 minutes

Objective: Explore the Backup and Restore program and set up a backup.

Required Tools/Equipment: Your classroom computer

Description: In this project, you run the Backup and Restore program in Windows 7 and explore the options for backing up your system. Then you create a small backup set and restore files from it.

1. Log on to your computer as an administrator.
2. Create a text document on your desktop named **testbackup**. This file will be included in the backup you create, and you'll delete and later restore this file. Double-click **testbackup** to open it in Notepad, and type your name so that the file isn't empty. Save the file and exit Notepad.
3. Click **Start**, point to **All Programs**, click **Maintenance**, and click **Backup and Restore**.
4. Click **Create a system image**. Read the description of a system image. Notice that you have the option to store the backup on a hard disk, a DVD, or a network location. Because you might not have a hard disk capable of storing the backup, click **Cancel**. An external USB hard drive is ideal for storing a system image.
5. Click **Create a system repair disc**. A system repair disc allows you to boot your computer in the event of a boot disk failure and contains recovery tools to solve problems

- or restore from a system image. You need a blank CD or DVD to create a system repair disc. Click **Cancel**.
6. Click **Set up backup**. In the Select where you want to save your backup window, click the **D:** drive. Click **Next**.
 7. In the What do you want to back up? window, click **Let me choose**, and then click **Next**.
 8. Leave the items under Data Files selected, but click to clear the **Include a system image of drives** check box (see Figure 9-33). The Data Files option backs up your document files, including those on your desktop. Click **Next**.

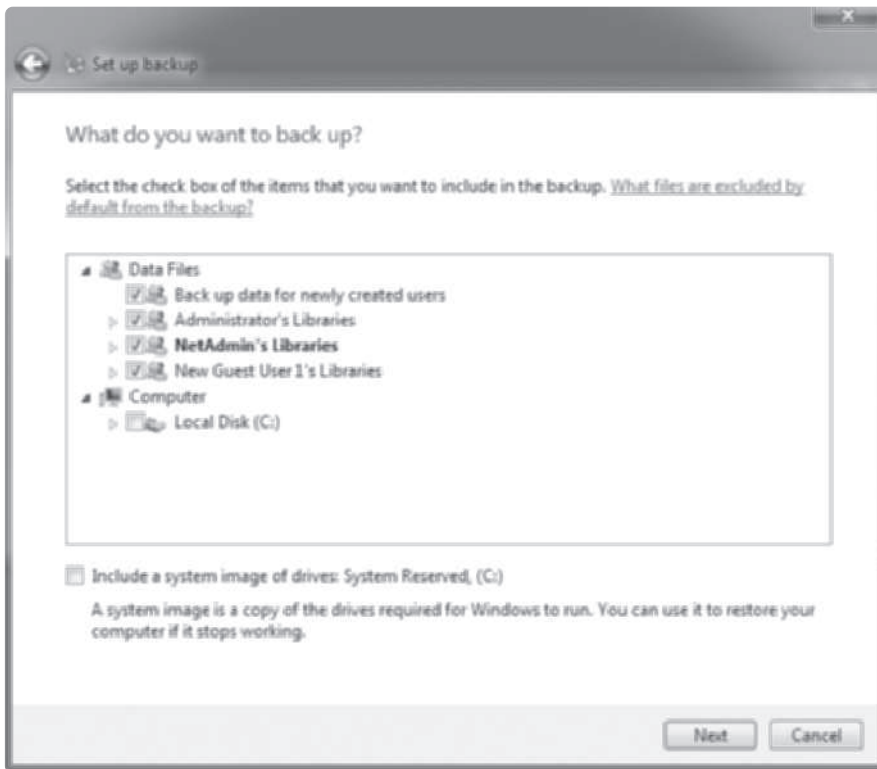


Figure 9-33 Choosing the files you want to back up

Courtesy of Course Technology/Cengage Learning

9. In the Review your backup settings window, notice that Backup and Restore automatically scheduled a backup to occur weekly. Click **Save settings and run backup**. To view the progress of the backup, click **View Details**.
10. When the backup is finished, you see a Restore section in Backup and Restore. Delete the **testbackup** file you created in Step 2. Empty the Recycle Bin so that the file is really deleted. Click **Restore my files** in Backup and Restore.
11. Click **Browse for files**. Click the logon name for your account. For example, if you're logged on as Bill Smith, you see a folder named "Bill Smith's backup" in the left pane

of the Browse the backup for files window. Double-click the **Desktop** folder in the right pane (see Figure 9-34).

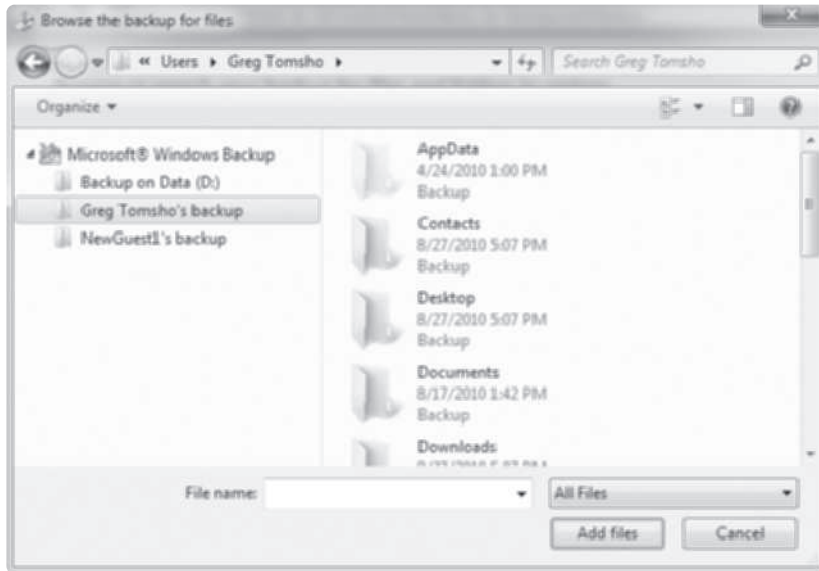


Figure 9-34 Browsing for files to restore

Courtesy of Course Technology/Cengage Learning

12. Click **testbackup** and click the **Add files** button. Click **Next**. Leave the default option “In the original location” selected, and click **Restore**. The file is then displayed on your desktop. Click **Finish**.
13. Close all open windows.

Protecting Data with Fault Tolerance

Some networks can’t afford downtime for their network servers. A disk crash, motherboard failure, or power outage can stop a server dead in its tracks if no fault-tolerant systems have been put in place. This section discusses three forms of fault tolerance that are common on networks and servers:

- Redundant power supply and uninterruptible power supply
- Redundant disk systems
- Server clustering

Redundant Power A computer requires a consistent, clean source of power. If power is interrupted for even the briefest moment, a computer is likely to reboot. On a desktop system, a user might lose some work, or if a file transaction was in progress, disk corruption could result. On a server, dozens of people could lose work, databases could become

corrupted, and the file system could be damaged. Power fluctuations can cause even worse damage because if the power spikes or sags, the motherboard and other hardware components can be damaged. To combat these potential problems, server systems and some desktop systems use redundant power supplies and uninterruptible power supplies.

A **redundant power supply** is essentially a second power supply unit in the computer case. Each unit is capable on its own of maintaining adequate power to the computer, so if one power supply fails, the other unit takes on the full load. When both units are operating, they share the power load, which also reduces heat and stress on each unit, further increasing reliability. Redundant power supplies can add considerable cost to a computer and are generally found only in servers.

An **uninterruptible power supply (UPS)** is a device with a built-in battery, power conditioning, and surge protection. A UPS is plugged into the wall outlet to charge the battery continuously, and the computer and monitor are plugged into outlets on the UPS. If the power fails, the UPS battery provides enough power to keep your computer and monitor running until main power is restored or you can shut down the computer safely. Most UPSs come with software and a USB or serial connection to the computer so that the UPS can communicate with the computer. The UPS can inform the computer when main power has been lost, when it has been restored, and when the UPS battery power is running low. You can configure the software so that if the amount of remaining battery time falls below a certain number of minutes, the system shuts itself down. UPSs also protect systems from power sags or brownouts, in which the main power voltage output falls below what's required to power off a system safely.

UPSs come in two main categories: online and standby. A standby UPS supplies power to plugged-in devices by passing power from the wall outlet directly to the device. In a power outage, a standby UPS detects the power failure and switches to battery power quickly. Unfortunately, if the switchover doesn't happen fast enough, the plugged-in devices might lose power long enough to reboot or cause a malfunction. An online UPS supplies power continuously to plugged-in devices through the UPS battery, which is recharged continually by the wall outlet power. In a power outage, there's no need to switch to battery power because the UPS is already supplying power from the battery. Overall, an online UPS is a far better solution for computer equipment but costs more.

Battery backup isn't the only advantage of UPSs. Power conditioning and surge protection are equally important to the sensitive components in computers. **Power conditioning** "cleans" the power, removing noise caused by other devices on the same circuit (such as fans, motors, and laser printers). **Surge protection** protects the computer from voltage spikes or surges—conditions that can be caused by lightning strikes, problems with the electric company's power transformer, or switching on large appliances, such as air conditioners.

Redundant Disk Systems Hard drives contain some of the few moving parts in a computer, making them more susceptible to failure than most other components. Redundant disk systems can prevent data loss in the event of a disk failure; in fact, a system with a redundant disk configuration can continue operating with no downtime. Redundant disk systems are based on the redundant array of independent disks (RAID) technology, introduced in Chapter 8. The two most common RAID configurations are disk mirroring (RAID 1) and disk striping with parity (RAID 5), discussed in the following sections.



RAID 1: Disk Mirroring Disk mirroring requires two disks. When data is written to one disk, it's also written to the second disk, thus creating a synchronized copy. If either disk fails, the system can continue operating because both disks have the same data. RAID 1's obvious disadvantage is that you have to purchase two disks of equal size but get the storage space of only one disk. Another disadvantage is somewhat slower performance, but it's a small price to pay for peace of mind. Write performance is slower because two disks must be written, but read performance is about the same as in a single-disk configuration. Both read and write performance can be enhanced by using two disk controllers: one for each disk, a configuration called "disk duplexing."

Most disk controllers on servers and even on many desktop systems can be configured for RAID 1. Windows Server 2008 supports RAID 1 in software if the disk controller lacks RAID support.



Disk performance depends on many factors, including the quality of the disk controller. On some systems, you might not see slower disk write performance with RAID 1, and you might see somewhat faster read performance compared with a single-disk configuration.

RAID 5: Disk Striping with Parity Disk striping with parity requires a minimum of three disks but is more space efficient than RAID 1. RAID 5 works by spreading data across multiple disks and using one disk in each write operation to store parity information. Parity information is generated by a mathematical calculation on data being written, so if one of the disks fails, this information can be used to re-create lost data from the failed disk.

For example, if a RAID 5 configuration consists of three disks with a cluster size of 64 KB and a file of 128 KB is written to the disk, 64 KB is written to the first disk, 64 KB is written to the second disk, and parity information is written to the third disk. Parity information isn't always written to the same disk, however. The next file written uses the second and third disks for file data and the first disk for parity (see Figure 9-35).




Figure 9-35 RAID 5: stripe set with parity

Courtesy of Course Technology/Cengage Learning

The number of disks in a RAID 5 configuration is theoretically unlimited, but the more disks there are, the more likely that more than one disk will fail simultaneously, and RAID 5 can recover from only a single disk failure. With RAID 5, you lose $1/n$ of your total disk space (with n representing the number of disks). So if your RAID setup consists of four 100

GB disks, you have 300 GB for disk storage and 100 GB for parity information. RAID 5 write performance is hampered because of calculating and writing parity information, but read performance is usually as good or better than with a single disk. Windows Server 2008 also supports RAID 5 in software if the disk controller doesn't.

Other RAID configurations exist, including RAID 0. RAID 0 is disk striping, requires two or more disks, and is used primarily for high performance. However, it provides no fault tolerance, so if any disk fails, all data is lost. Table 9-3 lists the RAID levels, some of which are no longer in use.

Table 9-3 RAID levels

RAID level	Description and use
RAID 0	Called disk striping, RAID 0 distributes data over two or more disks. RAID 0 has no fault-tolerance capabilities but does enhance disk performance. Available in Windows Server 2000/2003 as well as Linux.
RAID 1	Applies to disk mirroring and disk duplexing, in which two drives are exact copies of each other, and failure of the primary drive causes the secondary drive to take over automatically. Available in Windows Server 2000/2003 as well as Linux.
RAID 2	Uses separate check disks, in which data bits are striped on both data and check disks, to replace information from a damaged data or check disk in the array. Because check data requirements are high and require multiple separate drives for data, this form of RAID is seldom used. Not available for Windows OSs.
RAID 3	Uses a single check disk for parity information (sometimes called a parity disk for that reason) for each group of drives. Because the same size chunk of data is read or written each time the array is accessed, space allocation on these drives isn't efficient, especially for small files. Not available for Windows OSs.
RAID 4	Works much like RAID 3 but uses block or sector striping so that a single block or sector can be accessed at a time, instead of requiring all drives in the set to be accessed. Inefficient for writing data because check writes must occur immediately after data writes. Not available for Windows OSs.
RAID 5	Divides parity information across all drives in the RAID array so that each drive can be reconstructed from parity information stored on all other drives in the set. This array type, also called disk striping with parity, is available for Windows Server 2000/2003. Can withstand a single disk failure and continue to operate with no loss of data.
RAID 6	Like RAID 5, except two sets of parity information are written to two different disks. Can survive the failure of two disks rather than just one.
RAID 1+0	Combines RAID levels 1 and 0 and is sometimes called RAID 10. It uses a RAID 1 mirror set that's striped with another set of drives and can survive multiple disk failures in some circumstances. A minimum of four disk is required for this configuration.
RAID 0+1	Combines RAID levels 0 and 1 but starts with a stripe set and then mirrors it. RAID 0+1 can withstand two disk failures if both failed disks are from the same stripe set. A minimum of four disks is required.
Raid 5+0	A RAID 0 striped across two or more RAID 5 sets; also called RAID 50. Requires at least six disks and improves performance over traditional RAID 5. Can withstand up to four disk failures, depending on whether each failed disk is in a different RAID 5 set.

Server Clustering Windows Server 2008 supports server clustering in the Enterprise and Datacenter editions. A server cluster is made up of two or more servers that are interconnected and appear as a single unit. Two common types of clustering are failover and load-balancing. A failover cluster is used to provide fault tolerance in the event of a



complete system failure. It involves two or more servers sharing a high-speed link used to synchronize data continuously. One server is designated as the primary server, and the others are standby servers. If the primary server fails, a standby server takes its place as the primary server. A failover cluster is usually used to ensure high availability for applications such as database and messaging systems, where data is read and written frequently.

A load-balancing cluster consists of two or more servers that appear as a single unit to users. All servers in the cluster operate and share the server load. This type of cluster is used with read-intensive applications that exceed a single server's performance capabilities and when data on the server is comparatively static. The advantage of load-balancing clusters is that you still have the benefit of failover, in that if one server fails, the others just have to shoulder a bigger load.

Chapter Summary

- User accounts are the link between real people and network resources. They have two main functions in a network: Provide a way for users to authenticate themselves to the network, and provide detailed information about users.
- User accounts and passwords should follow naming conventions for their creation. Typical password restrictions include minimum length, complexity requirements, and how often they must be changed.
- Group accounts are used to organize users so that assignment of resource permissions and rights can be managed more easily than working with dozens or hundreds of user accounts. Groups in Active Directory are assigned a scope (domain local, global, or universal) and type (security or distribution).
- A user profile is a collection of a user's personal files and settings that define his or her working environment. A profile contains personal data folders a user maintains as well as files and folders containing user and application settings. A local profile is stored on the system the user logs on to. A roaming profile is stored on a network server and follows the user wherever he or she logs on.
- Locally attached storage is a device, such as a hard disk, that's connected to a storage controller on the server. Storage is divided into volumes or partitions. A file system is installed on the volume. NTFS is the file system used in Windows in most cases; it provides advanced features such as disk quotas, permissions, file compression, and encryption.
- Linux file systems, which support permissions and journaling, include Ext3, Ext4, ReiserFS, and XFS.
- SMB is the Windows default file-sharing protocol; NFS is the native Linux file-sharing protocol. Linux also supports SMB via Samba, and Windows supports NFS. Files can be shared in Windows by using the File Sharing Wizard, the Advanced Sharing dialog box, the Shared Folders snap-in, and the Share and Storage Management snap-in.
- Windows Server 2008 includes tools to manage and monitor server operation and resources, including Task Manager, Event Viewer, Performance Monitor, and Windows System Resource Manager.

- Regular backups provide a safety net to restore a system to working order in the event of a disk failure or file corruption. Fault tolerance ensures that a system can continue running after a failure has occurred. System failures can be power failures, disk failures, and entire computer failures. Backups can provide fault tolerance when data has been lost or corrupted, and RAID disk systems can provide fault tolerance in the event of a disk failure. UPSs can protect systems against the effects of power loss or fluctuations.

Key Terms

active partition A partition that can hold boot files the BIOS loads before it can start the OS.

baseline A record of performance data gathered when a system is performing well under normal operating conditions. The baseline can then be compared with data collected during peak resource demands to give you insight into your system's capabilities and limitations.

basic disk A disk configuration in which the space on the disk can be divided into one to four partitions.

boot partition The partition or logical drive holding Windows OS files.

data collector set A feature of Performance Monitor that specifies the performance counters you want to collect, how often to collect them, and the time period.

default groups Special groups with rights already assigned; created during installation in a Windows environment.

disk mirroring A fault-tolerant disk configuration in which data is written to two hard drives rather than one so that if one disk fails, the data isn't lost.

disk quotas A feature available in some file systems that allows an administrator to set a limit to how much disk space a user's files can occupy.

disk striping with parity A fault-tolerant disk configuration in which parts of several physical disks are linked in an array, and data and parity information are written to all disks in this array. If one disk fails, data can be reconstructed from the parity information written on the others.

dynamic disk A disk configuration in Windows that can be divided into one or more volumes. You can create up to 1000 volumes per dynamic disk (although no more than 32 is recommended). A dynamic disk offers features that a basic disk doesn't, namely RAID and disk spanning.

extended partition A partition type that can be divided into one or more logical drives, each of which can be formatted and assigned a drive letter.

local profile A user profile stored on the same system where a user logs on; created from a hidden profile called Default the first time a user logs on to the system. *See also* user profile.

Network File System (NFS) The native Linux file-sharing protocol.

NTFS permissions A feature in Windows NTFS that gives administrators fine-grained control over file and folder access for both network users and interactive users.

power conditioning A method of cleaning the power input, removing noise caused by other devices on the same circuit.



primary partition A partition type that can be formatted with a file system and assigned a drive letter or mounted in an empty folder on an existing drive letter; also called a volume. *See also* volume.

redundant power supply A second power supply unit in the computer case. Each unit is capable on its own of maintaining adequate power to the computer, so if one power supply fails, the other unit takes on the full load.

rights In Windows, they define the types of actions a user can perform, such as creating file shares or installing software.

roaming profile A user profile in a Windows environment that's stored on a server and can be accessed from any computer the user logs on to. *See also* user profile.

Server Message Block (SMB) The Windows file-sharing protocol.

special identity groups A type of group in Windows in which membership is controlled dynamically by Windows, can't be viewed or changed manually, and depends on how an account accesses the OS. For example, membership in the Authenticated Users group is assigned to a user account automatically when the user logs on to a computer or domain.

surge protection Power protection that evens out spikes or sags in the main current and prevents them from affecting a computer.

system partition The active primary partition storing the Windows boot loader.

uninterruptible power supply (UPS) A power protection device that includes a battery backup to take over if the main current fails; usually incorporates power conditioning and surge protection.

user profile A collection of a user's personal files and settings that define his or her working environment.

volume Part or all of the space on one or more disks that contains or is ready to contain a file system. In Windows, volumes with file systems are usually assigned a drive letter. In Linux, volumes are mounted in the file system and accessed as though they were just another folder.

Review Questions

1. Which of the following is a function of a user account? (Choose all that apply.)
 - a. Establishes a link between the user and the computer's IP address
 - b. Provides a method for user authentication
 - c. Provides information about a user
 - d. Authorizes a user to log on to network servers
2. Which of the following is true of group accounts? (Choose all that apply.)
 - a. They organize users for easier assignment of resource permissions.
 - b. They can be used only to assign permissions, not rights.
 - c. Each group has a password assigned.
 - d. You can select a group scope in Active Directory but not in Windows 7.

3. Which of the following is true of a user logon name in Windows?
 - a. It's case sensitive.
 - b. It's not case sensitive.
 - c. It must contain both uppercase and lowercase letters.
 - d. It must contain at least one number.
4. Which of the following is a group scope in Active Directory? (Choose all that apply.)
 - a. Domain local
 - b. Global
 - c. Distribution
 - d. Security
5. Which of the following is a default group in Active Directory? (Choose all that apply.)
 - a. Network Administrators
 - b. Backup Operators
 - c. Server Operators
 - d. User Managers
6. Which of the following is a special identity group? (Choose all that apply.)
 - a. Everyone
 - b. Logged-On Users
 - c. Authenticated Users
 - d. Creator Owner
7. Which of the following is a collection of a user's personal data folders and application settings that's available at any computer where the user logs on?
 - a. Local profile
 - b. Roaming profile
 - c. Mandatory profile
 - d. Network profile
8. Which command in Linux might be needed to perform user account management if you aren't logged on as the root user?
 - a. runas
 - b. superuser
 - c. passwd
 - d. sudo



9. Which command in Linux gives you extensive help on how to use a command?
 - a. help
 - b. more
 - c. man
 - d. guide
10. Which is true about Linux user accounts? (Choose all that apply.)
 - a. They must belong to at least one group.
 - b. The account names are case sensitive.
 - c. The full name is a required part of the user account.
 - d. The account names can't contain lowercase letters.
11. Which is true about partitions? (Choose all that apply.)
 - a. A partition is always a volume, too.
 - b. You can have up to four primary partitions.
 - c. An extended partition is assigned a drive letter by default.
 - d. Only a primary partition can be active.
12. Which file system supports a maximum file size of 2 GB?
 - a. NTFS
 - b. FAT16
 - c. FAT32
 - d. EXT3
13. Which NTFS feature should you configure if you want users to be able to revert to an older version of a file?
 - a. Disk quotas
 - b. EFS
 - c. Mount points
 - d. Shadow copies
14. Which is true about file compression? (Choose all that apply.)
 - a. A compressed file can't be encrypted with EFS.
 - b. It's a standard feature in Windows starting with FAT32.
 - c. Decompression occurs automatically when a file is accessed.
 - d. Compressed files can't have permissions assigned to them.

15. Which of the following is true about permissions in NTFS?
 - a. An Allow permission always overrides a Deny permission.
 - b. Permissions can be set only on folders, not on files.
 - c. By default, permissions are inherited automatically from the parent folder.
 - d. The last permission assigned is the only one that takes effect.
16. Which of the following is a permission in the Linux OS? (Choose all that apply.)
 - a. Read
 - b. Modify
 - c. Delete
 - d. Execute
17. Which of the following correctly describes how sharing and NTFS permissions work?
 - a. When a file is accessed over the network, sharing permissions are checked first.
 - b. When a file is accessed interactively, only sharing permissions are checked.
 - c. When both sharing and NTFS permissions are applied, the least restrictive permissions apply.
 - d. Sharing permissions can be assigned to separate files.
18. Which of the following is true of Windows printing terminology? (Choose all that apply.)
 - a. A printer is a physical device.
 - b. A print server is a computer sharing a printer.
 - c. A print queue stores jobs waiting to be printed.
 - d. A printer pool is two or more printers representing a single print device.
19. Which of the following tools is used to view information categorized as Information, Warning, or Error?
 - a. Event Viewer
 - b. Performance Monitor
 - c. Task Manager
 - d. Report Generator
20. Which is best described as a record of performance data gathered when a system is performing well under normal conditions?
 - a. Data collector set
 - b. Real-time monitoring
 - c. Baseline
 - d. Performance counters



21. Why would you choose to monitor a system's performance remotely?
 - a. You don't have permission to log on to the system.
 - b. The computer you want to monitor doesn't have enough memory to run Performance Monitor.
 - c. The system is running Windows 7, which doesn't have Performance Monitor installed.
 - d. You want to lessen the impact of the monitoring session on the computer.
22. Which is true about Windows Backup in Windows Server 2008 R2? (Choose all that apply.)
 - a. Backups can be stored on network drives.
 - b. The Volume Shadow Copy Service backs up open files.
 - c. You can back up another computer remotely.
 - d. Backups can be stored on tape.
23. Which UPS type is best for computer power backup?
 - a. Standby
 - b. Online
 - c. Offline
 - d. Always on
24. Which RAID level uses a minimum of three disks and provides fault tolerance?
 - a. RAID 1
 - b. RAID 0
 - c. RAID 5
 - d. RAID 1+0

Challenge Labs



Challenge Lab 9-1: Creating Users in Linux with the `newusers` Command

Time Required: 30 minutes

Objective: Create new users in Linux with the `newusers` command.

Required Tools/Equipment: A computer with a Linux OS installed or a Linux Live CD

Description: This lab can be done in groups. In this lab, you create new Linux users in batch mode with the `newusers` command, which accepts a text file as input. Use the man pages for the `newusers` command and create a correctly formatted file to use as input to the `newusers` command. Five new users should be created, and each user should be new in the system. The users' UIDs

should specified in the file and be in the range 5001 to 5005. The primary group name should be the same as the user's logon name. The user's full name can be whatever you like. The home directory should be at `/home/username`, and the user's default shell should be `/bin/bash`. After you're finished, print a copy of the input file you created and hand it in to your instructor.



Challenge Lab 9-2: Creating Users, Groups, and Shares to Represent a Company's Organization

Time Required: 1 hour

Objective: Create users, groups, and shares and assign permissions.

Required Tools/Equipment: Your classroom computer with Windows or Linux installed

Description: This lab can be done in groups and with Windows or Linux. You're the IT administrator for a new company with three departments: Marketing, Operations, and Support. Each department has three users and needs access to shared documents that the other departments don't have access to. Users in each department should be able to create, modify, and delete files in their department's shared folder. Create one shared folder named Public that all users have read-only access to, except one user in the Operations department, who should be able to create, modify, and delete documents. Administrators have full access to all shared folders. Perform the following tasks:

- Create users for each department, assigning an initial password to each user but requiring the user to change the password at first logon.
- Create groups for each department and assign user memberships appropriately.
- Create shared folders for each department and the Public share.
- Assign sharing and NTFS permissions for the departmental shared folders so that users in each department have the necessary access to their shared folders.
- Assign sharing and NTFS permissions to the Public shared folder.

After you're finished, diagram your solution, showing groups, group memberships, and shares and their associated sharing and NTFS permissions.

Case Projects



Case Project 9-1

You have created shared folders for all your companies departments and assigned the appropriate permissions. Everyone can access the shares as planned, but now you find that an extraordinary amount of disk space is being used on the server. What can you do to limit how much users can store on your servers?



Case Project 9-2

You have just purchased a new Dell PowerEdge R910 Rack Server. The server is equipped with two Xeon E7520 CPUs, 8 GB RAM, four 200 GB SAS hard drives, and four 750-watt power supplies. You have a 19-inch LCD monitor connected to your server. You want a UPS to keep your server running in the event of a power failure for up to 20 minutes. You prefer a rack-mounted solution. Use the UPS selector tool at APC's Web site (a well-known UPS vendor). The UPS selector is at www.apcc.com/tools/ups_selector/. Determine which model UPS will work well for this server, and state your reasons.



chapter

10

Introduction to Network Security

After reading this chapter and completing the exercises, you will be able to:

- Develop a network security policy
- Secure physical access to network equipment
- Secure network data
- Use tools to find network security weaknesses

You have learned how to make wise choices about which network technologies to use and how to configure networking protocols for your network, and you have a good handle on working with network operating systems. Your learning in becoming a network technician or network administrator is almost done, right? Wrong. You have plenty to do before you can bring a network online because of all the security risks that can make networks fail: Trojan programs, worms, spammers, denial-of-service attacks, spyware, network attackers, backdoors, and on and on. Understanding and preventing attacks that can infiltrate or disrupt your network are what network security is all about. This chapter gives you a solid foundation of knowledge and tools for protecting your network and its users.

Network Security Overview and Policies

Network security can mean different things to different people. To network users, network security is sometimes considered a necessary evil that takes the form of hard-to-remember passwords that must be changed frequently and cryptic terms, such as “VPN” and “IPSec,” to describe methods they use to access the network. To other users, network security means the comfort of knowing that if they erase their hard drives accidentally, the friendly network administrator will gladly restore their data from the most recent system backup.

Perceptions about network security also vary depending on the industry a person is in or the job a person does. A chemical engineer might perceive network security to mean that the compound he has just developed is safe from the competition’s eyes. A lawyer might describe network security as a means of safeguarding against illegal activities, such as unlawful distribution of copyrighted materials. To a network engineer, network security might simply mean she can get a good night’s sleep, secure in the knowledge that the network is safe from the latest threats—at least until tomorrow. To a chief security officer (CSO), however, network security means more than job-specific tasks. A CSO knows that the goal of network security is to protect the organization and its users, customers, and business partners from any threat to the integrity of information passing through or residing on the corporate network.

Ideally, network security should be as unobtrusive as possible, allowing network users to concentrate on the tasks they want to accomplish instead of how to get to the data they need to perform those tasks. Achieving this goal enables an organization to go about its business confidently and efficiently. In today’s security-conscious world, a company that can demonstrate its information systems are secure is more likely to attract customers, partners, and investors. That’s where a network security policy comes in.

Developing a Network Security Policy

The lofty goal of good, unobtrusive security encompasses many dimensions, so where do you start? With a security policy that reflects your organization’s attitude toward securing network resources. A network security policy is a document that describes the rules governing access to a company’s information resources, enforcement of these rules, and steps taken if rules are breached. The document should describe not only who can have access to which resources, but also how these resources are allowed to be used after they’re accessed. In addition, it should follow these basic guidelines:

- A security policy should be easy for ordinary users to understand and reasonably easy to comply with. If you make the policy too difficult to understand or follow, users

resist adhering to it. A policy requiring users to change their passwords every week, for example, is too difficult to follow. Users who must change their passwords too frequently often select easy-to-remember passwords that are based on common words and, therefore, are easy to crack.

- A security policy should be enforceable. A rule that can't be reasonably enforced will almost always be broken. For example, you shouldn't prohibit use of the Internet during certain hours of the day unless you have a method of monitoring or restricting this use.
- A security policy should clearly state the objective of each policy so that everyone understands its purpose. For example, a policy that states "Misuse of the network is forbidden" doesn't define misuse, making this policy useless because of its lack of specificity.

The preceding guidelines explain how a security policy should be written. Now you need to know what information should be included in a security policy.

Determining Elements of a Network Security Policy

Explaining all the elements of a security policy is beyond the scope of this book, but the following items give you a solid start:

- *Privacy policy*—Describes what staff, customers, and business partners can expect for monitoring and reporting network use.
- *Acceptable use policy*—Explains for what purposes network resources can be used.
- *Authentication policy*—Describes how users identify themselves to gain access to network resources. Logon names, password conventions, and authentication methods should be described.
- *Internet use policy*—Explains what constitutes proper or improper use of Internet resources.
- *Access policy*—Specifies how and when users are allowed to access network resources. Policies should exist for both onsite and remote access to the network.
- *Auditing policy*—Explains the manner in which security compliance or violations can be verified and the consequences for violations.
- *Data protection*—Outlines the policies for backup procedures, virus protection, and disaster recovery.



TIP

To learn more about security policies, refer to RFC 2196 at www.faqs.org/rfcs/rfc2196.html.

Your security policy might have other elements, depending on the type of organization it's being created for and the level of security required, but the preceding list is usually the minimum for most networks. Keep in mind that a well-thought-out security policy also protects the organization legally. If no policy exists, disciplining or prosecuting people who misuse or intrude on the network is more difficult. Unfortunately, after you create a security policy, your work isn't done. A security policy should be a constant work in progress, with modifications made as needed to reflect changing technology and business practices.

Understanding Levels of Security

Before starting to design a network security policy, you need to be aware of the relationship between the level of security imposed on a network and the cost and difficulty of supporting the network. Security doesn't come without a cost. If you're the network administrator for the security department of certain government offices, price is likely no object in determining the extent of security measures. However, if you're setting up a network for a small manufacturer of household items, you might need to scale back on security measures. Before determining the level of security your network needs, answer these questions:

- What must be protected? Is there information on the network that would compromise the viability of the company or its customers if it fell into the wrong hands?
- From whom should data be protected? Is the biggest threat from people inside or outside the company?
- What costs are associated with security being breached and data being lost or stolen?
- How likely is it that a threat will actually occur? Do you have a high-profile business, or do you have known competitors who are likely to want to sabotage your business or steal trade secrets?
- Are the costs to implement security and train personnel to use a secure network outweighed by the need to create an efficient, user-friendly environment?

Depending on your answers, you'll likely decide to implement one of these security levels or some combination of them: highly restrictive, moderately restrictive, and open.

Highly Restrictive Security Policies Highly restrictive security policies usually include features such as data encryption, complex password requirements, detailed auditing and monitoring of computer and network access, intricate authentication methods, and policies governing use of the Internet and e-mail. Some features of this type of policy might require third-party hardware and software. The high expense of implementing these restrictive policies comes in the form of design and configuration costs for software and hardware, staffing to support security policies, and lost productivity caused by a steep learning curve for users. However, if you need highly restrictive security, it's probably because the cost of a security breach would be more expensive than implementing the security policy.

Moderately Restrictive Security Policies Most organizations can probably opt for a moderately restrictive security policy. These policies require passwords for each user but not overly complex passwords. Auditing is geared toward detecting unauthorized logon attempts, misuse of network resources, and network attacker activity. Most network OSs contain satisfactory authentication, monitoring, and auditing features to carry out these policies. The network infrastructure can be secured with moderately priced off-the-shelf hardware and software, such as firewalls and access control lists. The costs of moderate security policies are mainly in initial configuration and support. This type of policy is used in a typical business setting, in which users have personal files that require moderate security and users in some departments are responsible for files that might need additional security measures, such as data encryption.

Open Security Policies A company that uses an open network security policy might have simple or no passwords, unrestricted access to resources, and probably no monitoring and auditing. This type of policy might make sense for a small company with the main goal of making access to network resources easy. The company might not want to spend additional funds for the employee training often required for more restrictive policies.

In an open security environment, Internet access probably shouldn't be possible via the company LAN because it invites too many possibilities for outside mischief or inside abuse. If companywide Internet access is available, a more restrictive policy is probably warranted. In an open security environment, sensitive data, if it exists, might be kept on workstations that are backed up regularly and physically inaccessible to other employees.

Common Elements of Security Policies No matter which type of security policy your company uses, some common elements should be present. Virus and other malware protection for servers and desktop computers is a must for every computing environment, and there should be policies aimed at preventing malware from being downloaded or spread. Backup procedures for all data that can't be reproduced easily should be in place, and a disaster recovery procedure must be devised and carried out. Remember: Security is aimed not only at preventing improper use of or access to network resources, but also at safeguarding the company's information, which is often more valuable than its physical assets. Before you turn to methods and practices for securing data, however, one often neglected aspect of security must be discussed: physical security of servers and network devices.

Securing Physical Access to the Network

A common guideline in discussing network security is "If there's physical access to the equipment, there's no security." This guideline applies to servers, desktop computers, network devices (such as routers and switches), and even network media. No matter how strong your logon name and password schemes are, if a person has physical access to a device, access to data isn't far behind.

There are numerous ways to break into an unprotected computer or networking device. A computer left alone with a user logged on is particularly vulnerable. A person walking by could access all the files the logged-on user has access to. If the computer is a server and an administrator account is logged on, a person has full reign of the network and can even give his or her account administrator control. Even if no user is logged on, people could log on to the computer with their own accounts and access files they wouldn't normally have access to. Failing that, the computer could be restarted and booted from removable media, thereby bypassing the normal OS security. Last, if a person is desperate, the entire computer or its hard drives could be stolen and later cracked. The following sections describe best practices for preventing a physical assault on your network.

Physical Security Best Practices

The following list is an overview of best practices to secure your network from physical assault:

- When planning your network, ensure that rooms are available to house servers and equipment. These rooms should have locks to prevent unauthorized access and be suitable for the equipment being housed, including having enough power receptacles,



adequate cooling measures, and an environment clear of electromagnetic interference (EMI) sources. In addition, rooms should be inaccessible through false ceilings.

- If a suitable room isn't available, locking cabinets, freestanding or wall mounted, can be purchased to house servers and equipment in public areas. Wall-mounted cabinets are particularly useful for hubs, switches, and patch panels. You must be certain cabinets have suitable ventilation for the devices they're housing.
- Wiring from workstations to wiring cabinets should be inaccessible to eavesdropping equipment. Wiring that's not concealed in floors or ceilings should be concealed in raceways or other channeling devices to discourage access.
- Your physical security plan should include procedures for recovery from natural disasters, such as fires or floods.

Physical Security of Servers Securing servers from physical access should be a high priority in any security plan. This goal can be achieved in a number of ways, and sometimes a combination of methods works best, depending on your environment. Many servers are stashed away in a lockable wiring closet along with the switch the servers are connected to. This setup is fine as long as the environment is suitable for the server, and the same people who have authority to access wiring and hubs also have authority to access the servers, although this isn't always the case.

Servers often require more tightly controlled environmental conditions than patch panels and switches do. They can generate a substantial amount of heat and, therefore, need adequate cooling. The lack of cooling can damage hard drives, cause CPUs to shut down or malfunction, and damage power supplies, among other consequences.

In addition to adequate cooling, server rooms should be equipped with power that's preferably on a circuit separate from other electrical devices. Enough power outlets should be installed to eliminate the need for extension cords. Because you'll be putting servers on uninterruptible power supplies (UPSs), you need to verify power requirements for UPSs. Some UPSs require special twist-lock outlet plugs rated for high currents. Nothing is more frustrating than getting a brand-new UPS and preparing to plug in your servers, only to find that the wall outlets are incompatible with the UPS requirements.

Sometimes putting your servers in a place that's accessible to people who should *not* have physical access to servers is unavoidable. For example, you might have different teams maintaining internetworking equipment and servers, and you don't want internetworking maintenance teams to have access to servers. If you don't have the facilities to separate the two types of equipment physically, however, you can still take some steps to provide a measure of physical security. Many servers come with locking cabinets to prevent access to the inside of the case. Some also have lockable covers that protect the drives and power buttons from unauthorized access. You can also place the keyboard, mouse, and monitor in an area separate from the actual server by using long-distance cable extenders or network-based keyboard, mouse, video (KVM) switches. Last, you can place the server in a freestanding locking cabinet.

If you're forced to place servers in a public access area, locking cabinets are a must. Even if no users have malicious intentions, someone is sure to kick the server, spill coffee on it, unplug it, or inflict some sort of accidental damage. You can purchase rack-mountable servers, which are designed to bolt to a standard 19-inch equipment rack. To conserve space, you can purchase a freestanding cabinet with a built-in 19-inch rack, allowing you to store several servers. Make

sure the cabinet you purchase is well ventilated or allows adding fans for ventilation. These cabinets typically start at about \$1000. Like everything else, security comes with a price.

Security of Internetworking Devices Routers and switches contain critical configuration information and perform tasks essential to your network's operation. A user with physical access to these devices needs only a laptop or handheld computer and a few easily discovered keystrokes to get into the router or switch, change passwords, and view or change the device's configuration. In addition, a person who has access to a switch port can attach a laptop with a protocol analyzer installed. From that point, it's simply a matter of waiting for the right data to be captured to gain access to critical or sensitive information.

Clearly, internetworking devices should be given as much attention to physical security as servers. These devices give potential network infiltrators access to the network and an opportunity to wreak havoc. Configuration changes made to routers and switches can have disastrous consequences. In addition, access to routers can reveal network topology information you might not want everyone to know. The more troublemakers know about a network's configuration, the more tools they have to break into the network or cause problems on it.

A room with a lock is the best place for internetworking devices, but a wall-mounted enclosure with a lock is the next best thing. These cabinets are usually heavy-duty units with doors that swing out and built-in 19-inch racks. Wall-mounted cabinets are expensive, so budget between \$300 and \$1000 for them, depending on the features and size you need. Some cabinets have a built-in fan or a mounting hole for a fan. The racks also come with convenient channels to run wiring.

Securing Access to Data

Physically securing your network assets is only one part of the security puzzle. Networks are designed to give users operating from remote locations access to data, whether the remote location is the next room or the other side of the world. Securing data on a network has many facets, some of which are discussed in more detail in the following sections:

- *Authentication and authorization*—Identifying who's permitted to access which network resources
- *Encryption*—Making data unusable to anyone except authorized users
- *Virtual private networks (VPNs)*—Allowing authorized remote access to a private network via the public Internet
- *Firewalls*—Installing software or a hardware device that protects a computer or network from unauthorized access and attacks designed to cripple network or computer performance
- *Virus and worm protection*—Securing data from software designed to destroy data or make computers and networks operate inefficiently
- *Spyware protection*—Securing computers from inadvertently downloading and running programs that gather personal information and report on users' Web browsing and computing habits
- *Wireless security*—Implementing measures for protecting data and authorizing access to a wireless network

The following sections discuss these areas of security and explore some features of network OSs that help secure a network.

Implementing Secure Authentication and Authorization

Authentication and authorization are security features that enable administrators to control who has access to the network (authentication) and what users can do after they're logged on to the network (authorization). Network OSs include tools that enable administrators to specify options and restrictions on how and when users can log on to the network. There are options for password complexity requirements, logon hours, logon locations, and remote logons, among others. After a user is logged on, file system access controls and user permission settings determine what a user can access on a network and what actions a user can perform (such as installing software or shutting down a system) on the network.

Configuring Password Requirements in a Windows Environment Administrators can specify whether a password is required for all users, how many characters a password must be, and whether the password should meet certain complexity requirements. Windows 7 allows passwords up to 128 characters, but a minimum of five to eight characters is typical. A password minimum length of zero means blank passwords are allowed, a setting that might be adequate for networks with open security policies but should never be used for networks requiring more security. A password policy with complexity requirements means that user passwords must have three of these four characteristics: lowercase letters, uppercase letters, numbers, and special (nonalphanumeric) characters.

Other password options include:

- *Maximum password age*—Specifies, in days, how often users must change their passwords.
- *Minimum password age*—Specifies the minimum number of days that must pass before users can change their passwords
- *Enforce password history*—Determines how many different passwords must be used before a password can be used again.

One word of caution on password settings: Don't make your password requirements so stringent that well-meaning users feel forced to write their passwords down so that they can remember them. Password policies should make it difficult for would-be attackers to gain access to the system, but not so difficult that your users have trouble adhering to the policies.

When a user fails to enter a correct password, a policy can be set to lock the user account, preventing that account from logging on. This account lockout option, used to prevent intruders from guessing a password, can be enabled or disabled. If it's enabled, the administrator can specify how many times an incorrect password can be entered before the account is locked. After it's locked, the administrator can require manual unlocking or automatic unlocking of the account after a certain amount of time has expired.

Password policies for a single Windows 7 or Windows Server 2008 computer can be set in the Local Security Policy console found in Administrative Tools. Figure 10-1 shows the Local Security Policy console with Password Policy selected. In a domain environment, password policies are set by using group policies on a domain controller. Password policies take effect immediately for all existing and new user accounts.

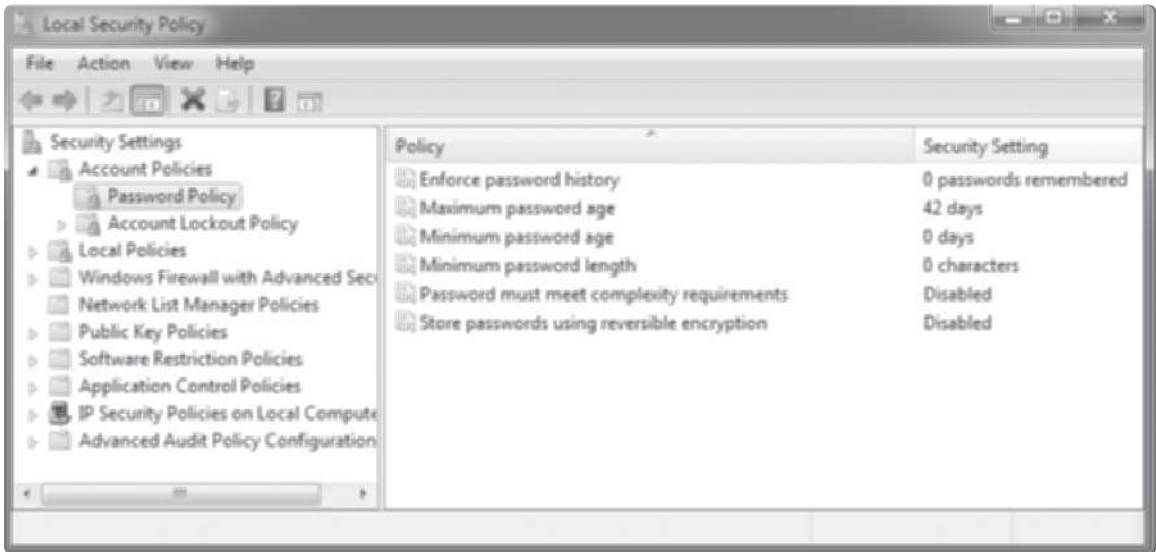


Figure 10-1 Password policy settings in Windows 7

Courtesy of Course Technology/Cengage Learning

Configuring Password Requirements in a Linux Environment Like Windows, Linux has a number of password options that can be configured. Whereas Windows password policies are set globally and affect all users, Linux password configuration can be done globally or on a user-by-user basis. Password options in a standard Linux environment include the maximum password age, minimum password age, and number of days' warning a user has before the password expires.

For these password options to be available, the Linux system must be using **shadow passwords**, a secure method of storing user passwords on a Linux system. The passwords are stored in an encrypted format in the shadow file located in the /etc directory; this file can be accessed only by the root user. Most Linux installations use shadow passwords by default. Password options are set by editing the /etc/login.defs configuration file. Only accounts created after the login.defs file has been edited are affected.

Other password options can be configured, including account lockout and password history, by using Pluggable Authentication Modules (PAM). PAM is the standard software service on many Linux distributions for authenticating users. One standard feature of PAM authentication is a password complexity test. When a new password is created for a user account, a database of common dictionary words is searched. If the password is found, the user is informed that the password isn't complex enough to be considered secure. The default configuration allows using weak passwords, but this feature can be changed to disallow a weak password. Password options controlled by PAM can be configured by editing the /etc/pam.d/common-password file.

Reviewing Password Dos and Don'ts Some general rules for creating passwords include the following:

- Do use a combination of uppercase letters, lowercase letters, and numbers.
- Do include one or more special characters, such as periods, dollar signs, exclamation points, and question marks.

- Do consider using a phrase, such as NetW@rk1ng !s C00l. Phrases are easy to remember but generally difficult to crack, especially if you mix in special characters and numbers.
- Don't use passwords based on your logon name, your family members' names, or even your pet's name. Users often use these types of passwords, but unfortunately, they're easy to guess after attackers discover personal information about users.
- Don't use common dictionary words unless they are part of a phrase, and substitute special characters and numbers for letters.
- Don't make your password so complex that you forget it or need to write it down somewhere.



Hands-On Project 10-1: Setting Password and Lockout Policies in Windows

Time Required: 15 minutes

Objective: Set password and lockout policies in the Windows Local Security Policy console.

Required Tools/Equipment: Your classroom computer

Description: This project shows you how to use the Local Security Policy console in Windows 7. You set a password policy that specifies the following:

- Users must use 10 different passwords before reusing a password.
- Users must change their password every 30 days.
- Users can't change their password more often than every seven days.
- The minimum password length is six characters.
- The password must contain three of these characteristics: uppercase letters, lowercase letters, numbers, or special nonalphanumeric characters.

You set the account lockout policy to enforce the following:

- User accounts are locked out after four invalid logon attempts.
 - Locked accounts are unlocked automatically after 60 minutes.
 - The counter is reset 15 minutes after each invalid logon attempt.
1. Log on to your computer as an administrator.
 2. Click **Start**, point to **Administrative Tools**, and click **Local Security Policy**. The Local Security Policy console opens.
 3. Click to expand **Account Policies** and click **Password Policy**.
 4. In the right pane, double-click **Enforce password history**. Set the value to 10 passwords remembered, and then click **OK**.
 5. Double-click **Maximum password age**, set the value to 30, and click **OK**.
 6. Double-click **Minimum password age**, set the value to 7, and click **OK**.
 7. Double-click **Minimum password length**, set the value to 6, and click **OK**.

8. Double-click **Password must meet complexity requirements**, click the **Enabled** option button, and click **OK**. The settings should look similar to Figure 10-2 when you're finished.

Policy	Security Setting
Enforce password history	10 passwords remember...
Maximum password age	30 days
Minimum password age	7 days
Minimum password length	6 characters
Password must meet complexity requirements	Enabled
Store passwords using reversible encryption	Disabled

Figure 10-2 Viewing password policy settings

Courtesy of Course Technology/Cengage Learning

9. Next, you set the account lockout policy. Under Account Policies, click **Account Lockout Policy**.
10. In the right pane, double-click **Account lockout threshold**, set the value to **4**, and then click **OK** twice. Windows fills in the other policies automatically with default values.
11. Double-click **Account lockout duration**, set the value to **60 minutes**, and click **OK**.
12. Double-click **Reset account lockout counter after**, set the value to **15 minutes**, and click **OK**. The Local Security Policy console should look similar to Figure 10-3 when you're finished.
13. You test these settings in a challenge lab at the end of the chapter. Close the Local Security Policy console, and log off Windows for the next project.

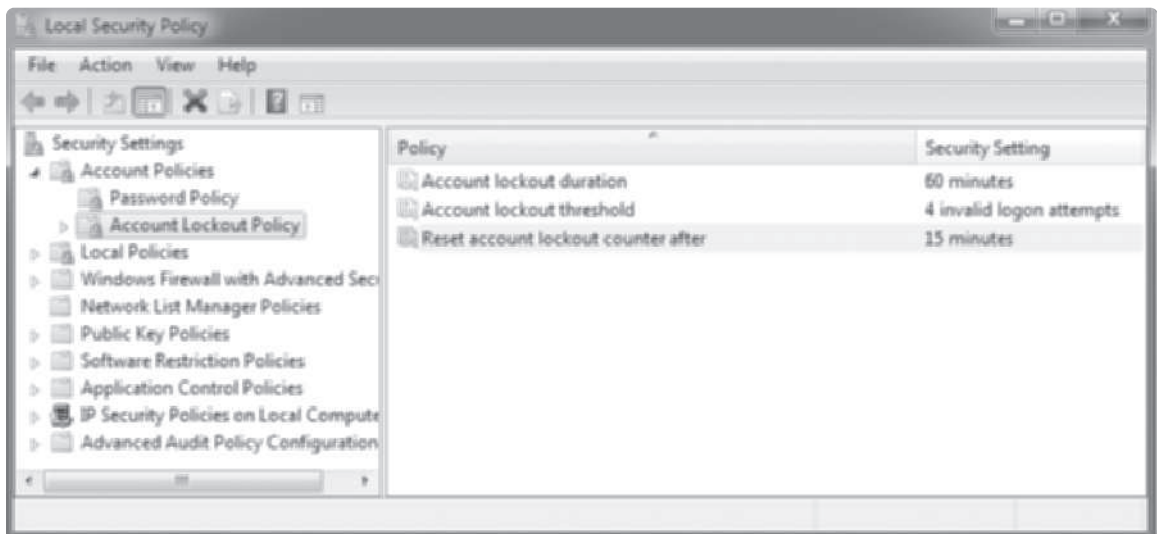


Figure 10-3 Account lockout policies

Courtesy of Course Technology/Cengage Learning



Hands-On Project 10-2: Setting Password Policies in Linux

Time Required: 15 minutes

Objective: Set password policies in Linux by editing the `login.defs` file.

Required Tools/Equipment: Your classroom computer with Linux installed or a Linux Live CD; this project uses Ubuntu Linux 10.04.

Description: This project shows you how to edit the `login.defs` file to set the Linux password policy as follows:

- Users must change their password every 30 days.
 - Users can't change their password more often than every 7 days.
 - Users are warned 4 days before their password expires.
1. Log on to Linux. Open a terminal window by clicking **Applications**, pointing to **Accessories**, and clicking **Terminal**.
 2. Type `cd /etc` and press **Enter**. To gain root privileges for editing the `login.defs` file with `gedit`, type `sudo gedit login.defs` and press **Enter**. Enter your password if prompted.
 3. Scroll down in the `login.defs` file until you find a line beginning with `PASS_MAX_DAYS` (probably between line 150 and 200). Change the value next to `PASS_MAX_DAYS` to 30. Change the value next to `PASS_MIN_DAYS` to 7 and the `PASS_WARN_AGE` value to 4.
 4. Click **Save**, and then close `gedit`.
 5. You set additional password policies and an account lockout policy in a challenge lab at the end of the chapter. Shut down Linux.

Restricting Logon Hours and Logon Location Some network administrators allow users to log on any time of the day and any day of the week, but if your security policy states otherwise, both Windows and Linux have solutions to restrict logon by time of day, day of the week, and location.

In a Windows domain environment, allowed logon times can be set for each user account, as shown in Figure 10-4. The default settings allow logon 24 hours per day, seven days a week. A common use of restricting logon hours is to disallow logon during system backup, which usually takes place in the middle of the night. In Figure 10-4, the dark boxes indicate times that the user can log on, and the white boxes indicate hours the user can't log on. In this example, logging on from 2 a.m. to 4 a.m. is not allowed. Note that the logon hours option is available only in a Windows domain environment.

Sometimes users log on to the network from computers that aren't their regular workstations. This practice might be allowed in your environment, but extending this option to users who have access to sensitive data can be dangerous. If a user logs on at a workstation in a coworker's office and then walks away from that machine, the coworker now has access to the sensitive data. To prevent this problem, users can be restricted to logging on only from particular workstations. Figure 10-5 shows the Windows user account settings for logon location; the user can log on only to the computers named `smiller01` and `engineering`. As with logon hours, this option is available only in a Windows domain environment.

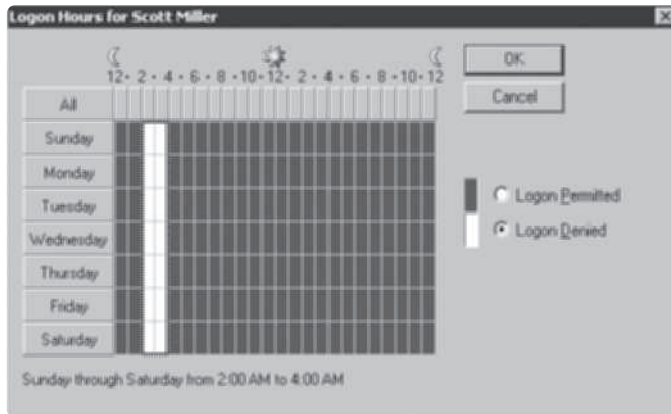


Figure 10-4 Setting logon hours for a user

Courtesy of Course Technology/Cengage Learning

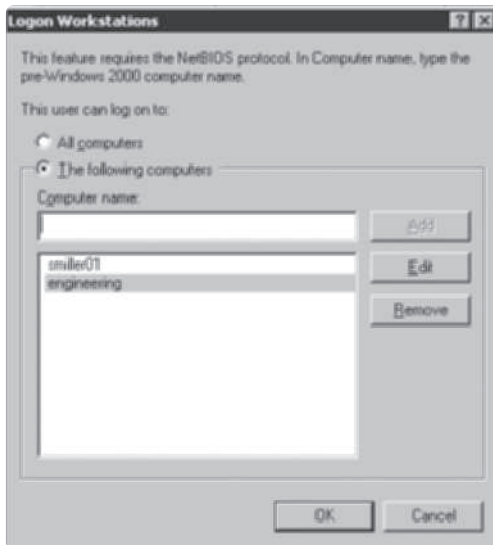


Figure 10-5 The Logon Workstations dialog box

Courtesy of Course Technology/Cengage Learning

The Linux OS offers similar features for logon time restrictions with timeout, which is usually installed as an add-on package in most Linux distributions. In general, standard Linux distributions don't include a GUI to configure these settings. Typically, text configuration files must be edited to enable and configure logon restrictions, and these details are beyond the scope of this book.

Giving employees remote access to networks has become essential in many businesses, especially with secure, high-speed network connections. However, remote access isn't an all-or-nothing proposition because you can restrict remote access by user or by creating remote access policies that permit or deny user access based on criteria such as a user's group memberships and time of day.

Authorizing Access to Files and Folders After users have logged on to a network or computer, they must be authorized to access network resources. A common network resource is shared files, which are controlled by the OS file system security. As discussed in Chapter 9, file system security allows administrators to assign file and folder permissions to users or groups of users.

As you know, Windows OSs have two options for file security: sharing permissions and NTFS permissions. Sharing permissions are applied to folders (and only folders) shared over the network. The files in a shared folder inherit the same permissions applied to the parent folder. Sharing permissions don't apply to files or folders if a user is logged on locally; they restrict only users accessing files across the network. Sharing permissions are the only file security option available in a FAT or FAT32 file system. In other words, if a user is logged on locally to a Windows system formatted with FAT or FAT32, that user has full access to all files on all drives.

NTFS permissions are more sophisticated than sharing permissions. Administrators can assign permissions to files as well as folders, so one level of permission can be assigned to a folder but a different level assigned to files in the folder, if needed. Additionally, NTFS permissions apply not only to access over the network, but also to file access by a locally logged-on user. Recall that Linux also supports file and folder security, albeit quite differently than with NTFS.

File and folder permissions are a necessary tool administrators use to make network resources secure and still give users appropriate access to the resources they're permitted to use. However, permissions don't protect data traversing the network media, nor do they protect data in files if file system security has been compromised—which is where data encryption comes in.

Securing Data with Encryption

Many network administrators use **encryption** technologies to safeguard data as it travels across the Internet and even within the company network. This security measure prevents people from using eavesdropping technology, such as a packet sniffer, to capture packets and use data in them for malicious purposes. Data stored on disks can also be secured with encryption to prevent someone who has gained physical access to the computer from being able to use the data.

Using IPSec to Secure Network Data The most widely used method for encrypting data as it travels network media is using **IP Security (IPSec)**, an extension to the IP protocol. IPSec works by establishing an association between two communicating devices. An association is formed by two devices authenticating their identities via a preshared key, Kerberos authentication, or digital certificates.

A **preshared key** is a series of letters, numbers, and special characters, much like a password, that both communicating devices use to authenticate each other's identity. A network administrator must enter the same preshared key in the IPSec configuration settings on both devices. **Kerberos authentication** is used in a Windows domain environment or on a Linux system to authenticate users and computers. It also uses keys, but the OS generates the keys, which makes this method more secure than having an administrator enter keys. **Digital certificates** involve a certification authority (CA). Someone wanting to send encrypted data

must apply for a digital certificate from a CA, which is responsible for verifying the applicant's authenticity. When an IPSec communication session begins, the communicating parties exchange certificates, and each party sends the certificate to the CA electronically to verify its authenticity. Windows servers can be configured as a CA when certificates are used on computers in a private network. Public CAs, such as Verisign, sell certificates to companies wanting to have secure communication sessions across public networks.

After the communicating parties are authenticated, encrypted communication can commence. Data sent across the network, even if it's captured by an eavesdropper, is unreadable to all but the intended recipient. Only the message recipient has the information needed to decrypt the message.

Although IPSec is an excellent way to secure data as it travels across a network, it doesn't secure data on disk drives if someone gains unauthorized access to the computer. Other security methods, discussed in the next section, are available for addressing this possibility.

On Linux systems, a simple method for encrypting files is using `gpg` (Gnu Privacy Guard), a command-line program. This program uses a password the user enters to encrypt the file specified as an argument to the `gpg` command, which then creates a new file with the encrypted contents of the specified file. The original file can then be deleted, and the only way to decrypt the data is to use the `gpg` command and supply the correct password.

Securing Data on Disk Drives Sometimes file system permissions aren't enough to stop an attacker who's determined to gain access to data on your system. If someone can access the hard disk where sensitive data is stored or compromise system security, your data could be vulnerable. Data stored on a disk drive can be encrypted, however, so that only the person who created the encrypted file can read the data, even if the hard disk is read sector by sector, therefore bypassing file system security.

In Windows OSs, Encrypting File System (EFS) is a standard feature on NTFS-formatted disks, as you learned in Chapter 9. To encrypt a file or the files in a folder, you simply select the "Encrypt contents to secure data" option in the Advanced Attributes dialog box. Encrypted files provide an extra layer of security that file permissions alone can't provide. Someone with physical access to a computer can boot the system into an OS on a CD/DVD, effectively bypassing normal file access controls. However, encrypted files are still inaccessible because an EFS certificate file matching the account of the user who encrypted the file must be available to the user trying to open the file. Even if someone is able to extract the file contents, he or she would see only encrypted gibberish.

EFS encrypts only files or folders with the Encrypt attribute set, and it can't be used on Windows system files. Windows also offers BitLocker for full disk encryption. This feature is available in the Ultimate and Enterprise editions of Windows 7 and Vista as well as Windows Server 2008. It can protect the entire system volume as well as other volumes and works in one of three modes:

- *Transparent mode*—Requires hardware with Trusted Platform Module (TPM) support. TPM hardware determines whether any changes have been made to the initial boot environment; if so, the user is prompted for a recovery key on a USB device or a recovery password. If no changes have been detected, the system boots normally. This method protects the system if someone tries to boot with a different OS.

- *USB key mode*—This mode is the most common method for booting a system configured with BitLocker that doesn't have TPM support. An encryption key is stored on a USB drive that the user inserts before starting the system.
- *User authentication mode*—The system requires a user password before it decrypts the OS files and boots. It's considered the fail-safe mode if TPM detects that the boot environment was compromised or if the USB key isn't detected.

BitLocker is a good security enhancement for laptops and servers that aren't physically secure. Third-party solutions, such as TrueCrypt, offer many of the features of BitLocker. TrueCrypt (www.truecrypt.org) is a free open-source product that works on Windows XP, Vista, 7, Mac OS, and Linux. Other open-source whole drive encryption solutions are available for the Linux environment, including Scramdisk 4Linux (<http://sd4l.sourceforge.net/>) and FreeOTFE (www.freeotfe.org).



A new feature in Windows 7, BitLocker To Go can protect the contents of removable storage, such as USB drives.

TIP

Securing Communication with Virtual Private Networks

A **virtual private network (VPN)** is a network connection that uses the Internet to give users or branch offices secure access to a company's network resources. VPNs use encryption technology to ensure that communication is private and secure, so while data travels through the public Internet, the connection remains private—hence the name “virtual private network.” Privacy is achieved by creating a “tunnel” between the VPN client and VPN server. A tunnel is created by encapsulation, in which the inner packet containing the data is encrypted and the outer headers contain the unencapsulated addresses that Internet devices need to route the packets correctly.

To use another mail delivery analogy, suppose you have an ultra-secure package to deliver, but you must use a courier. In a separate transaction, you deliver a key to the office manager at the package recipient's location. Next, you place the secret package containing the recipient's name in a lockbox. You put the lockbox inside an envelope and address the envelope to the office manager of the company where the recipient works. The courier can read the addressing on the envelope, but if the envelope is opened, the package contents are inaccessible without the key to the lockbox. The envelope is delivered, and the office manager removes the lockbox from the envelope and opens it with the key delivered earlier. The office manager can then deliver the package to the final recipient. In this analogy, the lockbox and outer envelope make up the VPN tunnel, and the office manager is the VPN server to which messages are delivered.

Figure 10-6 shows a VPN between a client computer and a corporate network. The tunnel connection is made between the client computer and the VPN server. After the VPN server opens the packet, the inner packet is decrypted (unlocked) and delivered to the resource the client requested.




Figure 10-6 A typical VPN connection

Courtesy of Course Technology/Cengage Learning

VPN servers can be configured on server OSs, such as Windows Server 2008 and Linux. They can also be in the form of a dedicated device with the sole purpose of handling VPN connections or as a software add-on solution to some routers. Whatever solution is used, the VPN server must have at least two network interfaces: one for the internal or company network and one that connects to the external or public network.

VPNs in a Windows Environment Windows server OSs include a VPN server solution with Routing and Remote Access Service (RRAS), a component of the Network Policy and Access Services server role. Windows Server 2008 supports three implementations of VPN:

- *Point-to-Point Tunneling Protocol (PPTP)*—A commonly used VPN protocol in Windows OSs with client support for Linux and Mac OS X, too. Most OSs that support VPN clients support PPTP.
- *Layer 2 Tunneling Protocol with IPsec (L2TP/IPsec)*—Developed in cooperation with Cisco Systems and Microsoft, L2TP with IPsec generally provides a higher level of security than PPTP. In addition to data security through encryption, L2TP/IPsec provides data integrity as well as identity verification.
- *Secure Socket Tunneling Protocol (SSTP)*—SSTP has the advantage of working behind most firewalls without firewall administrators needing to configure the firewall to allow VPN. It uses the standard port 443 used for SSL communication (HTTPS). SSTP is supported only on Windows clients starting with Vista SP1 and as a VPN server starting with Windows Server 2008. It requires the VPN server to have a valid digital certificate issued by a certification authority for server identification.

All three implementations are enabled by default when you configure Windows Server 2008 as a VPN server, so any type of client that tries to connect will be successful. VPN server configuration in Windows Server 2008 is fairly straightforward. After RRAS is installed, you configure it in the Routing and Remote Access console accessed via Administrative Tools. By default, RRAS isn't configured (see Figure 10-7); to do so, right-click the server icon in the console and click Configure Routing and Remote Access.

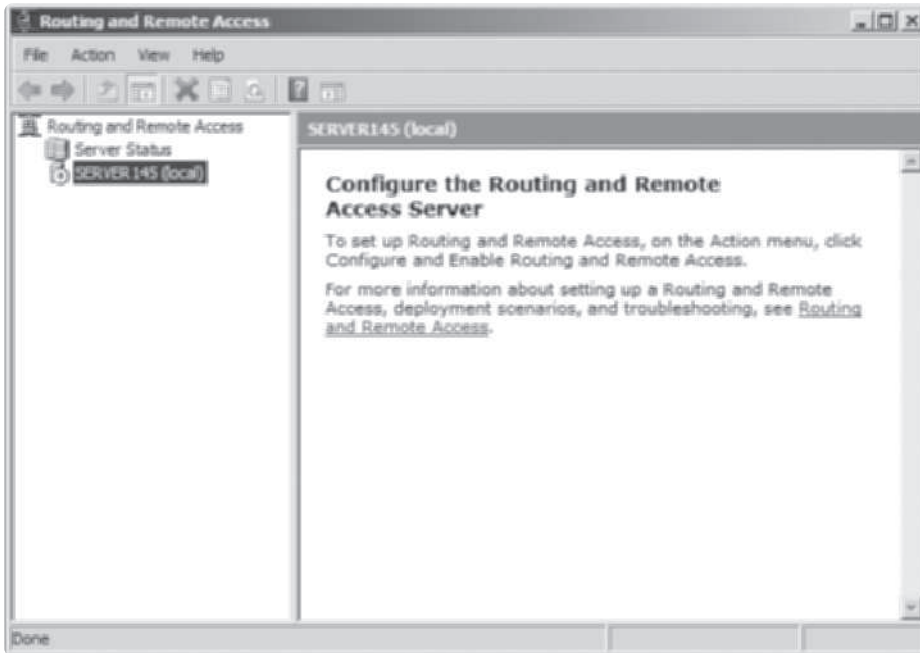


Figure 10-7 Configuring RRAS before it's enabled

Courtesy of Course Technology/Cengage Learning

VPNs in Other OS Environments Linux OSs also support VPN client and VPN server applications. Typically, they use PPTP or L2TP with IPsec. A popular free VPN solution for Linux is OpenSwan (available with documentation at www.openswan.org). OpenSwan uses IPsec as is the encryption protocol and is becoming the default in many Linux versions.

Mac OS X supports VPN client connections to Windows servers by using PPTP or IPsec. Mac OS X Server has a VPN server service that allows Mac OS X, Windows, and UNIX/Linux clients to connect to a LAN through the Mac OS X VPN server by using PPTP or L2TP.

One method of providing VPN services to connect remote sites is to use routers with VPN capability to form a router-to-router VPN connection. This type of VPN connection uses the Internet to connect remote sites with corporate headquarters or to connect corporate partners to form an extranet. In this arrangement, the VPN software is installed on the router rather than on client systems and is much less expensive than using leased or dedicated lines to connect remote networks.

VPN Benefits VPNs enable organizations to use the Internet as a private network to connect mobile users and telecommuters to corporate networks securely. Organizations can also use VPNs to connect headquarters to branch offices with permanent connections without the expense of leased lines and WAN links. To summarize, VPN benefits include the following:

- Enable mobile users to connect with corporate networks securely wherever an Internet connection is available.

- Allow multiple sites to maintain permanent secure connections via the Internet instead of using expensive WAN links.
- Reduce costs by using the ISP's support services instead of paying for more expensive WAN support.
- Eliminate the need to support dial-up remote access, which is a higher-cost solution requiring more personnel.

Protecting Networks with Firewalls

A **firewall** is a hardware device or software program that inspects packets going into or out of a network or computer, and then discards or forwards these packets based on a set of rules. A hardware firewall is configured with two or more network interfaces, typically placed between a LAN and the WAN connection. The WAN link can connect to an ISP, another LAN in another city, or even the network of a partner organization. The type of firewall you use and how it's configured are determined by what's at the other end of the WAN link. For example, you might want to allow a remote sales office to access the company's database, but you should deny this access to Internet users and perhaps restrict it for users in a partner network.

A software firewall is installed in an OS and inspects all packets coming into or leaving the computer. Based on predefined rules, the packets are discarded or forwarded for further processing. Software firewalls are an integral part of Windows, Linux, and Mac OS X, but third-party solutions are also available that sometimes have more features or more user-friendly interfaces.

A network administrator is courting disaster if a firewall isn't installed between the network and the Internet. Firewalls protect against outside attempts to access unauthorized resources, and they protect against malicious packets intended to disable or cripple a network and its resources. A second use of firewalls placed between the Internet and the network is to restrict users' access to Internet resources. This type of restriction is usually intended to prevent users from accessing offensive Web sites or bandwidth-heavy content, such as streaming audio or video, which might not be the best use of an employee's time or the network's bandwidth.

Firewalls installed on a network are usually dedicated devices with preinstalled software that must be configured by a knowledgeable administrator. This type of firewall, however, usually isn't suitable for home Internet users trying to protect their computers from would-be attackers. Because of the widespread availability of fast, always-on Internet connections for home users, personal firewalls were developed to guard a single workstation against Internet attacks. You install these software firewalls to guard your computer from attempts to access your resources and services through the Internet. Personal firewalls are not just for the home, however. Because many attacks occur inside networks, these lightweight firewalls can also be used in the office to prevent other users from infiltrating workstations or to prevent the spread of worms.

Firewall devices vary quite a bit in configuration details, but all are based on one premise: Rules are created to determine what type of traffic is allowed to enter and exit the network. A firewall, by default, is usually a closed device. After the firewall is installed and its interfaces are configured, it stops all incoming packets (and sometimes all outgoing packets, depending on the firewall). To configure a firewall, the network administrator must build

rules that allow only certain packets to enter or exit the network. The rules are based on a variety of packet properties, including source and destination addresses; protocols such as IP, TCP, ICMP, and HTTP; and sometimes even a packet's context.

Source and destination addresses can be examined to determine whether the packet is coming from an approved network or device to an approved network or device. For example, a network might have a restricted segment where no external traffic is permitted. The firewall can examine all incoming packets and discard those with a destination address of the network's restricted segment. The protocol in the packet can also be examined to determine whether it's a type that should be permitted into the network. For example, you might want to deny certain ICMP packets from entering the network. (ICMP packets are generated by the ping command, among others, and can be used to saturate a network's bandwidth or tie up a network server, thereby denying legitimate users access to the network. This is a denial-of-service [DoS] attack.) Firewalls can also attempt to determine a packet's context; this process is called **stateful packet inspection (SPI)**. SPI helps ensure that a packet is denied if it's not part of an ongoing legitimate conversation. Attackers can insert rogue packets into a data stream in an attempt to hijack a legitimate connection or tie up network services. Examining a packet's context can reduce the success of these attacks.



Firewalls perform other functions not covered here, but the functions discussed in this section are typically universal of all firewalls.

Using a Router as a Firewall Conceptually, a firewall is just a router with specialized software for creating rules to permit or deny packets. Many routers have capabilities similar to firewalls but with one key difference: Routers, by default, are open systems. After a router is first configured, all packets are permitted into and out of the network. Therefore, a network administrator must create rules, called access control lists (ACLs), that deny certain types of packets. Typically, an administrator builds ACLs so that all packets are denied, and then creates rules that make exceptions. ACLs can examine many of the same packet properties that firewalls can.



Routers intended for SOHO use that often combine a wireless AP, router, and switch into a single unit have a preconfigured firewall. These routers are designed to connect a small LAN to the Internet and configured to allow all traffic to leave the network and deny traffic coming into the network that isn't part of a communication session initiated by a computer inside the LAN.

Using Intrusion Detection Systems An **intrusion detection system (IDS)** usually works with a firewall or router with ACLs. A firewall protects a network from potential break-ins or DoS attacks, but an IDS must detect an attempted security breach and notify the network administrator. In some cases, an IDS can take countermeasures if an attack is in progress. These countermeasures include resetting the connection between source and destination devices or even disabling the link between inside and outside networks. An IDS is an invaluable tool to help administrators know how often their network is under attack and devise security policies aimed at thwarting threats before they have a chance to succeed.

Using Network Address Translation to Improve Security Network Address Translation (NAT) was discussed in Chapter 5 in the context of alleviating the IP address shortage. An additional benefit of NAT is that an internal network resource’s real address is hidden and inaccessible to the outside world. Because most networks use NAT with private IP addresses, devices configured with private addresses can’t be accessed directly from outside the network. In fact, when NAT is used, the only way an outside device can send a message to a device in the internal network is in response to a message from the internal device. That is, an external device can’t initiate a network conversation with an internal device, thus limiting an attacker’s options. NAT is usually an integral part of a network firewall positioned between the network and the Internet or another outside network.



Hands-On Project 10-3: Exploring Windows Firewall

Time Required: 15 minutes

Objective: Explore the Windows 7 firewall.

Required Tools/Equipment: Your classroom computer

Description: In this project, you learn how to configure the Windows 7 firewall.

1. Log on to your computer as an administrator.
2. Click **Start**, **Control Panel**, and click **System and Security**. Click **Windows Firewall**.
3. Figure 10-8 shows the Windows Firewall summary window, which displays the current state of the firewall for connected networks. Firewall rules are different for private networks (home or work networks) and public networks. The “Change notification settings” and “Turn Windows Firewall on or off” links in the left pane bring you to the Customize Settings window. In the left pane, click **Turn Windows Firewall on or off**.

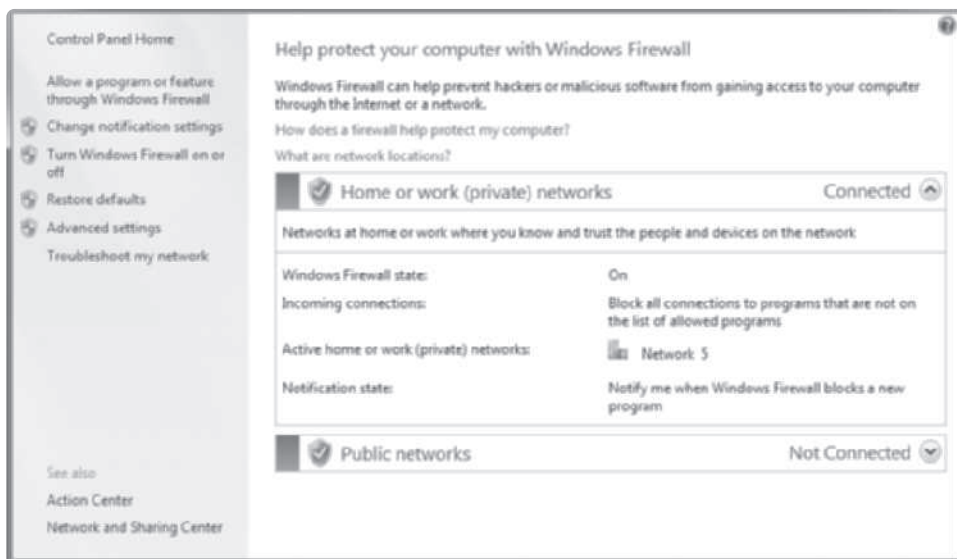


Figure 10-8 Windows Firewall

Courtesy of Course Technology/Cengage Learning

4. In the Customize settings for each type of network window, you can turn the firewall on or off for the public and private network. If the firewall is on, you can block all incoming connections and select whether you should be notified when the firewall blocks a program. Click **Cancel**.
5. Click **Restore defaults** to return your firewall settings to their original status. Click **Restore defaults** again, and click **Yes** in the Restore Defaults Confirmation message box.
6. Click **Advanced settings** in the left pane to open the Windows Firewall with Advanced Security console (see Figure 10-9). You can also open this console from Administrative Tools. The middle pane shows an overview of your current settings.

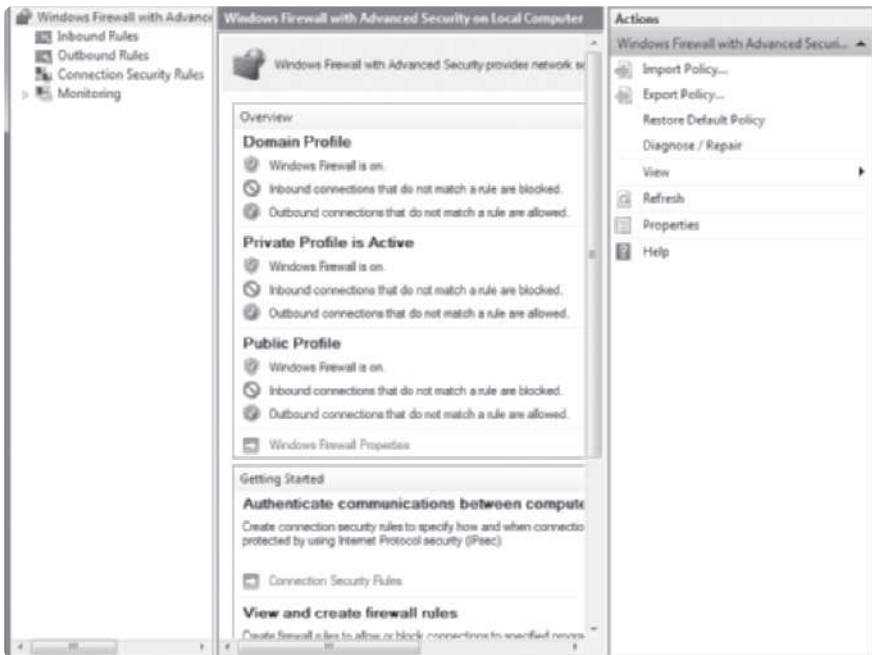


Figure 10-9 The Windows Firewall with Advanced Security console

Courtesy of Course Technology/Cengage Learning

7. In the Overview section, click the **Windows Firewall Properties** link. You see four tabs. The first three tabs are used to configure the Windows Firewall domain, private, and public profiles. The fourth tab has IPsec settings. Click each profile tab. Notice that inbound connections are set to Block and outbound connections are set to Allow. These settings allow any traffic out of your computer but deny unsolicited attempts from the outside to connect to your computer. Click the **Learn more about these settings** link at the bottom. Browse through the help window to learn more about the Domain, Private, and Public profile tabs. Close the help window, and click **Cancel**.
8. In the Windows Firewall with Advanced Security console, click **Inbound Rules** in the left pane. A list of rules is displayed in the middle pane. Each rule has a name describing

the type of network activity the rule pertains to, a group, a profile in which the rule is defined, a status (enabled or disabled), an action, and several other properties.

9. Scroll down the Inbound rules pane until you see File and Printer Sharing (Echo Request - ICMPv4-In). This rule controls whether another computer can ping your machine and is disabled by default. You should see the same rule for both the Private and Domain profiles. Ask a partner to try to ping your computer. The attempt should be unsuccessful because the default inbound setting is to block inbound traffic, such as Echo Requests (the type of packet Ping produces).
10. Double-click the **File and Printer Sharing (Echo Request - ICMPv4-In)** option with Private listed in the Profile column. Click **Enabled**, and then click **OK** to enable the rule, which is set to allow the connection. Have a partner try to ping you again. This time the ping should be successful. (On a domain-based network, you would enable the rule for the Domain profile.)
11. Windows Firewall has dozens of predefined rules, and you can create your own rules. In the Actions pane, click **New Rule**.
12. In the Rule Type window, click **Port**, and then click **Next**. In the Protocol and Ports window, leave **TCP** selected. In the Specific local ports text box, type **80**, and then click **Next**. This setting allows other computers to connect to your computer via TCP port 80, which is the standard Web server port.
13. In the Action window, leave **Allow the connection** selected, and then click **Next**. In the Profile window, leave all three profiles selected, and then click **Next**.
14. In the Name window, type **Web Server In** in the Name text box, and then click **Finish**. Scroll to the top of the Inbound Rules section to see your new rule. If you installed a Web server now, users would be able to connect to it.
15. Close all open windows, but leave your computer running for the next project.



Hands-On Project 10-4: Testing Your Firewall with ShieldsUP!

Time Required: 15 minutes

Objective: Connect to a Web site to run a free Internet vulnerability test.

Required Tools/Equipment: Your classroom computer with access to the Internet

Description: In this project, you connect to a Web site to run a free Internet vulnerability test. This Web site does a simple penetration test on your computer. If you're running it behind a NAT firewall that assigns a single public IP address to all outside connections, only one person at a time can run the ShieldsUP! test. The instructor might want to perform this project as a demonstration and encourage students to run this test on their home computers.

1. Log on to your computer as an administrator, if necessary.
2. Start your Web browser and go to **www.grc.com**. Click the **Shields UP!!** graphic.
3. Scroll down and click the **ShieldsUP!** link.

4. Read the information that's displayed, and then click the **Proceed** button. (Click **Continue** if prompted.)
5. In the ShieldsUP!! Services box (see Figure 10-10), click **File Sharing**. ShieldsUP! displays your IP address or the translated IP address, and then attempts to connect to your computer. Port 139 should not be available; if it is, you have a serious security problem because this port is used to access files on your computer.

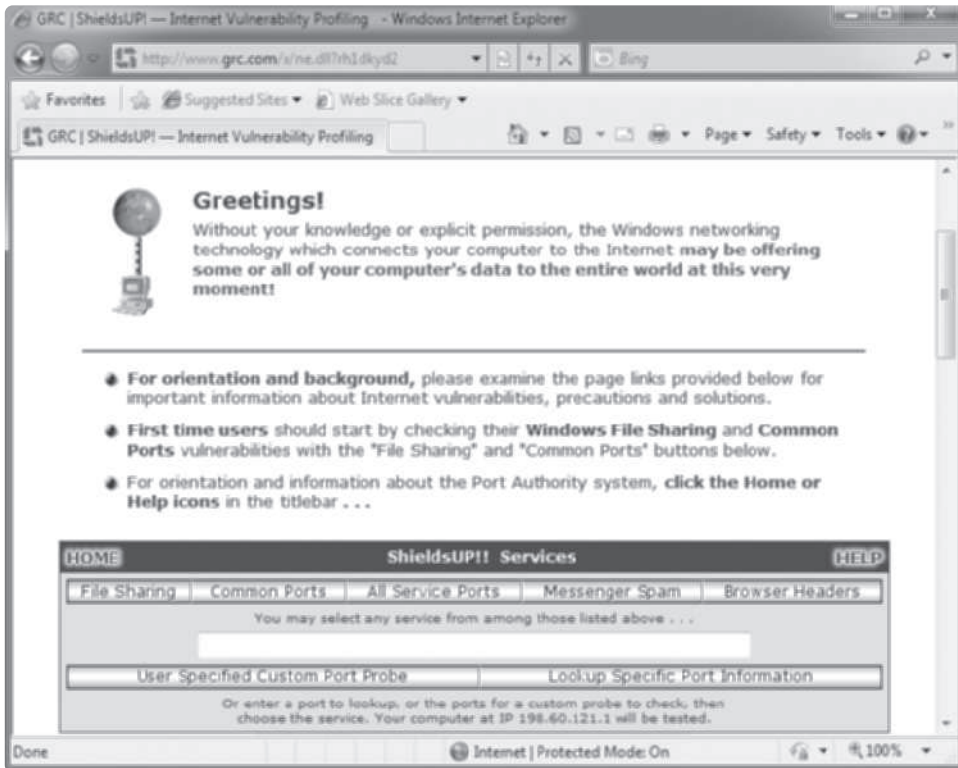


Figure 10-10 ShieldsUP! vulnerability testing

Courtesy of Course Technology/Cengage Learning

6. Next, click **Common Ports** and see whether any other ports are available through the Internet.
7. Continue clicking available ShieldsUP! options and write a summary of your findings on the following lines. If ShieldsUP! finds any security vulnerabilities, it makes recommendations for securing your computer.

8. Exit your Web browser, and leave your computer running for the next project.

Protecting a Network from Worms, Viruses, and Rootkits

In Internet-connected networks, virus and worm attacks are a constant threat. Users download programs, bring disks from home, and open e-mail attachments. All these actions are normal computing activities, but they can also bring viruses into the network. A **virus** is a program that spreads by replicating itself into other programs or documents. Its sole purpose is to disrupt computer or network operation by deleting or corrupting files, formatting disks, or using large amounts of computer resources. A **worm** is similar to a virus in that it's self-replicating, but a worm doesn't attach itself to another program; it's a self-contained program. Worms are now more common than viruses because with the Internet and widespread network connectivity, they don't need help to spread. A virus requires a user to run the program containing it to operate, and then copy the file to spread, but a worm can do its work without any help and spread via an available network connection. Some insidious actions a worm can perpetrate include using large amounts of network bandwidth, deleting files, sending e-mails, and creating backdoors into computers. A **backdoor** is a program installed on a computer that permits access to the computer, bypassing the normal authentication process. A common use of backdoors is to allow spammers to send e-mail from computers on which they're installed, thereby hiding a spammer's true identity.

Viruses, worms, and rootkits are part of a broader category called **malware**, which is any software designed to cause harm or disruption to a computer or perform activities on a computer without the owner's consent. To help prevent the spread of viruses and worms, every desktop and server should have virus-scanning software running. Most virus-protection software is also designed to detect and prevent worms. A virus scanner residing in memory should be used so that every program file or document that's accessed is scanned. Documents should be scanned if the document type might contain macros, and servers should run virus-protection software that scans every file read or written to and from servers' drives. If a server file accessed by other users on the network gets infected, the malware can spread through the network in a matter of seconds.

Viruses and worms that spread through e-mail attachments have been common for years. They're simple to avoid; just don't open any e-mail attachments sent by someone from whom you're not expecting a message. Even if you know the sender, be cautious; malware can use an e-mail program's address book to send messages, causing you to believe the message is safe. Most virus scanners actually detect a virus or worm contained in an e-mail message and often delete the attachment before it ever reaches your inbox, but if the virus is very new, it might not be detected.

Another type of malware that's not technically a virus because it's usually not self-replicating is a **Trojan program**. It appears to be something useful, such as a free utility, but in reality contains some type of malware. What's unfortunate about a Trojan program is that users willingly run the software and don't even know it's causing problems on their systems.

Rootkits are a form of Trojan programs that can monitor traffic to and from a computer, monitor keystrokes, and capture passwords. They're the ultimate backdoor into a system and are among the most insidious form of Trojan programs because they can mask that the system has been compromised by altering system files and drivers required for normal computer operation. Rootkits aren't specific to an OS and can be found for Windows, Linux, and forms of UNIX. They're notoriously difficult to detect because they hide themselves so well and integrate into the OS they have infected. Typically, detection requires restarting the

system and booting to another medium, such as a CD or flash drive, with tools that can scan for and detect a rootkit's presence. Removal is even more difficult because rootkits often alter system files and drivers the system depends on to run normally. Many experts agree that the time and effort required to remove a rootkit is better spent backing up critical data files, reformatting the disk, and reinstalling the OS.

A **hoax virus** is one of the worst kinds of viruses. With a hoax virus, someone sends an e-mail proclaiming that Microsoft, the government, or another well-known entity has just discovered a new virus that reformats your hard drive or performs another nefarious deed. The hoax message goes on to say that you should forward this e-mail immediately to everyone you know to inform them of this terrible virus. The flood of e-mail from people actually falling for this hoax *is* the virus! It clogs e-mail servers, decreases productivity, and generally wastes time. Although a hoax virus doesn't actually destroy data, it uses a tactic called **social engineering**, in which attackers get users to do their bidding without being aware of the consequences. If you're concerned that the warning might be real, check the Web site of the organization the message references or the Web site of your antivirus software. (You do have antivirus software installed, don't you?) If the supposed virus isn't mentioned at these sites, stop this type of virus in its tracks and delete the e-mail without forwarding it to innocent friends and acquaintances.

**TIP**

In the Computer section of the www.snopes.com Web site, you can find a list of real and hoax viruses.

Virus and worm protection can be expensive, although many quite capable freeware or shareware packages are available. However, the loss of data and productivity that result when a network becomes infected is much more costly. Remember that virus-protection software must be updated because developers of viruses and worms are always looking for new and clever ways to wreak havoc on networks.

Protecting a Network from Spyware and Spam

Spyware and spam aren't that similar in function, but both affect your privacy, and their primary goal is to get you to buy something or get taken by a fraud. **Spyware** is a type of malware that monitors or controls part of your computer at the expense of your privacy and the gain of some third party. The result of spyware is usually a decrease in computer performance and an increase in pop-up Internet messages and spam. The goal is to monitor your Internet activity, such as which Web sites you visit and how often. The data the spyware gathers is then used by advertisers, spammers, and perhaps even more malicious third parties for the sole purpose of extracting money from your wallet.

Unlike a virus or worm, spyware isn't usually self-replicating. Typically, it's installed on a system when a user installs some legitimate software or is too quick to click OK when a message pops up on a Web site offering to install a program. Many free peer-to-peer file-sharing applications install spyware on your computer as a condition of being free. Nonetheless, millions of users install the software (*and* the spyware) because the prospect of being able to download free music and software is just too compelling.

Many antispyware programs are available, and some are bundled with popular antivirus programs. Microsoft offers Windows Defender with Windows 7 and Vista and as a free

download with older versions. Microsoft Security Essentials, a free download, provides anti-virus and antispyware software. Antispyware programs scan your computer for known spyware and remove it, and some can offer real-time protection to stop spyware from being installed in the first place. Note, however, that some freeware programs installed on your computer might stop working if the spyware that came packaged with them is removed.

Spam, like spyware, is more a nuisance than a threat to your computer. It's simply unsolicited e-mail. Although spam doesn't delete files or format disks, it's a thief of e-mail storage space, network bandwidth, and, most important, people's time. For those naive enough to click where spam leads, it can also be a thief of your hard-earned cash if you end up purchasing products or fall for frauds. Like spyware and virus protection, spam detection and prevention are uphill battles because for every rule or filter antispyware software places on your e-mail account, spammers find a way to get around them.

Probably one of the best ways to avoid spam is to not give your e-mail address to anyone but trusted parties. If you must register on a Web site with an e-mail address, use one from a free e-mail service that you never use for personal mail. That way, you can simply log on to the free e-mail Web site periodically and delete all the messages. Unfortunately, this method still doesn't guarantee your protection from spam, as even legitimate organizations you communicate with regularly can sell their e-mail lists or have them stolen. In addition, worms and spyware can use the address books of people you know to get access to your e-mail address.

It should be clear by now that the Internet, with its wealth of information and its avenues of entertainment and business, is also a dangerous place. The best advice, in lieu of pulling the plug on your Internet connection, is to be dutiful in keeping anti-everything software up to date, and use common sense when opening e-mails or responding to Web-based solicitations. Network security is effective only when users understand the risks of installing and using certain types of software and have a solid understanding of the organization's security policies. A well-educated workforce is a safe workforce.

**TIP**

Before purchasing or downloading anti-malware software, it's best to read reviews published by reputable print or Internet publishers. Some software billed as anti-malware is actually malware that can disable your existing anti-malware software and thwart your attempts to uninstall it or install legitimate anti-malware software.

**HANDS-ON PROJECTS**

Hands-On Project 10-5: Installing Microsoft Security Essentials (Optional)

Time Required: 20 minutes

Objective: Install Microsoft Security Essentials.

Required Tools/Equipment: Your classroom computer with access to the Internet

Description: If you already have up-to-date antivirus and antispyware software that works well, you might not want to do this project. However, if you have older software that can no longer be updated, or none at all, you should do this project.

1. Log on to your computer as an administrator.
2. Start your Web browser and go to www.microsoft.com/security_essentials/.

3. Click **Download Now**. When prompted, choose your language and click the link for the OS you're running. (If you're using Internet Explorer, you might not get this prompt.) Microsoft Security Essentials supports Windows XP 32-bit and 32-bit and 64-bit versions of Windows Vista and Windows 7.
4. Run the file or save it and run it when the download is finished. Follow the prompts to install the software. In the final installation window, make sure the **Scan my computer for potential threats after getting the latest updates** check box is selected, and then click **Finish**.
5. Microsoft Security Essentials starts. Click the **Update** button if you get the message "Virus & spyware definitions status - Out of date." Microsoft Security Essentials then begins scanning your computer.
6. While the scan is running, click the **Settings** tab. Set the scanning schedule to a time when your computer is turned on but you aren't using it heavily. Once-a-week scans are probably enough unless you regularly visit risky Web sites, such as some gaming sites and "free" music and software download sites, among others.
7. In the left pane, click **Real-time protection**. Make sure real-time protection is on. Browse through the other settings to see the software's configuration options.
8. After the scan is finished, exit the program. (It continues to run in the background.) Close all open windows, and leave your computer running for the next project.

Implementing Wireless Security

The explosion of wireless networking devices creates a new problem for network administrators. Because wireless signals aren't bound by physical cables, an attacker doesn't need physical access to network cabling to compromise a network. Anyone with a wireless scanner and some software who gets within range of your wireless network's signals can intercept data or access wireless devices. What's worse, because most wireless networks eventually tie into a wired network, an attacker potentially has access to your entire network infrastructure while sitting in a car outside your building. Attackers who drive around looking for wireless LANs (WLANs) to intercept are called **wardrivers**.

To foil would-be wireless attackers, wireless security must be enabled on all your networking devices by using one or more of the following methods:

- *Service set identifier (SSID)*—An SSID is an alphanumeric label configured on the access point (AP) that identifies one WLAN from another. Each client must configure its wireless NIC for that SSID to connect to the AP. A private WLAN should set the SSID to a value that's not too easy to guess, and the SSID shouldn't be set to broadcast. When an SSID is broadcast, wireless software can scan the network to look for available SSIDs, allowing an attacker to gain the first piece of information required to access your WLAN. Hiding the SSID doesn't stop a seasoned wardriver, but it can at least discourage neighbors from accidentally or purposely trying to connect to your network. For true Wi-Fi security, always use the strongest encryption protocol available.
- *MAC address filtering*—If your wireless network is fairly small and only certain computers are to have access to the network, you can use the **MAC address filtering** feature on APs to restrict network access to computers with specific MAC addresses. This security measure isn't viable in a large or nonstatic network where new laptops

and PDAs access the network frequently. In addition, it shouldn't be used alone because experienced wardrivers can thwart it fairly easily.

- *Wired Equivalent Privacy*—This option must be set at the AP and on the wireless client. **Wired Equivalent Privacy (WEP)** provides data encryption so that a casual attacker who gains access to your wireless signals sees only encrypted data. However, WEP has its flaws, and a determined attacker can crack the encryption code fairly quickly. Nonetheless, WEP has the advantage of being available in just about all wireless equipment, so you can have some security without buying anything new. However, if your devices support only WEP, it's time to consider an upgrade.
- *Wi-Fi Protected Access*—**Wi-Fi Protected Access (WPA)**, the successor to WEP, has enhancements that make cracking the encryption code more difficult. WEP uses a static encryption key, but WPA alters the key periodically and automatically, so even if an attacker does determine the key, the key soon changes and the attacker must start over. WPA can also use an enhanced authentication method, with a centralized server maintaining the database of users permitted to access the wireless network.
- *802.11i*—The **802.11i** standard, an extension to 802.11, was ratified in 2004 and is the latest standard that defines wireless security. 802.11i is usually referred to as WPA2 because it incorporates much of the WPA standard. Its advantage over WPA, however, is that it uses more advanced encryption standards and a more secure method of handling encryption keys. If your devices support WPA2, use it, as it offers the best security for your Wi-Fi network.

A final word about wireless networking security: Using a strong encryption protocol doesn't mean you're completely safe. Determined attackers can and will get into any network if they want to. With that in mind, you should follow some policies that further protect your network:

- Do a site survey and try to position your APs so that only required areas are covered by the signal; limit signal access outside the building whenever possible.
- If you're using WEP, change the encryption key manually on a regular basis.
- When possible, use APs that can filter MAC addresses, allowing only known addresses access to the network.

Remember that the benefit of a wireless network is having easy access to a corporate or home LAN without being tethered to a cable, but this access also applies to attackers wanting to harm your network.

Using an Attacker's Tools to Stop Network Attacks

If you want to design a good, solid network infrastructure, hire a security consultant who knows the tools of the network attacker's trade. The terms "black hats" and "white hats" are sometimes used to describe people skilled in breaking into or disabling a network. A black hat is, as the analogy implies, the bad guy, and a white hat is the good guy. White hats often use the term **penetration tester** for their consulting services. In fact, a certification has been developed for white hats called Certified Ethical Hacker (CEH; www.eccouncil.org). This section approaches the subject of network security from the white hat's perspective. The goal is to see what type of holes exist in a network's security and close them.





The term **cracker** is sometimes used to describe a person who attempts to break in, disable, or otherwise attack a network. Sometimes confused with a hacker, a cracker attempts to compromise a network or computer for personal gain or to cause harm. Contrast this term with **hacker**, which has had several meanings over the years. It's sometimes a derogatory term to describe an unskilled or undisciplined programmer. It can also mean someone who's highly skilled with computer systems and programs and can use some of the same tools crackers use to poke around networks or systems, but not for evil purposes. For simplicity's sake, the term attacker is used in this book to mean a person who tries to compromise a network for nefarious purposes.

Discovering Network Resources Before attackers can gain access to or cause problems with your network, they must get information about the network configuration and available resources. Some tools they use are command-line utilities, such as Ping, Trace Route, Finger, and Nslookup. These commands can help you determine which devices are available, identify name information for these devices, and possibly learn user information. Ping, as you have learned, can be used to determine whether a particular computer is responding on the network. Because you can ping a computer by name and have its IP address returned, it can also be used to resolve a computer name to an IP address. Trace Route provides information about the route a packet takes from one computer to another, which can help determine an internetwork's structure. With Finger, you can query a computer and determine who's logged on to it and its address. Nslookup is used to query DNS servers. Depending on how well a DNS server is secured, you can use this command to retrieve a list of all computer names and mail servers on a domain.

Other tools of the trade include ping scanners and port scanners. A **ping scanner** is an automated method for pinging a range of IP addresses. A **port scanner** determines which TCP and UDP ports are available on a computer or device. With a ping scanner, you can enter a network address, and the program queries all IP addresses in that network (or a range of IP addresses). Many ping scanners also look up the DNS name of any computer that responds. Attackers use this information to see what computers are available on a network, and the DNS name can provide useful information because most network administrators name devices to describe their purpose or location, such as naming a database server SQL-Server or a router Router-3rdFloor. Figure 10-11 shows the results of a ping scan.



A ping scan can be done from the Windows command line by using the command `FOR /L %I IN (1,1,254) DO PING -n 1 -w 100 192.168.1.%I > pingresults.txt`. The results are sent to the pingresults.txt file.

A port scanner, by determining which ports are active, can tell you what services are enabled on a computer. Figure 10-12 shows running a port scan on a computer with the IP address 198.60.123.101. Most services are closed, but several are open. A network administrator should use this information to be sure ports listed as open are necessary for the computer's operation. Any unnecessary ports should be closed, which usually involves stopping a service or an application from running or configuring Windows Firewall to deny access to the port.

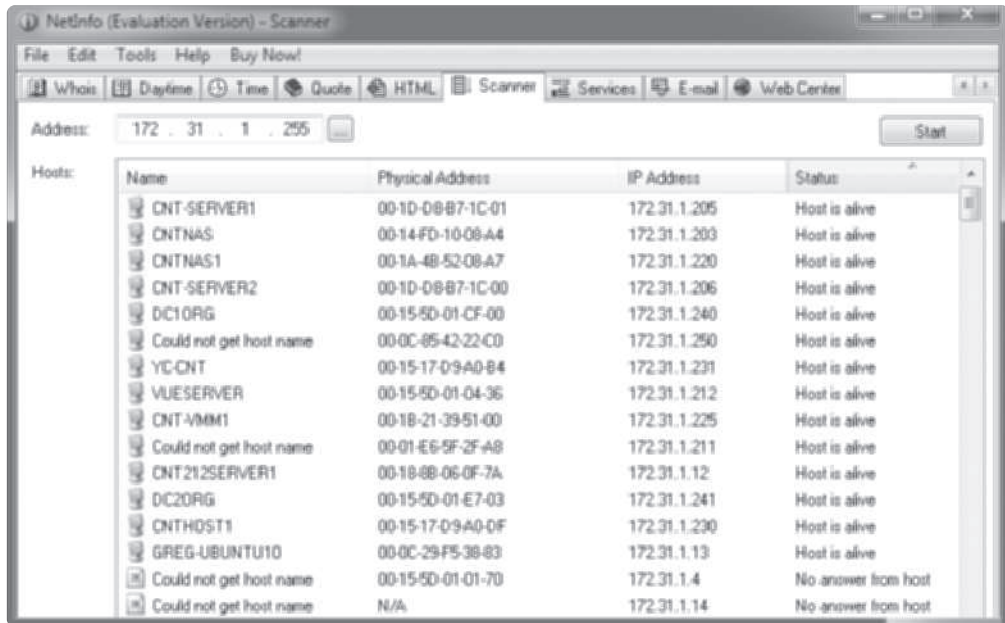


Figure 10-11 The results of a ping scan on a network

Courtesy of Course Technology/Cengage Learning

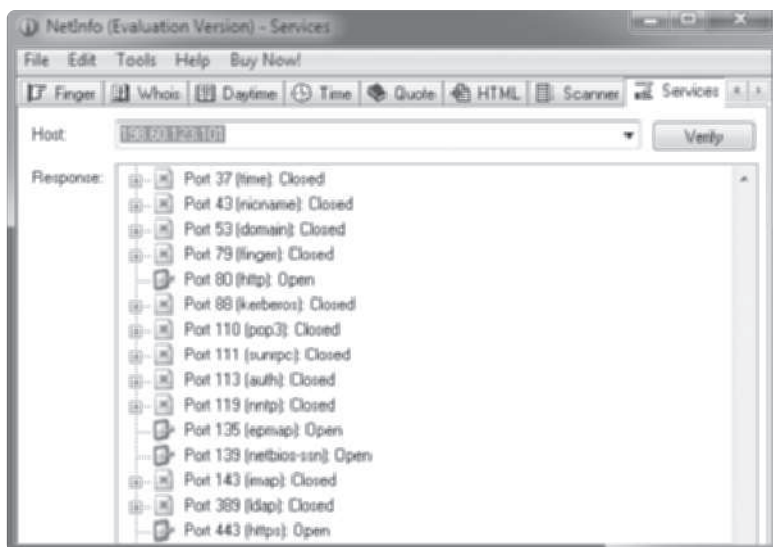


Figure 10-12 The results of a port scan on a computer

Courtesy of Course Technology/Cengage Learning

Whois is a handy utility for discovering information about an Internet domain. You can find the name and address of the domain owner, contact information for the domain, and the DNS servers managing the domain (see Figure 10-13). The information that can be gathered from a Whois query includes IP address information and names and addresses of DNS

servers used by that domain. DNS servers can also be queried to determine names and addresses of computers in that domain.

NetworkSolutions | LEARNING CENTER | PRODUCTS & SERVICES | ACCOUNT MANAGER | CUSTOMER SUPPORT

Back to Home

WHOIS Search Results

WANT A DOMAIN NAME THAT'S ALREADY REGISTERED?
 Use our Backorder service to get that name as soon as it expires. [» TRY IT NOW](#)

WHOIS Record For

course.com

Certified Offer Service - Make an offer on this domain
 Backorder - Try to get this name when it becomes available
 Private Registration - Make personal information for this domain private
 SSL Certificate - Make this site secure
 Site Confirm Seals - Become a trusted Web Site

Registrant: Thomson Learning
 5191 Nabors Blvd
 Mason, OH 45040-7845
 US

Domain Name: COURSE.COM

Administrative Contact, Technical Contact:
 Thomson Learning
 domain_admin@thomsonlearning.com
 5191 Nabors Blvd
 Mason, OH 45040-7845
 US
 Phone: 513-228-1101
 Fax: 513-229-1002

Record expires on 28-Jul-2007
 Record created on 31-Jul-1997
 Database last updated on 28-Jun-2005

Domain servers in listed order: [Manage DNS](#)

NS1.THOMSONLEARNING.COM	199.80.140.11
NS2.THOMSONLEARNING.COM	199.80.131.2

[Show underlying registry data for this record](#)

Current Registrar: NETWORK SOLUTIONS, LLC.
IP Address: 199.80.146.38 (ARIN & RIPE IP search)
IP Location: US/UNITED STATES/MASSACHUSETTS-BOSTON
Record Type: Domain Name
Server Type: NS 5
Lock Status: REGISTRAR-LOCK
Web Site Status: Active
DMOZ: 2 listings
Y! Directory: see listings
Secure: No
E-commerce: Yes
Traffic Ranking: Not available
Data as of: 21-Oct-2005

View Order

BUY THE AVAILABLE EXTENSIONS FOR THIS DOMAIN NAME

COURSE .biz
 COURSE .org
 COURSE .net
 COURSE .info
 COURSE .us

[Continue >](#)

SEARCH AGAIN

Enter a search term:

e.g. network-solutions.com

Search by:
 Domain Name
 NIC Handle
 IP Address

[Search >](#)

Figure 10-13 Results returned from a Whois query

Courtesy of Course Technology/Cengage Learning

Protocol analyzers are also useful for resource discovery because they enable you to capture packets and determine which protocols services are running. They require access to the network medium and are, therefore, effective tools only if the attacker is an internal user or has gained access to the internal or wireless network.

To protect your network from some of these utilities, you can take a variety of actions. Some utilities, such as Finger, can be rendered useless if they're turned off on all devices that support them. Some Linux and UNIX systems as well as some routers often leave the Finger service on by default. A port scan should be run on all network devices to see which

services are on, and then services that aren't needed should be turned off. This process is a white-hat use of a port scanner.

Access lists on routers and firewalls, including personal firewalls, can block pings to prevent the use of ping scanners. To protect the network from internal users of protocol analyzers, all hubs and switches should be secured to prevent an unauthorized user from hooking up a laptop or other device to the network.

Gaining Access to Network Resources After an attacker has discovered the resources available on a network, the next step might be gaining access to these resources for the purposes of viewing, stealing, or destroying data. One of the easiest resources to open is one for which no password is set. Believe it or not, this situation happens more often than you think, and often numerous routers and switches with no passwords set are available through the Internet or on a company network. The remedy to this problem is, of course, to check all devices supporting Telnet, FTP, e-mail, and Web services, verify that passwords are set on them, and disable any unnecessary services.

Often an attacker runs into a resource that requires a username and password. Finger can be used in some cases to discover usernames, and Linux and Windows servers have default administrator names that are often left unchanged—a fact that an attacker with a password-cracking tool can exploit easily. Some password-cracking tools use a systematic method of guessing passwords from a dictionary of words or from an algorithm that uses all combinations of letters, numbers, and symbols. This type of tool can be extremely time and CPU intensive. If passwords are strong, these tools are often impractical because guessing complex passwords can take days. Using a password-cracking tool on your own system is recommended to see whether your passwords are complex enough.



For an extensive list of security and hacking tools, including password crackers, visit www.securiteam.com/tools/.

TIP

Disabling Network Resources A denial-of-service (DoS) attack is an attacker's attempt to tie up network bandwidth or network services so that it renders resources useless to legitimate users. Some attackers launch a DoS attack for fun; others do it to satisfy a grudge or even gain a leg up on the competition. Three common types of DoS attacks focus on tying up a server or network service: packet storms, half-open SYN attacks, and ping floods. Programs that can create these attacks are readily available for download.

Packet storms typically use the UDP protocol because it's not connection oriented. One packet storm program called Pepsi5 sends a stream of UDP packets that have spoofed host addresses, causing the host to be unavailable to respond to other packets. A **spoofed address** is a source address inserted into the packet that isn't the sender's actual address.

Half-open SYN attacks use the TCP three-way handshake to tie up a server with invalid TCP sessions, thereby preventing real sessions from being started. The attacker sends a series of packets with a valid port number and a request to start a conversation. These packets, called SYN packets, cause the server to respond. The original SYN packet contains a spoofed source address, resulting in the server waiting for the final packet in the three-way handshake until it times out. If enough SYN packets are sent, the server uses all available

connections and, therefore, can't respond to legitimate attempts to make a connection. Several programs that create this type of attack are available.

A ping flood is exactly what it sounds like. A program sends a large number of ping packets to a host. They cause the host to reply, which ties up CPU cycles and bandwidth. A variation is the smurf attack, in which pings are sent to a broadcast address. All the requests contain the spoofed source address of the host to be smurfed. When computers respond to the broadcast ping, they send their replies to the single host whose address is spoofed. The host is then flooded with ping responses, causing it to slow down or even freeze while it processes all the packets.

Distributed denial-of-service (DDoS) attacks use many systems to attack a single network or resource. Often the attacking systems are unaware they're involved because the attack software is installed as malware and set to activate on a certain date and time. MyDoom, the fastest-spreading worm ever at the time, is a well-known example of a DDoS attack; it targeted *www.sco.com*, among other sites.

There's no end to the methods for wreaking havoc on a network. Becoming familiar with the tools and methods that can be used against your network is essential so that you can prepare defenses against network attacks. You can also use these tools to test the integrity of your network security. Firewalls, access lists, virus scanners, and strong OS security are some ways to prevent these attacks or reduce their effects. In addition, using an IDS helps you analyze attempts to breach network security and track down and close potential holes in your security measures. Regardless of your tools, you should always start by devising a sound security policy that maps out your overall network security plan and contains provisions for auditing and revising the policy as your needs and technology change. Implementing your policies and using the tools available to protect your network keep your data safe and keep you sleeping well at night.



Network security is a complex topic. You're encouraged to take the knowledge you have learned in this book and study network security in more depth by reading books or taking a security class. Several network security certifications can be earned, such as the Security+, Systems Security Certified Practitioner (SSCP), and many others.



Hands-On Project 10-6: Using NetInfo to Collect Computer and Network Information

Time Required: 15 minutes

Objective: Install the NetInfo program to collect information about a network.

Required Tools/Equipment: Your classroom computer with Internet access

Description: In this project, you download and install an evaluation version of NetInfo, used to scan computers for open ports and scan a network for IP addresses. The download is a zip file containing a Microsoft installer (.msi) file that must be extracted to run it. Alternatively, the instructor can download the file, extract the installer, and make it available to students on the local network. The software can also be preinstalled on students' computers, and students can begin this project at Step 3.

1. Log on to your computer as an administrator, if necessary.
2. Start your Web browser and go to <http://netinfo.tsarfin.com>. Click the **download** link and then the **Download Now** link. You're directed to the Cnet download page. Click **Download Now** again. Follow the instructions to download and install NetInfo.
3. To find out about other computers on your network, open a command prompt window, and then type **net view** and press **Enter**. Write down several of the computer names this command returns:

4. Choose one of the computer names you wrote down in Step 3. Type **ping *computername*** (replacing *computername* with the name of the computer you chose) and press **Enter**.
5. Write down the IP address of the computer returned by the **ping** command:

6. To start NetInfo, double-click the **NetInfo** desktop shortcut or click **Start**, point to **All Programs**, point to **NetInfo**, and click **NetInfo**. Click **No** in the message box stating you can try NetInfo for 30 days, and then click **I Agree**. If the Tip of the Day dialog box opens, click to clear the **Show tips at startup** check box, and then click **Close**.
7. Click the **Services** tab. In the Host text box, type the IP address you wrote down in Step 5 and click **Verify**. (If you're doing this project yourself, you can type 127.0.0.1 in the Host text box to scan your own computer).
8. Find ports that show a status of Open (see Figure 10-14). Open ports represent network services the computer offers but can also represent vulnerabilities that attackers can exploit. Write the name and number of these ports:

9. You'll use the information from Step 8 in a Case Project at the end of this chapter. Clear the output from the last command by right-clicking in NetInfo, pointing to **Clear**, and clicking **All**. Next, you'll scan a range of IP addresses to see which computers are available on the network.
10. In the Address text box, type the first three octets of the IP address you used in Step 7, followed by a 0 for the last octet. For example, if the address you used was 192.168.1.55, type 192.168.1.0. This setting scans all addresses from 192.168.1.0 through 192.168.1.255.
11. Click **Start**. Write down the name and address of the first three computers for which NetInfo indicated the status "Host is alive." (You can sort the results by clicking the **Status** column.)

12. In the Name column, right-click one of the computers, point to **Send To**, and click **Services**. In the Services tab, click **Verify** to see a list of services this computer provides.



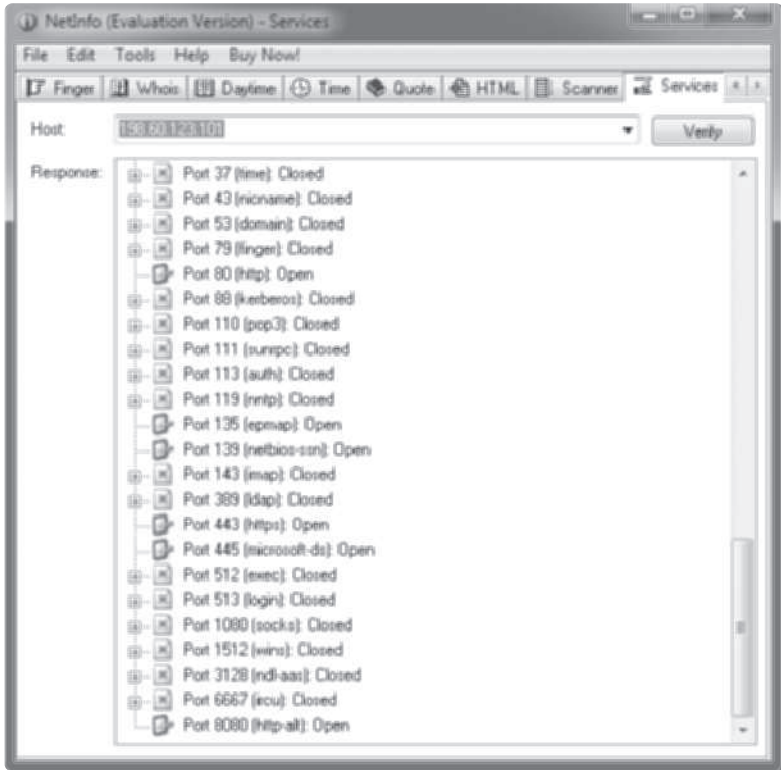


Figure 10-14 NetInfo scans for open ports or services

Courtesy of Course Technology/Cengage Learning

13. Write a short explanation of how NetInfo’s Scanner and Services features could help an attacker:

14. Close NetInfo and all other open windows.

Chapter Summary

- A network security policy is a document that describes the rules governing access to a company’s information resources. A security policy should be easy to understand and enforce and should state each policy objective clearly.
- A security policy should contain these types of policies: privacy policy, acceptable use policy, authentication policy, Internet use policy, auditing policy, and data protection policy.

- Securing physical access to network resources is paramount. Separate rooms or locking cabinets should be available to house network servers and equipment. Wiring should be inaccessible to eavesdroppers. Physical security includes procedures to recover from natural disasters.
- Securing access to data includes authentication and authorization, encryption, VPNs, firewalls, virus and worm protection, spyware protection, and wireless security.
- VPNs are an important aspect of network security because they secure remote access to a private network via the Internet.
- Firewalls, a key component of any network security plan, filter packets and allow or deny packets based on a set of defined rules.
- Malware encompasses viruses, worms, Trojan programs, and rootkits. Malware protection should be a required element on every computer and network.
- Wireless security involves configuring a wireless network's SSID correctly and configuring and using wireless security protocols, such as WEP, WPA, or 802.11i.
- Tools that attackers use to compromise a network, such as ping scanners, port scanners, and protocol analyzers, can also be used to determine whether a network is secure.
- Denial of service is one method attackers use to disrupt network operation. Three types of DoS attacks include half-open SYN attacks, ping floods, and packet storms.

Key Terms

802.11i A security extension to 802.11 and a successor to Wi-Fi Protected Access; currently the strongest security protocol for wireless networks. *See also* Wi-Fi Protected Access (WPA).

backdoor A program installed on a computer that permits access to the computer, thus bypassing the normal authentication process.

cracker Someone who attempts to compromise a network or computer for personal gain or to cause harm.

denial-of-service (DoS) attack An attempt to tie up network bandwidth or services so that network resources are rendered useless to legitimate users.

digital certificates Digital documents used in encryption and authentication protocols that identify a person or computer and can be verified by a certification authority.

encryption A technology used to make data unusable and unreadable to anyone except authorized users of the data.

firewall A hardware device or software program that inspects packets going into or out of a network or computer and then discards or forwards packets based on a set of rules.

hacker Sometimes a derogatory term to describe an unskilled or undisciplined programmer; it can also mean someone who's highly skilled with computer systems and programs and can use some of the same tools attackers use to poke around networks or systems, but not for evil purposes.

hoax virus A type of virus that's not really a virus but simply an e-mail announcement of a made-up virus. Its harm lies in people believing the announcement and forwarding the e-mail on to others.

intrusion detection system (IDS) Usually a component of a firewall, a hardware device or software that detects an attempted security breach and notifies the network administrator. An IDS can also take countermeasures to stop an attack in progress.

IP Security (IPSec) An extension to the IP protocol suite that creates an encrypted and secure conversation between two hosts.

Kerberos authentication An authentication protocol used in a Windows domain environment or on a Linux system; uses OS-generated keys, which makes this protocol more secure than having an administrator enter keys.

MAC address filtering A security method often used in wireless networks, in which only devices with MAC addresses specified by the administrator can gain access to the wireless network.

malware Any software designed to cause harm or disruption to a computer system or otherwise perform activities on a computer without the consent of the computer's owner.

penetration tester A term used to describe a security consultant who detects holes in a system's security for the purpose of correcting these vulnerabilities.

ping scanner An automated method for pinging a range of IP addresses.

port scanner Software that determines which TCP and UDP ports are available on a computer or device.

preshared key A series of letters, numbers, and special characters, much like a password, that both communicating devices use to authenticate each other's identity.

protocol analyzers Programs or devices that can capture packets traversing a network and display packet contents in a form useful to the user.

rootkits Forms of Trojan programs that can monitor traffic to and from a computer, monitor keystrokes, and capture passwords. They're among the most insidious form of malware because they can mask that the system has been compromised by altering system files and drivers required for normal computer operation. *See also* malware.

shadow passwords A secure method of storing user passwords on a Linux system.

social engineering A tactic attackers use to get users to perform an action, such as opening an infected e-mail attachment, sending a hoax virus, or providing a password, without being aware that they're aiding the attacker. *See also* hoax virus.

spam Unsolicited e-mail. The harm in spam is the loss of productivity when people receive dozens or hundreds of spam messages daily and the use of resources to receive and store spam on e-mail servers.

spoofed address A source address inserted into a packet that's not the sender's actual address.

spyware A type of malware that monitors or controls part of your computer at the expense of your privacy and the gain of some third party. *See also* malware.

stateful packet inspection (SPI) A filtering method used in a firewall, whereby packets aren't simply filtered based on packet properties but are checked for the context in which they're being transmitted. If a packet isn't part of a legitimate, ongoing data conversation, it's denied.

Trojan program A program that appears to be useful, such as a free utility, but in reality contains some type of malware. *See also* malware.

virtual private networks (VPNs) Temporary or permanent connections across a public network that use encryption technology to transmit and receive data. *See also* encryption.

virus A malicious program that spreads by replicating itself into other programs or documents; usually aims to disrupt computer or network functions by deleting and corrupting files.

wardrivers Attackers who drive around with a laptop or PDA looking for wireless LANs to access.

Wi-Fi Protected Access (WPA) A wireless security protocol that's the successor to Wired Equivalent Privacy. It has enhancements that make cracking the encryption code more difficult. *See also* Wired Equivalent Privacy (WEP).

Wired Equivalent Privacy (WEP) A wireless security protocol that encrypts data so that unauthorized people receiving wireless network signals can't interpret the data easily.

worm A self-replicating program, similar to a virus, that uses network services such as e-mail to spread to other systems. *See also* virus.

Review Questions

1. Your friend creates a shared folder on her computer for several coworkers to use. She assigns the password "00xqH}ml2-wO" to the folder. Is it an example of a good password or a bad password? Explain.
2. List at least three techniques you can use to help secure a wireless network.
3. Which of these protocols is used for VPNs? (Choose all that apply.)
 - a. PPTP
 - b. WPA
 - c. SSTP
 - d. L2TP
 - e. UDP
4. How do VPNs accomplish the "private" part of a virtual private network?
5. Which of the following terms refers to attacking a Web server by forcing it to respond to a flood of ping packets so that the server can't respond to normal traffic?
 - a. DDR
 - b. ICMP
 - c. DoS
 - d. Worm
6. Which of the following is a guideline for creating a security policy?
 - a. A security policy should be cryptic so that attackers can't understand it.
 - b. A security policy should be general enough so that rules can be added as needed.
 - c. A security policy should be enforceable.
 - d. A security policy should have different provisions depending on the user.

7. Which of the following is a component of a security policy? (Choose all that apply.)
 - a. Authentication policy
 - b. Privacy policy
 - c. Network configuration policy
 - d. Computer specification policy
8. List two questions that must be answered before determining what level of security a network requires.
9. Which of the following should be a common element in any level of security policy? (Choose all that apply.)
 - a. Complex passwords
 - b. Backup procedures
 - c. Data encryption
 - d. Virus protection
10. Choose two words from the following list that best complete this sentence: If there's access to the equipment, there's no _____.
 - a. Physical
 - b. Network
 - c. Data
 - d. Security
11. Which of the following is a requirement for rooms housing network servers?
 - a. Separate heating system
 - b. Adequate cooling
 - c. False ceilings
 - d. Shared electrical circuit
12. The procedure that specifies what resources users can access and the tasks they can perform on a network is referred to as which of the following?
 - a. Authentication
 - b. Auditing
 - c. Authorization
 - d. Logon
13. If you want to allow a blank password in a Windows XP system, which of the following do you set the password minimum length to?
 - a. blank
 - b. 0
 - c. -1
 - d. nothing

14. If you want to prevent password guessing to foil intruders, you should enable which of the following?
 - a. Account lockout
 - b. Password expiration
 - c. Password disabling
 - d. Account policies
15. Which of the following is a secure method of storing passwords on a Linux system?
 - a. PAM
 - b. The login.defs file
 - c. Shadow passwords
 - d. Reverse encryption passwords
16. Which of the following is a good password?
 - a. astronomical
 - b. FluffEE
 - c. L0sT!n@Z
 - d. BillSmithJr
17. Which of the following is a method IPsec uses to authenticate the identity of communicating devices? (Choose all that apply.)
 - a. Multishared key
 - b. Kerberos
 - c. PAM
 - d. Digital certificates
18. To encrypt data stored on a hard drive on a Windows Server 2003 computer, you should use which of the following?
 - a. EFS
 - b. DFS
 - c. NTFS permissions
 - d. gpg
19. Firewalls can filter packets based on which of the following? (Choose all that apply.)
 - a. Source address
 - b. Protocol
 - c. OS
 - d. Context
20. If network administrators want to be informed when an attempt has been made to compromise the network, what should they use?



Challenge Labs



Challenge Lab 10-1: Using Windows Firewall with Advanced Security to Limit Access to Your Computer

Time Required: 15 to 30 minutes

Objective: Limit access to your computer's Remote Desktop service to other computers in your network and allow computers outside your network to send ICMP Echo Request messages (Ping).

Required Tools/Equipment: Your classroom computer running Windows 7 or Vista

Description: This lab can be done in groups. You want to be able to allow Remote Desktop connections from users in your network but you want your firewall to block attempts from outside your network. *Note:* If you want to test this setting, you need to enable Remote Desktop connections from the System control panel. You also want users in your network and outside your network to be able to ping your computer. You need to edit existing rules in the Windows Firewall with Advanced Security console. After you have edited the rules, answer the following questions:

- Which rule did you edit to allow Remote Desktop connections to your computer?

- What did you do to limit Remote Desktop access to your network only?

- Which rule did you edit to change the ICMP Echo Request settings?

- What did you do to allow pings from outside your network?



Challenge Lab 10-2: Setting Additional Linux Password Policies by Editing the Common-Password File

Time Required: 30 minutes or longer

Objective: Edit the common-password file in Linux to set additional Linux password policies.

Required Tools/Equipment: Your classroom computer with Linux installed or a Linux Live CD

Description: This lab can be done in groups. In this lab, you set Linux password policies by editing the `/etc/pam.d/common-password` file (or on some

systems, `/etc/pam.d/system-auth`). Set the password requirements to specify that passwords must be at least eight characters and have at least one uppercase letter, one lowercase letter, one digit, and two special characters. Users have only four opportunities to pick a good password when they're changing their passwords, and new passwords must differ by at least four characters from old passwords. You'll probably need to research password settings in Linux if you aren't already familiar with the procedure.



TIP

For more information on these settings, see www.kernel.org/pub/linux/libs/pam/Linux-PAM-html/old/pam-6.html.

Case Projects



CASE PROJECTS

Case Project 10-1

DoS attacks are one of the easier attacks to perpetrate on a network, so they're often used by people who have a grudge against a company or are out to commit acts of vandalism. To read about some well-known DoS attacks, go to www.grc.com/freepopular.htm and look for the Original DDoS Report, which discusses an attack perpetrated on that Web site by a 13-year-old. Read the report and write a synopsis of how the attack was carried out and what could have been done to prevent it.

Case Project 10-2

Using the information on open ports you found with NetInfo in Hands-On Project 10-6, research these ports to determine their function and whether leaving them open is safe. A Google search is a good place to begin your research. Write a summary of what you found and list which open ports pose a security risk.

Case Project 10-3

Search for security policy templates on the Internet. A good place to start is the SANS Institute (www.sans.org). Using one or more of the templates you find, develop a security policy for your school or a business, and present it to the class. This project can also be assigned to groups of students.

Case Project 10-4

A small research company in Pittsburgh is working to develop a new method of mass storage to replace current hard drive technology. Four engineers and an office manager work there. The engineers are highly skilled professionals, and the office manager is a capable computer user. The company has an always-on Internet connection because employees must conduct research frequently. The employees have hopes of making a breakthrough and bringing the company public within the next two years. You have been hired as a security consultant to assess the company's needs. Write a paper recommending

what type of security policy should be used (open, moderately restrictive, or highly restrictive) and what security technologies should be used. On what areas should the security policy focus (physical security, data security, auditing, passwords, and so forth), and what technologies should be used to secure these areas?

Case Project 10-5

An architectural firm of eight employees, each with a networked desktop computer, wants you to develop a security policy for the company. Management has emphasized that ease of use is paramount, and little time is available for training. Working in small groups, each group should write a list of questions aimed at getting enough information for developing the policy. After determining the questions, each group should interview another group, with the other group posing as the architectural firm and answering the list of questions. What level of security should the policy reflect? Use one of the templates you found in Case Project 10-3 to develop a policy based on the answers the other group supplies.



chapter

11

Supporting a Small-Business Network

After reading this chapter and completing the exercises, you will be able to:

- Explain how to address the needs of a small business network
- Identify network equipment requirements for small businesses
- Identify requirements for small business applications
- Describe issues in supporting a small business

Once an overlooked sector of users of information technology, small businesses are spending on information technology at a rapid rate. In the United States, small businesses spent more than \$230 billion on IT in 2008 and are expected to spend \$280 billion on IT products and services in 2012. IT companies and publishers of IT books and certifications have often overlooked this large market. This chapter covers some technology issues small businesses face to give you more insight into addressing a small business's computer and networking needs.

Addressing the Needs of Small-Business Networks

What exactly is a small business? The U.S. government has, in typical fashion, multiple definitions, but a small business is often defined as one that's independently owned and operated, doesn't dominate its field of operation, and has revenues of less than \$500,000 and/or fewer than 500 employees. For the purposes of this chapter, a small business can be defined as one with fewer than 200 computers, only one or two locations, and modest technology needs. Modest technology needs have been included as a characteristic because this chapter is geared toward the entrepreneur, consultant, or small computer company that can design, install, and support a small business network without having to become an expert in more advanced computing technologies, such as minicomputers and mainframes, complex WAN environments, and so forth.

Small businesses usually have more modest requirements of networks than large corporations do. Most want to share files and printers, have a networked application that applies to their business, and networked Internet access. Most small businesses don't require a complex, highly restrictive security policy, data encryption, or advanced WAN technologies. That being said, there are plenty of exceptions. You should be aware that one size doesn't fit all, and the most important factors in being successful in supporting small businesses are listening to their requirements and designing a solution that works for them. Small business owners can be a frugal bunch, and part of the challenge a network designer or installer faces is providing a solution that gets the job done at a reasonable price.

Data and Application Sharing in a Small Business

One of the first decisions to make before determining how to set up a data-sharing scheme is whether the network should be peer-to-peer or server based. When possible and when funds allow, a server solution is almost always the best way to go, particularly if you're supporting the network after it's installed. A peer-to-peer network is fraught with problems, particularly when a user untrained in managing a networked computer is left in control of a computer that's sharing resources. On a peer-to-peer network, users can shut down their computers, unknowingly severing other employees' access to shared files or printers, which can cause data loss and corruption. In addition, a user controlling a network resource might not understand the company's security policy or how to follow it and could make sensitive data available to unauthorized users. If you're forced to use a peer-to-peer scheme, you should limit the number of computers hosting network resources to minimize potential problems.

Whether you're using a server-based or peer-to-peer scheme, the simplest file-sharing solution is usually the best. Designate as few computers as possible as file-sharing computers. A common practice is for each user to have a home directory on the server, thereby making backups easier and giving each user a place to save most of his or her files. In Windows, this arrangement can

be set up by using roaming profiles and folder redirection. Depending on the security policy, other users might have read access (but usually not write access) to each other's home directories to facilitate file sharing. If the policy is more stringent, users have access only to their own home directories, with select managers also having access as needed.

In addition to home directories, a typical practice is having one or more common folders that the entire company has access to or perhaps departmental folders shared among department members. Having common folders is a convenient way to distribute master documents without employees having to know which user maintains a document. When changes to a document are made, the document developer can post a new version in the common folder. In most OSs, permissions can be set on a single document so that the developer can allow only read access to it; in this way, a user can't change or delete the master document inadvertently. Users can copy the file to their own home directories and make changes to the copy, if necessary.

Applications can also be shared across a network. Many applications can be installed on a network file server and run from workstations via a shortcut installed on the desktop. Some applications have an installation program that creates a shortcut and sets up any Registry information the application requires, such as the location of data files. In other cases, an application allows sharing data across the network but must be installed in its entirety on each workstation running it. In either case, multiple computers having access to the same data is a big advantage compared with storing multiple sets of data or having only one computer with access to the application.

Configuring a Windows 7 HomeGroup Network Windows HomeGroups is a new peer-to-peer networking feature starting with Windows 7. A **homegroup** simplifies the process of sharing files and printers between multiple Windows 7 computers operating in a peer-to-peer network configuration. Be aware that the HomeGroups feature isn't backward-compatible with Windows Vista or XP; traditional file-sharing methods are required to share files in these OSs.

Homegroups simplify file and printer sharing by making it unnecessary to create user accounts on every computer that shares resources. In addition, configuring permissions and finding files on the network have been reduced to child's play. Homegroups are usually suitable only for small office/home office (SOHO) networks with fewer than 10 computers that operate under a fairly open security policy. For networks where homegroups are suitable, this feature might be just what small business owners have been looking for to simplify networks and increase productivity.

Creating and Joining a Homegroup The basic requirement for creating a homegroup is Windows 7 Home Premium, Ultimate, Professional, or Enterprise. Windows 7 Home Basic and Starter editions can join but not create a homegroup. Next, the network must be designated as a Home network rather than Work or Public. If you choose Home as your network location, you're prompted to create a homegroup. You can choose what type of files you want to share and whether you want to share printers. If your network is already designated as a Home network, you can create or join a homegroup by going to Control Panel and clicking "Choose homegroup and sharing options" under Network and Internet.

Homegroups are password-protected by a password that's generated randomly when the homegroup is created (see Figure 11-1). Users must enter this password when they join a homegroup. To view or change the password, open HomeGroup in Control Panel.



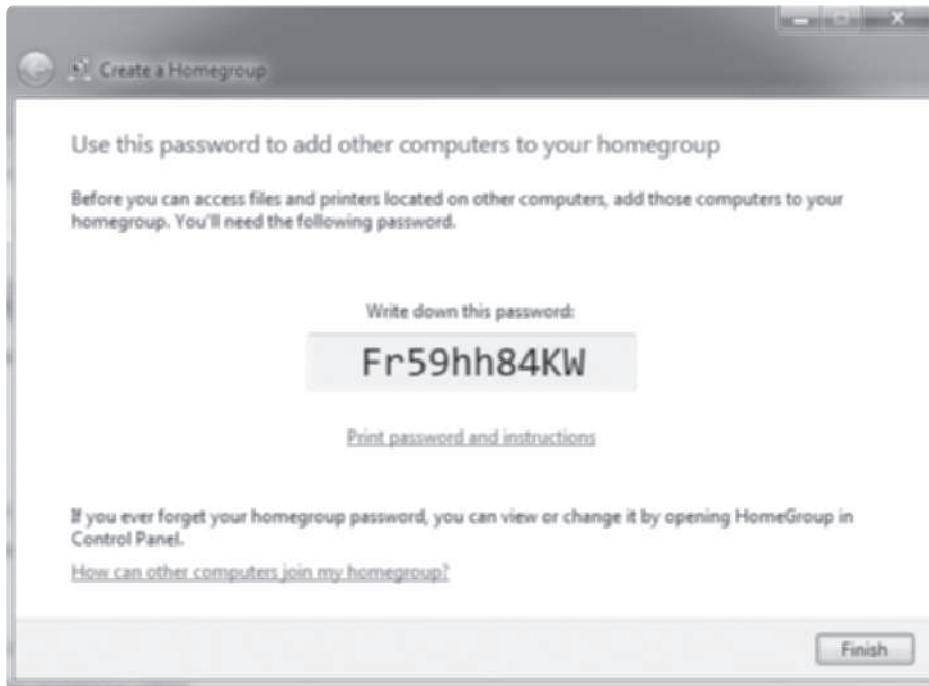


Figure 11-1 A homegroup password is generated randomly

Courtesy of Course Technology/Cengage Learning

After a homegroup is created, other computers can join it by changing their network locations to Home or by going to HomeGroup in Control Panel. When a computer joins a homegroup, the same sharing options configured when the homegroup was created are available, allowing different computers to have different sharing options. After sharing options have been chosen, the homegroup password must be entered. When a computer joins a homegroup, all users logging on to the computer can access shared resources on the homegroup, but each user controls access to the files he or she shares. Users access shared files by opening the Computer window and clicking Homegroup in the left pane to display a list of other users currently logged on with their computer names in parentheses (see Figure 11-2).

Users can share or unshare a folder or library by right-clicking it in Windows Explorer and pointing to Share with (see Figure 11-3). A folder or library can be unshared (by selecting the option to share with Nobody), shared with the homegroup for read or read/write access, or shared with specific users whose computers may nor may not be members of the homegroup.

To access files shared by another user, click the user to see a list of available shared libraries and folders. If you access another user's shared library frequently, you can add it to your library, thereby making the files accessible directly from your Libraries folder in Windows Explorer.

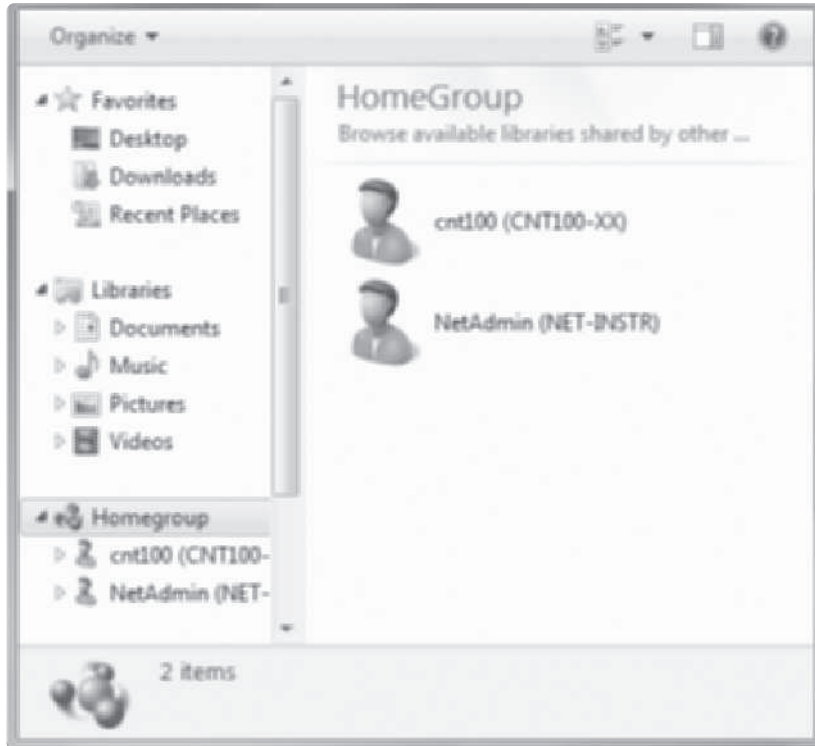


Figure 11-2 Viewing members of a homegroup

Courtesy of Course Technology/Cengage Learning

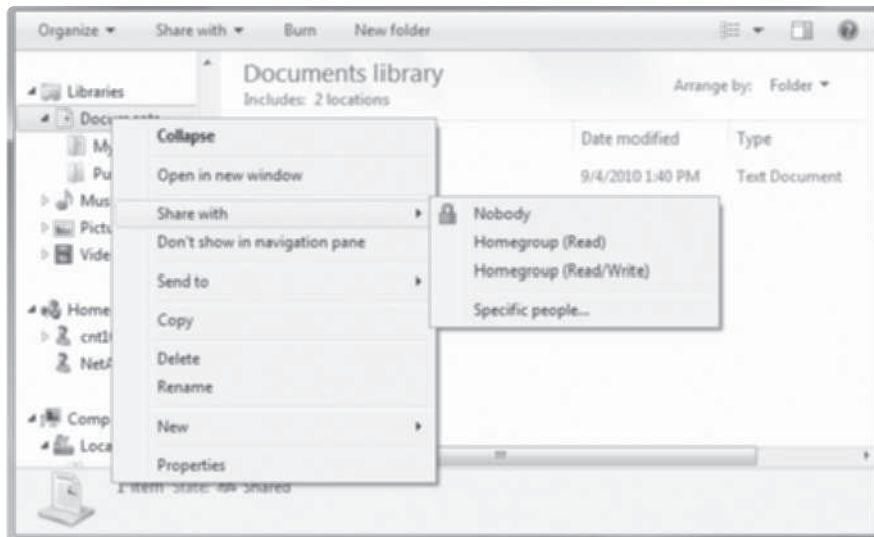


Figure 11-3 Changing homegroup sharing options

Courtesy of Course Technology/Cengage Learning

Troubleshooting HomeGroups At times, you might not see all computers that are members of a homegroup, or you might have problems creating homegroups. In these situations, you can use the HomeGroup Troubleshooter to solve some of these problems. To open it, right-click Homegroup in Windows Explorer. HomeGroup Troubleshooter attempts to verify the firewall and permission settings required for homegroups to work and tries to detect network problems that might prevent homegroups from working correctly. If it doesn't solve the problem, check the following:

- *Routers*—Homegroups don't operate across routers, so all computers must be on the same subnet.
- *IPv6*—IPv6 must be enabled to create or join homegroups.
- *Clock settings*—All computers must have the same time settings, including the correct time zone.
- *Third-party firewalls*—HomeGroups sets Windows Firewall when you create or join a homegroup, but if you're running a third-party firewall or antivirus program that performs firewall functions, you might have to disable or configure it. Microsoft publishes a document describing how to set firewalls to allow HomeGroup functionality.
- *Network Discovery*—Network Discovery must be enabled. To access this setting, click Advanced sharing settings in the Network and Sharing Center.

Homegroups can simplify file and printer sharing when your security needs are modest and you want files to be shared between several computers in a peer-to-peer network. For more advanced sharing options in a peer-to-peer network, use traditional file sharing and permissions, discussed in Chapter 9. For an all-around better file-sharing solution with centralized account management, use a domain-based network.



Hands-On Project 11-1: Sharing Files with Windows HomeGroup

Time Required: 20 minutes

Objective: Configure Windows 7 computers to join a homegroup.

Required Tools/Equipment: Your classroom computer running Windows 7

Description: In this project, you configure Windows 7 to use the HomeGroup file-sharing feature. The instructor should lead this project by creating a homegroup first (or designate one student to create the homegroup), using the steps outlined in the previous section. You then join the homegroup by following the steps in this project. You need the homegroup password.



The homegroup creator can change the randomly generated password to make it simpler.

1. Log on to your computer as an administrator.
2. Click **Start, Documents**. Create a text document in your Documents folder named **DocOnNet-XX** (replacing **XX** with your student number).

3. Open the Network and Sharing Center. Under View your active networks, click the **Work network** link. (You set your network type to Work network in Hands-On Project 5-1; if it's set to Public network, click this option instead.)
4. In the Set Network Location window, click **Home network**. If the homegroup hasn't been created yet, the next window prompts you to create a new homegroup. Your homegroup should already have been created, so click **Cancel**, wait until the instructor creates the homegroup, and start this step again.
5. In the Join a Homegroup window, click the **Documents** check box so that your documents are also shared (see Figure 11-4). Click **Next**.



Figure 11-4 Joining a homegroup

Courtesy of Course Technology/Cengage Learning

6. In the next window, enter the password your instructor has supplied for the homegroup, and then click **Next**.
7. The next window states that you have joined the homegroup. Click **Finish**.
8. To view other computers in your homegroup, click **Start, Computer**. In the left pane, a list of users in the homegroup, with their computer names in parentheses, is shown under Homegroup, similar to Figure 11-2 shown previously.
9. Click a user listed under Homegroup to see the list of shared folders. (If you don't see a list of users, right-click **Homegroup** and click **Start the HomeGroup troubleshooter**.)
10. Click **Documents** and verify that you can see the document created in Step 2.
11. Change your network location back to **Work network**. Close all open windows.

Sharing Files in a Windows Domain Environment If you require more than the minimum security and sharing files is a major part of your business network, using a file server with a centralized user database is the best way to go. In Windows, these requirements mean installing a domain controller. After installing a Windows server configured as a domain controller, user accounts need to be created only on the domain controller. All desktop computers and other servers simply need to be made domain members. After user accounts are created, they, along with group accounts, can be used to assign file and folder permissions on any computer in the domain. You can share folders on the domain controller, other servers, and even desktop computers, and permissions can be set by using accounts created on the domain controller. This centralization of accounts vastly simplifies resource management and improves network security.

Sharing Files in a Linux Environment Typically, you have two choices for sharing files in a Linux environment. One is using Samba, discussed in Chapter 9, for compatibility with a Windows environment. The other choice is using Network File System (NFS), which might be a good choice if the environment consists of mostly Linux computers. In both cases, the Linux client computer mounts a shared folder in its own file system and accesses the shared folder as though it were a local resource. Permissions are set in the Linux file system, as discussed in Chapter 9. The advantage of using Samba rather than NFS is Samba's compatibility with Windows file sharing.

Using an NAS to Share Files As discussed in Chapter 8, an NAS is a server dedicated to sharing files. A device sold as an NAS has its own user interface, usually accessed via a Web browser for creating user accounts, groups, and shared folders. Many NASs can also be configured to integrate with a Windows domain controller, so user accounts need to be created only on the domain controller. An NAS can be a good compromise between a peer-to-peer network and a domain-based network. NAS configuration is usually straightforward and doesn't normally require the expertise that Active Directory in Windows Server 2008 does. In addition, an NAS can cost much less than a server with Windows Server 2008 installed. Some NASs have a few slots for installing hard drives and a network interface and are no larger than a thick book. So if space is at a premium and simplicity in configuration is preferred, an NAS might be the ideal solution.

Equipment Sharing in a Small Business

A printer is the most common piece of equipment shared in a network. A typical issue in small businesses is sharing personal printers attached directly to a user computer's USB port. Sharing printers in this manner is challenging because the user has control over this printer's operation. The user could shut down the computer, turn off the printer, or take some other action that prevents network users from printing to the shared printer. Nonetheless, printer sharing is an important requirement of most small business networks.

One way to facilitate printer sharing is to connect the printer to the network rather than to a user's desktop. Some printers come equipped with a network interface or a slot to plug in a network interface, thereby allowing you to connect the printer directly to the network. If this option isn't available, some companies make small print server boxes that plug into the network on one end and plug into the printer via a USB port. Whether the network interface is built into the printer or is an add-on device, these print servers can be assigned an IP address and accessed by most OSs.

Scanners can also be shared. Scanners that can be shared on a network come with their own sharing software and can't be shared by using the same method printers use. Hewlett-Packard (HP), for example, has a utility that runs on the computer to which the scanner is attached. This utility shares the scanner and allows you to specify a password, if needed. Other computers must have the HP scanning software installed and run the remote scanning software supplied by HP. Because high-end scanners can be expensive and take up a lot of desktop space, sharing them among several users makes sense.

Other devices that can be shared on a network include external hard drives that connect via a USB interface and card readers that read Secure Digital (SD) and Compact Flash (CF) cards, such as those in digital cameras and PDAs. These external hard drives are a good solution for backing up data in lieu of tape or removable media, such as CDs or DVDs. Because hard drives have so much capacity, they can be set to back up data files of many users over the network quickly and easily.

Equipping Small-Business Networks

Most TV and print advertisements for network equipment are aimed at large enterprise network administrators. The equipment needs of most small businesses are far more modest. A rack full of blade servers is overkill for most small businesses, unless their business is Web hosting. A typical small-business environment might consist of one or two servers, some workstations, a few switches, and a router to connect to the Internet.

However, don't be enticed by low prices and the seemingly equivalent functionality of network devices sold as consumer products. Consumer products are intended for residential users and might lack important security or management features found in more robust products targeted at small-businesses. Most well-known equipment manufacturers have a small business or small and medium business (SMB) line of products with features that are often worth the extra cost.

Servers and Desktops

Most computer manufacturers, such as Dell and Hewlett-Packard, have small-business solution centers with products focused on the needs of small businesses. Go to any of these companies' Web sites, and you'll see a link to their small-business offerings. Usually, you can purchase a server fully loaded with a small-business edition of an OS and features you can choose, such as e-mail, Web, and database servers. Many companies give you an option of which OS you want preinstalled, or you can install your own. Common choices are Windows Server 2008 Standard Edition, Windows Small Business Server 2008, and Linux. Several server manufacturers offer a buying guide listing features and servers supporting these features.

A general rule of thumb when purchasing a server for a network is to buy as much hardware as the budget allows that will meet the company's estimated needs for the next two to three years. Buying hardware with expandability features that you can't foresee using in more than two to three years makes little sense because by that time, upgrading to a new computer might make more sense than upgrading the hardware on an existing one—assuming you *can* upgrade. An example of buying too much expandability is purchasing a server with one CPU that can be upgraded to four CPUs. If one CPU meets your needs today, it's unlikely you'll need four CPUs



in a year or two. In addition, because CPU technologies change so quickly, acquiring the additional CPUs you need to upgrade might be difficult. Besides, in two to three years, a single CPU will probably be able to do as much work as four older CPUs for less money. That being said, you should consider being able to upgrade to a faster CPU or add a processor because these needs are likely to come up within a year after your initial purchase.

Memory expansion and storage expansion are critical design features to look for in a server. You might think a server with 2 GB of RAM and 500 GB of hard drive space is enough, but when this simple file and print server turns into a database server and Web server, the OS might be starving for resources.

Another feature to look for is fault-tolerant storage solutions. That usually means a RAID disk system, which makes it possible for the server to continue operating even if a disk drive fails. A common disk configuration is using RAID 1 (disk mirroring) on the drives containing the OS and applications and RAID 5 (disk striping with parity) on data drives. Disk mirroring requires two disk drives of the same size because everything written to one disk is written to the second disk automatically. If one disk fails, the other disk has a complete up-to-date copy of the system, and the server can continue running as though nothing happened. Disk striping with parity requires at least three disks. When data is written to a RAID 5 disk system, it's spread evenly over two of the disks, and parity information is written to the third disk. With this arrangement, if a disk fails, the data on the failed disk can be reconstructed from the data and parity information on the remaining disks.

Desktop computers for a small business usually differ from a computer designed for home use in the software installed and some hardware components. In most home computers, multimedia and entertainment components and software are emphasized, but in most business computers, the focus is on productivity software and manageability. For example, many home computers have a home edition of the OS installed, whereas a business computer is better off with a professional, ultimate, or enterprise edition. Computers running one of these editions can be part of a Windows domain, and this OS offers more management and security features than home editions do.

Networking Equipment

Another decision to make before you select networking equipment for a small business is where to put the equipment. Most small businesses don't have a large wiring closet, so you might need to get creative. For example, in a business consisting of only eight peer-to-peer computers, an existing cabinet can be used with an eight-port switch bolted to the cabinet wall. Care has to be taken to ensure adequate ventilation, but a small switch doesn't require much. In a small business, the space used for network equipment is often shared with another function, such as phone and alarm system equipment. Common sense must be used to make sure the space is adequate to the job. Servers and some network switches can generate a lot of heat, so cooling is essential. If the system gets too warm, hard disks and motherboard components can fail or become corrupted. As with servers and desktops, look for the SMB line of products from network equipment manufacturers, such as NetGear, Linksys, and D-Link.



Linksys is owned by Cisco Systems. You can find Linksys products at www.linksys.com or <http://home.cisco.com/en-us/wireless/linksys/>.

Making a Wired Connection In a small network of only a few computers, simply running cable from computers to the hub or switch might be tempting, but don't give in to this temptation. The biggest problem with running a single cable from station to switch is that when you move the computer, you might not have enough cable slack to reach the new location. Even if you have only a few computers to connect, you should have network jacks at the work area wired to a patch panel in the wiring closet near the switch. Then make the connection from the jack to the computer's NIC and from the patch panel to the switch with patch cables. (This cabling arrangement was shown previously in Figure 4-9.) Category 5e or 6 cable should be used, and after it's installed, it should be tested. If you're working with an existing cable plant, test all cable runs before you begin your work, and replace or reterminate any cables that fail or have suspect terminations.

Switches should be used to connect workstations; if you're upgrading a network, replace hubs with switches. When choosing a switch, you should consider the following options:

- *Switch speed*—100 Mbps and Gigabit Ethernet switches are the norm today. In a server-based environment, an asymmetrical switch is recommended, with most ports being 100 Mbps ports and one or two being Gigabit Ethernet ports. Servers should be attached to the Gigabit Ethernet ports.
- *Managed or unmanaged*—A managed switch has several advanced configuration options and can sometimes gather network data on a per-port basis. The extra functionality comes at a price—usually 5 to 10 times that of an unmanaged switch. In most cases, an unmanaged switch is adequate. A smart switch (discussed in Chapter 7) might be a good compromise if you need some of the advanced features it offers. A 24-port 100 Mbps switch with two Gigabit Ethernet ports can cost from \$100 up to \$400 or more.
- *Support for multiple media types*—Higher-end switches might have provisions for both copper and fiber-optic connections. In some cases, the fiber-optic connections come as an optional plug-in module, called a gigabit interface converter (GBIC) or mini-GBIC. This type of switch is ideal when you have to connect to computers or another switch exceeding the distance limitations of UTP cable or if the cable must pass through an electrically noisy environment, such as a manufacturing floor.

As you plan and install wiring for the network, be sure to keep in mind the company's security policy. Although many small businesses don't have a defined security policy, physical security issues should be discussed before equipment is installed so that the business owner and you can make informed decisions.

Making a Wireless Connection The availability of wireless equipment at a good price with good performance makes going unwired an attractive option, especially for new installations with a fairly small number of computers. Vendors don't usually provide information on the maximum wireless connections an access point (AP) can handle, but independent testing has shown that most consumer products (such as an AP you find at an office supply store) max out at around 40 connected computers. More expensive commercial products, such as those manufactured by Cisco Systems, can likely handle more than 100 connections with little data loss. These restrictions don't mean that you can't use wireless in a larger network; they simply mean that multiple APs might be needed, especially if computers are spread over a large area.

That being said, wired connections are recommended when the environment accommodates running wires and the computers are stationary. If the environment consists mainly

of users running around with laptops or handheld computers, a wireless infrastructure is definitely an advantage. Of course, nothing prevents you from using a combination of wired and wireless networking. In fact, a combination is the most common design. A wired infrastructure is used for all desktop computers and servers, and a wireless AP is set up for mobile users. The wireless AP connects to the wired network so that the two networks can communicate, as shown in Figure 11-5.




Figure 11-5 Wired and wireless networks coming together

Courtesy of Course Technology/Cengage Learning

Working with a wireless network involves special considerations that aren't a factor in wired networks. For example, outside (and unauthorized) people can access the wireless network simply by being close enough to the AP to receive a signal. This security concern is why using the wireless security practices discussed in Chapter 10 is imperative. Using a form of encryption is critical when installing a wireless network for a business; otherwise, data is at risk. Another issue with wireless networks is unique forms of interference. For example, cordless phones can use the same frequency as an 802.11 network. If a phone is in use, users can lose their connections with the network. This problem might not be too severe in a home network (although it's annoying), but it can grind business to a halt. When designing a wireless network, identify any sources of potential interference, such as

cordless phones, microwave ovens, nearby wireless networks, and other radio frequency sources, and be sure they're on frequencies different from the wireless standard you plan to use. Also, make sure you test during normal business operations rather than after hours because equipment might be in use during business hours that you would be unaware of after hours.

When selecting wireless equipment, one of the most important considerations is the security standard supported. Whatever security standard (WPA or WPA2, for example) the AP supports must also be supported by the wireless NICs you use. Another feature to consider is whether the AP can be bridged to other access points. If your wireless network extends beyond the reach of one AP, multiple access points might be required, and you need to make sure all access points can communicate with one another.

Communicating with the Outside World

Even small businesses need to communicate with the outside world, particularly to access the Internet. Besides Internet connections, many businesses need employees to be able to access the company network from home or while away on business. Some issues discussed in this section include Internet access, dial-up connections, and virtual private network (VPN) connections.

Accessing the Internet Most small businesses require Internet access. For 10 to 20 computers, a broadband cable or DSL connection is usually enough. A typical setup includes a cable or DSL modem connected to a router and the router connected to one of the switches, as shown in Figure 11-6.




Figure 11-6 A network with a cable modem

Courtesy of Course Technology/Cengage Learning

In this type of network setup, the router is usually the network firewall, too. However, depending on your security needs, investing in a dedicated firewall that sits between the router and the switch might be worthwhile. In this configuration, the firewall stops any attacks before they reach the router. The router also usually plays the role of DHCP server and network address translator, handing out private IP addresses to internal computers and translating them via Port Address Translation (PAT) to the address the ISP provides when Internet requests are made. The router also gives each workstation a default gateway address (the router's address) and the DNS server addresses workstations need to translate domain names to IP addresses. The router usually acquires DNS server addresses from the ISP.

The type of router used in this situation is typically an inexpensive device (less than \$100) with a Web browser interface for configuration. If a wireless network is set up, a router supporting both wired and wireless connections can be used. These devices can be purchased from computer retailers and office supply stores and are often billed as SOHO routers.

The default operation of these devices is to allow all packets through if an internal computer initiates the communication session. However, no unsolicited packets are allowed into the network from the outside. If communication needs to originate from the outside (for example, the business is running its own Web server), the router configuration can be changed to allow this communication. Figure 11-7 shows the virtual server configuration window for allowing this type of communication. The configuration option is sometimes called **port forwarding** because you're configuring the router to forward requests that apply to only certain TCP ports. In this figure, port 80 (the port for HTTP or Web communications) is being forwarded to an internal Web server at address 192.168.2.10.

Router Setup Utility Home | Help | Logout Internet Status: Connected

Firewall > Virtual servers

This function will allow you to route external (Internet) calls for services such as a web server (port 80), FTP server (Port 21), or other applications through your Router to your internal network.
[More info](#)

Clear Changes Apply Changes

Add Active Worlds Add

Clear entry 1 Clear

	Enable	Description	Inbound port	Type	Private IP address	Private port
1.	<input checked="" type="checkbox"/>	Web server	80 - 80	TCP	192.168.2.10	80 - 80
2.	<input type="checkbox"/>		-	BOTH	192.168.2.	-
3.	<input type="checkbox"/>		-	BOTH	192.168.2.	-

Figure 11-7 Configuring port forwarding on a router

Courtesy of Course Technology/Cengage Learning

Of course, the company security policy should be taken into account when configuring access to the outside world because this type of connection is where most trouble originates. As soon as an Internet connection is established, equipping all servers and workstations with antivirus and antispyware software is doubly important.

VPN Remote Access A VPN remote access connection can be made as long as both parties are connected to the Internet. As discussed in Chapter 10, a VPN creates a private communication channel between two parties via the public Internet. Two VPN modes are available with most VPN devices:

- *Gateway-to-gateway*—In the **gateway-to-gateway VPN mode**, a VPN connection is established between two routers that support VPNs. No software needs to be installed on the computers using the VPN. This mode is used mostly between offices connected to the Internet through a router that supports VPNs. In this setup, all communication between the two offices is private, even though data travels across the public Internet.
- *Client-to-gateway*—The **client-to-gateway VPN mode** establishes a VPN connection between a single client computer and a VPN device. This mode requires configuring a VPN client on each computer participating in the VPN and a VPN device that clients connect to. This mode is best for providing private communication to the company network for employees working from home or employees who must connect to the network while traveling. In this setup, users connect to the Internet through their ISPs, and then run the VPN client software to create a private connection with their company network.

Many SOHO equipment manufacturers, such as Linksys, NetGear, and D-Link, have fairly inexpensive VPN routers that support either VPN mode. Prices for this equipment range from under \$100 for a VPN router that supports eight or fewer VPN connections to under \$500 for a VPN router that supports as many as 30 or 40 connections. Be aware of terminology when purchasing a router for VPN connections. Some routers claim to support VPN “pass-through,” which enables a VPN client to connect to a remote VPN device but doesn’t actually create a VPN connection. This type of router is best for users who have a small network at home and want to connect to the company network by using client-to-gateway VPN mode. The home router doesn’t participate in the VPN connection; it simply allows the VPN connection to pass through it.

When outfitting your business with a router that supports VPN connections from remote clients, look for one that supports VPN endpoints. The number of endpoints or tunnels the VPN router supports tells you how many VPN connections can be established. In client-to-gateway VPN mode, one VPN endpoint per user connection is required. In gateway-to-gateway VPN mode, in which a connection is made from LAN to LAN, one endpoint per LAN connection is required.

Identifying Requirements for Small-Business Applications

The application needs of small businesses range from ho-hum simple to quirky and complex. On the simple side, some businesses need only an office application suite, such as Microsoft Office. On the complex side, a business might use a custom program that has little or no support for networking and requires the network administrator to be creative in making it work

with the office network. This chapter doesn't delve into industry-specific applications; instead, it concentrates on some standard business applications that have a place in many small businesses.

Before discussing specific types of software, two issues for network applications should be addressed. The first is that not all software is designed to operate over a network with multiple users accessing the data. Many applications, such as account and sales management software, sell both single-user and multiuser versions of their programs. It might be possible to network a single-user version, but you're usually limited to one user at a time accessing the data. If you know that multiple users need to access the application simultaneously, you probably need to purchase the multiuser version. The second issue is software licensing. Just because an application has been purchased doesn't necessarily mean it can be installed on the network or on multiple computers. The **end user license agreement (EULA)** should be consulted before purchasing any software for use by multiple users.

Accounting Software

Many accounting or bookkeeping applications are tailored to small businesses, such as QuickBooks by Intuit, AccountEdge, Sage Software's Simply Accounting, and Peachtree Accounting. A few free products are worth a look, such as Freshbooks and Outright. Most of these applications offer a choice of version, depending on the complexity of the company's needs. Typically, accounting packages come in basic versions to support common business needs, such as invoicing, check writing, inventory tracking, and payroll. For more complex needs, many packages have a professional or advanced version that supports multiple users, bill-of-materials handling, time and billing management, and other advanced features. As a network technician or administrator, your job isn't to support the function of these applications but to help ensure that the network is set up to run them adequately and make sure data is backed up. Some issues a network technician faces in supporting these applications include the following:

- Should the software be accessible by multiple users in the company? If so, should users be able to access the application and make changes to data simultaneously?
- How should the application be secured from users who shouldn't have access?
- Older software packages might not integrate with Windows well or support networking directly. You might need to use mapped drives to specify network locations or use printer redirection to print to network printers.
- How is the software and its data backed up? Some applications have their own backup program for data files. Whether you use a third-party backup program or the backup program included with the application, you must devise a scheme to support backing up and restoring data.

These issues are just a few of the many to consider in supporting accounting software. Because accounting is of utmost importance to a business's operation, care must be taken before performing any actions that might corrupt or destroy data. When possible, consult the application vendor's technical support when configuring or troubleshooting these applications.

Sales and Contact Management Software

Although many small business owners are aware of accounting software, they might not be familiar with sales and contact management software. Today's sales and contact management

software offers features that are leaps and bounds ahead of yesterday's rolodex. Notes on customer conversations can be tracked and accessed by multiple users, promotional mailings can be targeted and automated, and even customers' birthdays and their pets' names can be tracked.

Products such as Goldmine from FrontRange Solutions, Maximizer by Maximizer Software, and ACT! by Sage Software have been industry standards in this area for years. These packages fall into the category of **customer relationship management (CRM)** and go far beyond simple electronic phone books. Shared calendars and to-do lists, sales forecasting, extensive client history management, and integration with smartphones are just a few features these packages offer. Many have versions targeted for specific industries, such as real estate and financial services.

CRM software has some of the same support issues you might encounter with accounting software. Again, although you might not want or even need to be an expert in this type of software, a network technician working with small businesses is often expected to be a jack-of-all-trades in computer and network support, including recommending and supporting small-business applications. Your knowledge of the products available for handling typical small-business tasks will help on the job, not to mention getting the job in the first place.

Windows Small Business Server

Although Windows Small Business Server (SBS) isn't a small-business application, it's bundled with many of the network services and add-on products a small business might need. SBS was designed with small businesses in mind and includes an administrative console with all main configuration options in one central user interface (see Figure 11-8).

SBS offers the following application features beyond the Windows Server 2008 OS:

- *E-mail server*—Microsoft Exchange Server 2007 comes standard with SBS 2008.
- *Shared calendar*—It includes Microsoft Exchange Server's calendaring functions.
- *Intranet resource sharing*—Sharepoint Services is available to organize documents and collaborate with colleagues.
- *E-mail security*—Microsoft Forefront Security is included to protect e-mail from viruses, worms, and spam.
- *Database*—Microsoft SQL Server comes with SBS 2008 Premium Edition.

SBS provides support for up to 75 users and is designed to be run by employees who don't have an extensive IT background. For larger networks, Windows Essential Business Server supports up to 300 users and has most of the same features as SBS plus security and network management options. However, networks of that size and complexity usually require onsite IT staff. Windows SBS is a good alternative to standard Windows Server editions if you know your business needs its built-in services and you want to be able to manage most services by using in-house employees instead of hiring IT consultants.

Hosted Applications

An alternative to installing and supporting business applications on servers and desktops is using hosted applications, which are one component of cloud computing (discussed more in Chapter 12). Google Apps, for example, provides a full suite of office applications, e-mail,






Figure 11-8 The Small Business Server 2008 administrative console

Courtesy of Course Technology/Cengage Learning

document collaboration, and calendaring that can be accessed through a Web interface. The Standard Edition is free and suitable for fewer than 50 users in environments where storage and security needs are modest. The Premier Edition has a nominal yearly fee for unlimited accounts, secure access, support, and considerable storage space. Using hosted applications is a viable alternative for businesses with fast, reliable Internet connections that don't need the advanced features of office suites, such as Microsoft Office, and don't want to support them.

Is Linux a Viable Desktop Alternative to Windows?

Much has been written about Linux in the workplace, with most of it focused on using Linux as a server OS. However, Linux has made inroads into desktops, too. Using an open-source Linux distribution is certainly less expensive than Windows or Mac OS in software costs. The biggest questions about using a Linux desktop are **total cost of ownership (TCO)** and application support. TCO is the cost when you factor in intangibles, such as support costs and productivity gain or loss. Linux has come a long way since the early 1990s, when only the savviest computer technician was willing to take on its difficult installation and configuration. Today, Linux installation for most distributions is no more difficult than a Windows installation. Although postinstallation configuration can still be challenging, Linux is definitely a viable option as a base OS.

In application support, Linux-based Web browsers and e-mail clients are similar to those in Windows and Mac OS, and powerful (and often free) office application suites are available. OpenOffice.org (www.openoffice.org) is an open-source office application suite that runs on Windows, Linux, Sun Solaris, and Mac OS. It's available as a free download and consists of a word-processing program, a spreadsheet program, presentation software, a database application, and a drawing and diagramming application. OpenOffice has good compatibility with Microsoft files. KOffice (www.KOffice.org) includes a full-featured word-processing program, a spreadsheet program, presentation software, a database application, and many other add-on applications. It runs on Linux and some UNIX versions and offers some Microsoft Office compatibility. A quick look at the Ubuntu Linux desktop and Applications folder (see Figure 11-9) shows that Linux is making good progress toward being your next desktop OS.

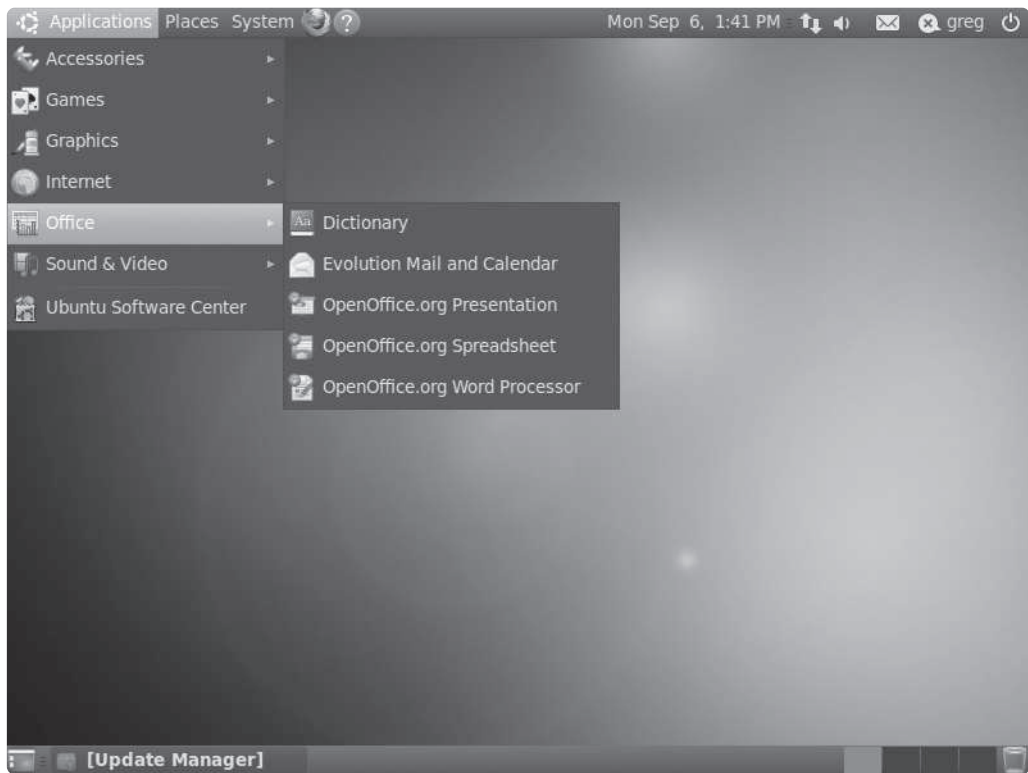


Figure 11-9 The Applications folder in Ubuntu Linux

Courtesy of Course Technology/Cengage Learning

Although Linux has a number of office productivity packages, support for industry-specific applications is lacking. If a small business simply needs what an office productivity suite includes, Linux is a viable option, particularly if support is available from the company's network technician. However, if the company needs to run industry-specific applications, you must find out whether Linux is supported. If a business is sold on Linux because of its open-source licensing and improved security over some competitors, there are solutions for running Windows applications in Linux. Windows Emulation (WINE) is one

that's available free for all Linux distributions when a Windows environment is simulated, allowing an application to run as though it were running on Windows. VMware and VirtualBox, discussed in Chapter 8, are virtualization environments that enable you to run Windows on a Linux OS and vice versa. So if Linux sounds attractive but a critical application requires Windows, you can have the best of both worlds by using virtualization. Linux has come a long way, but its suitability for desktops depends on the expertise of the technicians who set it up and the compatibility of applications the business needs to run.

Supporting a Small Business

The job of supporting small businesses in their IT needs can sometimes be more difficult than supporting a large business because small businesses rarely have in-house specialized expertise. Large businesses might have a technician who specializes in supporting the network infrastructure, another who supports servers, one who supports desktop OSs, and still another who supports specialized applications; small businesses usually count on their hired consultant to do it all. This is where you, the entrepreneur, come in.

Entrepreneurs Wanted

Although it's true that many equipment manufacturers have started catering to small businesses, many computer consulting companies still have a large-business mentality. They're used to large jobs with large budgets, and their job is to install the network and leave the rest to the internal staff. Small businesses don't have large budgets or internal staff to support a network after it's installed. If you decide you want to specialize in working with small businesses, you need to understand their needs and be able to explain what technology can do for them. That's why knowing the choices available to these businesses for accounting software, CRM software, and office suites is important. You also have to understand and respect the way companies do business yet be able to gently nudge them toward solutions you know will make sense for them when they understand their options. Working with small businesses can be financially rewarding and create a sense of achievement because you can make a difference in a company's success by helping it use technology to increase and maintain its business. Before you can start working with small businesses, however, you must convince small-business owners to place their information technology in your hands.

Getting the Job Most small-business owners who are looking for a computer or network consultant usually request proposals from multiple vendors. If you're called on to develop a proposal, the most important thing you can do is listen to the company's requirements. Find out what kind of business a company is in, how it's currently managing day-to-day information, and where it wants to be in the next several years. Talk not only to the owner or manager, but also to the people actually using the computers, working with customers, handling accounting, and so forth. The more you know about how the business works, the more tailored and detailed you can make your proposal.

Most small businesses are customer friendly, and they expect the same from their vendors. Typically, they delve into technology cautiously because they don't have internal expertise and are hesitant to place the future of their business in a consultant's hands. So above all, be responsive when customers or potential customers call you. They need to know they can

count on you to be available when they need you. Large businesses often don't expect to talk to a live person when they call for support because they often use call automation. If a small-business owner gets your voicemail or an automated answering service, you might lose that company as a customer.

When developing a proposal for a consulting job, be detailed about what's included in the price you quote. Many businesses want to have choices, so give them multiple quotes when appropriate, and spell out the advantages of the higher-priced option. Each line item of a proposal should specify what need it fills. That way, owners don't have to wonder whether they really need what you're quoting and whether all their needs are addressed adequately.

Securing a Small-Business Network

One aspect of your proposal to a small business that you shouldn't neglect is security. Spell out in your proposal how you plan to secure the network and data. First, you need to determine what type of security will work best with the business: an open security policy, for example, or a highly restrictive policy. When discussing a business's needs, be sure to emphasize the trade-offs between an open policy and a more secure policy. A more secure policy safeguards data better but at the expense of requiring more user training and perhaps lowering productivity. You must factor in how this company currently does business and whether tight security is truly a requirement for its business.

Passwords and Backup Don't automatically assume that every business should have password policies that require frequent changes and complex passwords—or any passwords, for that matter. Some businesses simply don't need or want this level of security. Perhaps all that's required is antivirus and antispyware software, an easy-to-follow backup scheme, and a simple disaster recovery plan. It's the consultant's job to make technology work for a business, not against it. As long as you explain the ramifications of an open security policy, and it's what the business wants, you just need to carry it out. If, on the other hand, a business does want a secure network, it's your job to know how to construct it. Chapter 10 covers many security issues you should be aware of and tools for setting up security measures.

Regardless of the security policy, one of the first security-related items on your agenda should be an easy-to-use backup strategy. Every business needs a backup scheme, and unless you're going to be available every night to run a backup, the process should be clear and concise so that any of the business's users can do it. Tape backup is still a favorite method for backing up large amounts of data, especially when hundreds or thousands of gigabytes must be backed up regularly. However, in a small business, tape backup might be unnecessary and unnecessarily complicated. Depending on the amount of data to back up, using removable media, such as DVD-RWs, might make sense; if there's more data, USB hard drives are a good solution. Some USB hard drives have a one-touch backup solution. For data on users' computers, document files can be backed up to a network hard drive, which can be backed up periodically to removable media for offline storage, if needed. The OS and applications don't usually require regular backup. One convenient method of backing up a user's OS and applications is creating a drive image of the computer's hard drive periodically and backing up the image to a network location. Many software packages are available to create an image of a hard drive, such as Symantec Ghost and Acronis True Image. A drive image allows you to recover from drive failure simply by restoring the image to a new drive. The user's computer is then returned to the same state it was in when the image was last recorded.





Microsoft's Backup utility in Windows Vista, Server 2008, and later can create a drive image, as can a free utility called Disk2VHD, available from Microsoft's Web site at <http://technet.microsoft.com/en-us/sysinternals/default.aspx>.

However you decide to back up, one essential step that many IT technicians skip is testing the backup and restore process to ensure that the backup scheme is actually working and data can be restored when disaster strikes.

Security from the Outside World Antivirus/antispyware software, as mentioned in Chapter 10, is a must for any computer with an Internet connection. Beyond that, a firewall should be in place for most businesses that share a connection to the Internet, such as through a cable or DSL modem. If a Windows computer is used to share an Internet connection, Windows Firewall is in place, but a dedicated router is preferable because it offloads extra traffic from a workstation and usually has more firewall features.

If a router is used to provide Internet access to multiple computers, it should be equipped with a firewall. Most inexpensive commercial routers are designed to block incoming traffic unless it's part of an existing conversation with an internal computer. However, for more complete protection, opt for a router described as a firewall router. These routers have firewall features that protect a network from external threats, such as DoS attacks and IP spoofing. They can also be set up to filter Web sites based on URLs and block cookies and scripting languages that can install spyware on company computers. In addition, these firewalls can be set up to allow Internet access only during certain times of the day and block unproductive bandwidth-heavy content, such as streaming media and peer-to-peer file sharing.

If you're running a wireless network, take extra care to ensure that wardrivers can't break into your wireless network and gain free reign of its resources. The wireless security precautions discussed earlier in this chapter and in Chapter 10 should be used. In addition, using an AP that allows adjusting the signal strength might be worthwhile. With some APs, you can adjust the wireless signal's strength so that only devices in close proximity can hear the signal. You can adjust the strength so that all your wireless devices receive the signal, but someone walking by outside can't.

Managing a Small-Business Network

Unlike a large business with its own IT staff, a set-it-and-forget-it approach doesn't usually work for a small-business network. There are hard drives to defragment, virus scanners to update, OS patches to install, and many other tasks. For this reason, working out a maintenance schedule and contract is usually a good idea. Some tasks can be automated, but others, such as software updates and disk cleanup, aren't as easy to automate. Setting up a weekly or monthly visit for maintenance keeps you in front of the small-business owner, inspiring confidence as well as making you the prime choice of vendor when more work needs to be done.

In managing a small-business network, there's nothing like personal contact. However, sometimes onsite visits are impractical or unnecessary. In these situations, remote access might be the best way to solve a problem quickly and easily. The following list describes some ways to set up remote access to a network:

- *VPN*—If the business is connected to the Internet through a broadband connection, a VPN is probably the best method for accessing and supporting the network remotely. You can connect to the company network securely from wherever you have an Internet connection and establish a remote desktop session with servers and workstations or monitor and update devices on the network. Windows has a built-in remote desktop application for remote control of a computer's desktop. Virtual Network Computing (VNC; www.realvnc.com) is another remote control/remote desktop application that works across several platforms, including Linux, UNIX, and Windows.
- *Dial-up*—Dial-up access is another option for accessing a network remotely, although it's less convenient. You can use Windows Server 2008 to configure remote dial-up access to a Windows network, or third-party products can be used on servers or desktops. Dial-up is sometimes used as a last resort when the network is down and VPN access isn't possible. If the situation warrants, having dial-up access to a network as a backup can save you a long trip.
- *Telnet*—Telnet is one way to gain command-line access to a computer or network device. It should be used when a secure connection has already been established, such as through a VPN. Telnet isn't a secure protocol, so usernames and passwords are sent across the network in unencrypted plaintext. Telnet is best used to access Linux or UNIX systems and command-line-based routers and switches. In a pinch, Telnet can also be used to manage a Windows system.
- *Secure Shell (SSH)*—SSH is a secure method of gaining command-line access to a computer or network device, and when it's supported, it should be used in place of Telnet because communication is encrypted. SSH is available for Windows and Linux systems.
- *Windows Remote Assistance*—For user help, Windows Remote Assistance is an option that doesn't require a VPN, as access to the user's computer is by invitation only. In addition, the user must be sure remote control is enabled. Remote control is enabled by default when Remote Assistance is enabled, and you can configure it in the Remote tab of the System Properties dialog box. The company firewall must also be configured to allow connections to the remote desktop port, which is 3389 by default.

However remote access to the network is set up, it must be done securely. Even a business following an open security policy shouldn't apply the policy to remote connections.

The technology needs of small businesses can be varied and complex. Like the people who own the business, each has its own quirks and requirements. Rarely can you devise a one-size-fits-all solution for a small-business network, but if you're prepared for something new every time you visit a new business, you'll soon have an arsenal of tools, tips, and tricks that make supporting small-business networks easier and help make the businesses you support more successful.

Chapter Summary

- Most small businesses have modest network requirements that don't require advanced WAN technologies, data encryption, or highly restrictive security policies.
- A server-based solution is often the best solution, but a peer-to-peer network is an option. Either way, it's best to design the simplest file-sharing solution that meets the organization's requirements.



- Windows 7 has a new simple file-sharing option called HomeGroup, which is used to create a password-protected file and printer sharing network without needing to set up user accounts and permissions on each computer in a homegroup.
- The two most common choices for file sharing in a Linux environment are Samba and NFS. Samba is compatible with Windows file sharing; NFS works best when most computers run Linux or UNIX.
- Most computer manufacturers maintain small-business solution centers offering equipment that focuses on the needs of small businesses. When purchasing servers, buy as much hardware as the budget allows that will meet needs for the next two to three years.
- When choosing network equipment, you need to decide between a wired and wireless network. In most cases, a wired solution works best for stationary systems, and wireless can be used for laptops and mobile users. In either case, selecting from the hardware vendor's SMB products is best.
- Internet connections and remote access usually require a broadband connection and a router. Some routers have built-in VPN capability.
- Small-business application requirements can range from simple and straightforward to very complex. Both single-user and multiuser versions are available for many applications.
- Typically, small businesses need office suites, accounting software, and sales and contact management software. Windows Small Business Server is a good alternative to Windows Server because it includes built-in e-mail, calendaring, and database applications as well as security features. Linux can be an alternative to Windows desktop OSs if there are no Windows-specific applications that all employees must run.
- Working with small businesses requires excellent communication skills and the ability to make proposals that business owners can understand. Security shouldn't be neglected, and devising a reliable backup scheme is a must.
- To manage a small business, remote control options should be considered, including Remote Desktop through a VPN connection, dial-up, Telnet or SSH, and Windows Remote Assistance.

Key Terms

client-to-gateway VPN mode This VPN mode establishes a VPN connection between a single client computer and a VPN device.

customer relationship management (CRM) A category of software designed to help businesses manage customers and sales prospects.

end user license agreement (EULA) A license that governs how an application can be used. It specifies how many users are allowed to use an application, how many times it can be installed, and whether the software can be copied, among other things.

gateway-to-gateway VPN mode This VPN mode establishes a connection between two routers that support VPNs.

homegroup A peer-to-peer networking feature introduced in Windows 7 that simplifies sharing files and printers between computers.

port forwarding The process by which a router forwards a request for a TCP or UDP port to a specified computer.

total cost of ownership (TCO) The cost of a product or service when intangibles, such as support costs and productivity gains or losses, are factored in.

Review Questions

1. Which of the following is one of the most important aspects of supporting a small business?
 - a. Finding the most inexpensive solution
 - b. Using a “canned” solution you can apply to all small businesses
 - c. Listening to the company’s requirements and designing a solution
 - d. Using your experience to tell the business what it needs
2. Which of the following is a possible problem with a peer-to-peer network solution? (Choose all that apply.)
 - a. Users can unknowingly sever access to shared files or printers.
 - b. A failure on one computer could cause the network to crash.
 - c. Sensitive data could be made available to unauthorized users.
 - d. The centralized server could cause security leaks.
3. Which of the following is true of the HomeGroup feature?
 - a. Homegroups are password-protected.
 - b. The network type can be Work or Home.
 - c. A Public network joins a homegroup automatically.
 - d. New user accounts must be created on each machine.
4. Which is true about sharing files in a Windows domain?
 - a. The user database is distributed among all domain members.
 - b. At least two domain controllers are required.
 - c. The user database is centralized.
 - d. All computers except the server are workgroup members.
5. Which of the following is true about an NAS? (Choose all that apply.)
 - a. It usually has its own user interface.
 - b. It works only in a Windows domain.
 - c. Users and groups can be created on an NAS.
 - d. Most NASs are too big for small-business networks.
6. Which is true about Windows SBS? (Choose all that apply.)
 - a. It includes an e-mail server.
 - b. It requires more IT support than server OSs.

- c. It supports 300 users.
 - d. It provides intranet functionality.
7. Which of the following is true about using Linux as a desktop OS? (Choose all that apply.)
- a. Office applications are available that are compatible with Microsoft Office.
 - b. Virtualization can be used to run Windows applications.
 - c. It's difficult to install.
 - d. Open-source licensing makes it more expensive than Windows.
8. A domain controller simplifies resource management by doing what?
- a. Distributing account creation
 - b. Providing larger hard drives
 - c. Eliminating the need to log on
 - d. Centralizing accounts
9. What are the most common options for sharing files in a Linux environment? (Choose all that apply.)
- a. NetBIOS
 - b. NFS
 - c. TCP
 - d. Samba
10. Which of the following is a common fault-tolerant disk configuration for servers?
- a. RAID 0 for the OS and RAID 1 for the data
 - b. RAID 5 for the OS and RAID 3 for the data
 - c. RAID 1 for the OS and RAID 5 for the data
 - d. RAID 3 for the OS and RAID 0 for the data
11. Which of the following is a consideration when purchasing a switch? (Choose all that apply.)
- a. Switch speed
 - b. Support for multiple media types
 - c. Support for multiple Network-layer protocols
 - d. Whether it's managed or unmanaged
12. Which of the following is a consideration when selecting wireless network equipment? (Choose all that apply.)
- a. Support for the 802.5 protocol
 - b. Category 6 cable connections for gigabit transfers
 - c. Security standards supported
 - d. Interference from other wireless devices

13. To run a Web server on a network protected by a SOHO router, you must enable what feature?
 - a. Address translation
 - b. Port forwarding
 - c. Port filtering
 - d. Address filtering
14. Which of the following is a VPN mode? (Choose all that apply.)
 - a. Client-to-endpoint
 - b. VPN-to-router
 - c. Gateway-to-gateway
 - d. Client-to-gateway
15. What legal document should be read carefully before purchasing software for multiple users to run?
 - a. ELAN
 - b. EULA
 - c. Readme file
 - d. User's manual
16. Which of the following describes the cost of a product when intangibles are factored in?
 - a. CRM
 - b. VPN
 - c. TCO
 - d. EULA
17. Which network remote access method provides a secure connection over the Internet?
 - a. Dial-up
 - b. Telnet
 - c. UDP
 - d. VPN
18. Which of the following is a feature a firewall router provides? (Choose all that apply.)
 - a. Protection against DoS attacks
 - b. Web site filtering
 - c. Cookie blocking
 - d. Protection against IP spoofing

Challenge Labs



Challenge Lab 11-1: Creating a Peer-to-Peer Small-Business Network

Time Required: 1 to 2 hours or more

Objective: Create a small-business network according to specifications.

Required Tools/Equipment: Four computers with a Windows desktop and/or Linux OS installed, one printer, and a wireless router/AP

Description: This lab can be done in groups. Set up a peer-to-peer small-business network for a company named SmallBiz. Three computers are wired, and one is wireless. If possible, one computer should run a Linux OS, such as Ubuntu. The network should look similar to Figure 11-10. One computer will run a Web server that should be accessible from the outside. Build this network to meet these requirements:




Figure 11-10 The Challenge Lab network layout

Courtesy of Course Technology/Cengage Learning

- The HomeGroup feature can't be used.
- Name these computers Marketing, Manager, Operations, and Mobile (the wireless computer).
- The wireless computer doesn't share any files.

- IP addressing should be static, and you can choose a suitable addressing scheme.
- The three wired computers have a dedicated folder set up to share files. The shares should be named appropriately.
- Users should be created as follows: MktUser, MgrUser, OpsUser, and MobUser.
- The shared documents on each computer should have at least read/write access by the user who uses the computer. For example, MktUser might be assigned Full Control access to the Marketing files. The other users should have Read access to the shared files. An Administrator account has Full Control access to all shared files.
- The router should be configured so that a Web site on the Marketing computer can be accessed from the outside.
- The wireless network should be secure and use an appropriate SSID.
- A printer should be configured on the Manager computer and shared. All other computers should be able to print to this printer.
- Develop a backup scheme stating how backup should be done. (You don't need to set up the Backup utility.)

Write a report documenting the network that could be used by an outside consultant who might need to troubleshoot or expand the network. In addition, write a paragraph or two stating whether a domain-based or other centralized server-based (such as an NAS) network might be a better solution.



Case Projects



Case Project 11-1

This project can be done in groups. You're going into business as a computer networking consultant, and you want to be sure all your potential clients get the same service. Devise a questionnaire that you and your other employees can use when interviewing a client about computer and networking requirements. Be sure to cover as many bases as you can think of, including but not limited to number of users, security, resource sharing, Internet access, applications, budget, existing cabling and equipment, and support needs. Save your questionnaire for use in the next project.

Case Project 11-2

Your instructor will concoct a fictitious small business for the purposes of this project. Each group should use the questionnaire designed in Case Project 11-1 to interview the instructor about the business's networking requirements. After the interview, each group should develop a proposal to submit to the business. The proposal should specify only solutions to the business's requirements and shouldn't include pricing yet. Each proposal should be presented to the entire class. Groups

can revise their proposals based on feedback from the class and the instructor's suggestions. A final proposal should then be submitted to the instructor.

Case Project 11-3

Based on the final proposal submitted in Case Project 11-2, each group should create a detailed quote for equipment and services. Good sites to find information on pricing include *www.tigerdirect.com*, *www.newegg.com*, *www.lanshack.com*, and *www.cyberguys.com*, but your group can use other resources to determine costs. Be sure to include labor costs at \$65 per hour (to keep labor rates consistent for all proposals). All items in the quote must be tied to the proposal submitted in Case Project 11-2. All quotes and final proposals should be presented to the class. The instructor will select a vendor based on the proposal's completeness and the price quote.



chapter

12

Wide Area Network Essentials

After reading this chapter and completing the exercises, you will be able to:

- Describe the fundamentals of WAN operation and devices
- Discuss the methods used to connect to WANs
- Configure and describe remote access protocols
- Describe the three major areas of cloud computing

Wide area networks connect LANs to LANs over a large geographic area. The Internet is the ultimate WAN, essentially connecting every LAN to every other LAN across the entire globe. Whether a company wants to connect its LAN to the Internet or to another corporate LAN across town or across the country, WAN technologies are used. This chapter discusses some devices used in WANs and the main methods of making a WAN connection. WAN technology has its own terms and acronyms, and this chapter explains some of the language used. Wide area networking is a vast topic because so many technologies can be used for WAN connections. This chapter gives you a brief introduction to some of these technologies to prepare you for further study of this complex topic.

Remote access is an extension of WAN communication, in which a network provides methods for remote users to connect to the LAN, and this chapter discusses the most common methods of remote access. Finally, this chapter gives you an overview of a growing trend in networking called cloud computing, which uses WANs to offer companies a range of services, including application services and infrastructure services.

Wide Area Network Fundamentals

Large, and sometimes even small, businesses often have multiple sites. For example, a company might have sales offices in New York and Los Angeles and a manufacturing plant in Chicago. Facilitating communication between geographically dispersed sites requires a WAN, which is simply an internetwork spanning a large geographical area.

From a user's perspective, WANs provide access to network resources the same way LANs do, albeit sometimes slower depending on the technology used. As you have learned, both internetworks and WANs can be described as two or more LANs connected. The most obvious difference between internetworks and WANs is the distance between the LANs being connected. However, aside from distance, WANs differ from internetworks in two other critical areas:

- WANs use the services of carriers or service providers, such as phone companies and ISPs, for network connections, whereas internetworks are confined to a building or campus where the internetwork's owner owns and operates all the technology.
- WANs use serial communication technologies that can span distances measured in miles compared with typical LAN technologies that span distances measured in hundreds of meters.

WAN technologies provide the same function in a network as LAN technologies, such as Ethernet. Looking at the two layered architectures discussed in this book, WAN technologies operate at the Network Access layer of the TCP/IP model and the Data Link and Physical layers of the OSI model. The top layers are unchanged when WAN technologies are part of the network. Network protocols and applications are unaware that data's traveling over a WAN connection across thousands of miles versus over a 1 GB Ethernet connection across the room.

WAN Devices

Because WANs operate at the layers involving media access and WAN signals must traverse long distances, it follows that WANs use different methods for accessing network media and transmitting bit signals. Several types of devices are likely to be used in WANs for media access, signal transmission and reception, and connecting a WAN to a LAN:

- Modems
- Channel service units/data service units
- Routers

Modems A **modem** is a device that allows a computer, which works with digital information, to communicate over lines that use analog signals. For example, the telephone system and cable Internet networks use analog communication. A modem converts a digital signal from a computer into an analog signal to be transmitted over phone or cable lines. This conversion is called “modulation.” A modem modulates the digital signal into an analog signal, and at the other end of the line, another modem demodulates the analog signal back to digital. The term “modem” is just a shortened form of “modulator/demodulator.”

A **digital signal** is a series of binary 1s and 0s represented by some type of signal that has two possible states. For example, on copper media, two voltage levels, such as 5v and 0v, might be used, with 5v representing a 1 bit and 0v representing a 0 bit. On fiber-optic media, a 1 bit can be represented as a pulse of light and a 0 bit by the absence of light. Figure 12-1 shows a digital signal represented as a square wave.




Figure 12-1 A digital signal represented as a square wave

Courtesy of Course Technology/Cengage Learning

An **analog signal** is a signal that varies over time continually and smoothly. Whereas a digital signal is 0v or 5v but no value in between, an analog signal in the same voltage range transitions smoothly from 0v to 5v and every voltage value in between. You can look at a digital signal as a light bulb controlled by a typical light switch: It’s either off, giving no light, or on at full brightness. An analog signal, on the other hand, is like a light bulb controlled by a dimmer switch that transitions from off to full brightness after passing through every other level of brightness. An analog signal is represented by a sine wave, as shown in Figure 12-2.




Figure 12-2 An analog signal varies continually

Courtesy of Course Technology/Cengage Learning

Modems used in WAN communications are most closely associated with dial-up communication over the public switched telephone network (PSTN) or with cable or DSL modems. Dial-up networking over the PSTN is now used most often to give mobile users and telecommuters remote access to a network and is being replaced rapidly by VPNs over Internet connections.

CSUs/DSUs A **channel service unit/data service unit (CSU/DSU)** is a device (actually two devices that are usually combined) that creates a digital connection between a LAN device, such as a router, and the WAN link from the service provider. The WAN link is usually a T-carrier technology, such as a T1 or T3 (discussed later in “Leased Lines”). A CSU/DSU performs a function similar to a modem’s—converting WAN signals into a form a LAN can use and vice versa—but in this case, all the signals are digital, so a CSU/DSU converts one type of digital signal to another type of digital signal.

Routers As you know, a router is responsible for getting packets from one network to another. In a WAN, it’s usually the device connecting a LAN to the WAN service provider. In most cases, it connects to the modem or CSU/DSU, which then connects to the link from the WAN provider. Figure 12-3 shows a typical arrangement in which a LAN connects to the WAN through a router and a CSU/DSU or modem.




Figure 12-3 A typical LAN-to-WAN connection

Courtesy of Course Technology/Cengage Learning

The three types of devices discussed in this section are instrumental in making WAN connections of all types, and one or more of these devices is a part of any type of WAN. Next, take a look at the main methods of making a WAN connection.

WAN Connection Methods

Numerous WAN technologies are available that differ in speed, level of security and reliability, and cost. The methods used to make a WAN connection often dictate the technologies that can be used and the connection’s properties. The four most common connection methods are as follows:

- Circuit-switched
- Leased line

- Packet-switched
- VPN over the Internet

Each method has its strengths and weaknesses and works with different underlying WAN technologies. To make an informed decision, a network administrator must weigh each method's cost, speed, reliability, level of security, and available underlying technologies. The following sections explain these methods along with the technologies they work with.

Circuit-Switched WANs

A **circuit-switched WAN** creates a temporary dedicated connection between sender and receiver on demand. The prime example is nothing fancier than a phone line connection from the PSTN, also known as plain old telephone service (POTS). When you pick up the phone, there's no connection until you dial the destination. If the destination is available, a circuit is created that's maintained until one party ends the communication session. A major drawback of using POTS for a network connection is that at least part of the connection is always analog, which limits the connection speed. Another circuit-switched technology that's all digital is Integrated Services Digital Network (ISDN).

Both POTS and ISDN have lost out to faster technologies for network access, but they're still in use in parts of the world where faster technologies are unavailable or too expensive. So both these circuit-switched technologies are covered in the following sections to give you some familiarity with how they work and what their limitations are.

Plain Old Telephone Service Circuit-switched connections over POTS are limited in bandwidth, partly because of the phone system's analog nature. Even when most of the connection is digital, a digital-to-analog conversion is still done by modems and usually the phone carrier's network. The conversion process degrades signal quality and limits data transfer speeds over POTS to about 56 Kbps. If both ends of the connection use a traditional modem, which requires another set of analog-to-digital conversions, the connection speed is limited to about 33 Kbps.

The most common modem standard for connecting to the Internet or a remote office is V.92, which allows connection speeds up to 56 Kbps by eliminating one of the modulation/demodulation steps in traditional modem communication. As shown in Figure 12-4, traditional modem communication converts a computer's digital data into analog data. The analog signal travels over phone lines until it reaches the telephone company ("telco"), where the signal is converted to digital. The telco then must convert the signal back to analog for the receiving




Figure 12-4 Modem communication with two analog-to-digital conversions

Courtesy of Course Technology/Cengage Learning

modem at the ISP, which converts the signal to digital for the Internet. This two-way conversion limits transfer speeds to 33.6 Kbps because each conversion degrades signal quality.

V.92 modems assume that the network from the telco to the ISP and then to the Internet is all digital. Therefore, instead of modulating analog data into digital data as it's received from the telco, a V.92 modem uses a technique called pulse code modulation (PCM) that digitizes analog signals. It introduces less noise into the signal than traditional modulation/demodulation techniques do, so it boosts the total number of bits per second at which data can be transferred. As shown in Figure 12-5, there's only one analog connection—from the home to the telco. From the telco to the ISP and then to the Internet, the signal is digital.




Figure 12-5 Modem communication using the V.92 standard

Courtesy of Course Technology/Cengage Learning

Two caveats with V.92 communication: There must be only one analog circuit between the modem and the Internet, and 56 Kbps communication works in only one direction—the download direction. This means data from the modem to the ISP travels at about 48 Kbps, but data from the ISP to the modem travels at the V.92 speed of 56 Kbps. Although V.92 technology is capable of data transfer speeds up to 56 Kbps, actual transfer rates depend on line conditions.

Integrated Services Digital Network Integrated Services Digital Network (ISDN) is a digital communication technology developed in 1984 to replace the analog phone system. Not as widely deployed as expected, it's available in many U.S. metropolitan areas as well as most of Western Europe. The ISDN specification defines communication channels of 64 Kbps. When dial-up was the standard access method during the 1980s, ISDN seemed attractive, but with broadband access widely available at speeds up to 10 Mbps and faster, ISDN offers no major benefits in most places. If not for private and SOHO use to establish Internet connections, ISDN might never have been deployed, although it was once used widely in corporate WANs as a backup line. Because ISDN charges are often based on connect time, a company can use ISDN in standby mode so that the WAN connection is established only if the primary connection fails.

ISDN offers speeds two to four times that of a standard POTS modem—not an overwhelming increase in speed but a vast improvement for SOHO users when faster technologies, such as DSL or cable modem, aren't available. ISDN is available in two formats or rates:

- **BRI**—The **Basic Rate Interface (BRI)** format consists of two B-channels (64 Kbps) and a D-channel (16 Kbps). Each B-channel can transmit and receive voice or data independently of the other or bonded together at a speed of 128 Kbps. This arrangement

allows a user to talk on the phone with one B-channel and stay connected to the Internet with the second channel. The D-channel is used for call setup and control.

- **PRI**—The **Primary Rate Interface (PRI)** format consists of 23 B-channels and a D-channel. Each B-channel can be used independently or aggregated to provide up to 1.544 Mbps. As with BRI, the B-channels are 64 Kbps channels, but the D-channel is also 64 Kbps, compared with only 16 Kbps in a BRI.

ISDN connects to a network by using a device called a terminal adapter, which performs similar functions as a CSU/DSU. ISDN creates a circuit-switched network in a similar manner to POTS, but the call setup time is less than a second compared with several seconds for POTS. Sometimes you hear the term “ISDN modem,” but because the signals are digital, no modulation/demodulation of analog signals takes place, so “modem” doesn’t apply. ISDN is still somewhat important for telco networks, but DSL and broadband services have largely supplanted it as an Internet access technology.

In general, circuit-switched networks for WAN access have been passed over by faster, more reliable technologies. Their biggest advantage is low cost, and as much as you trust your phone line, they’re secure. However, bandwidth is low, and because they run over existing phone lines that may or may not be in good condition, their reliability is suspect.

Leased Lines

A leased line provides a dedicated point-to-point connection from the customer’s LAN through the provider’s network and to the destination network. It provides permanent, secure, and dedicated bandwidth limited only by the provider’s technology and how much the customer is willing to spend. In essence, customers own the network connection in the same way people “own” apartments they rent. That is, a line is leased to the customer, who has the full benefits of a dedicated connection to the remote network until the customer (or provider) terminates the lease.

Leased lines are the most expensive way to get WAN connectivity because of the dedicated nature of the link. However, a leased line should be considered in these circumstances:

- When high-quality, 24/7 access is needed
- For mission-critical applications
- When fast upstream as well as downstream communication is required

Leased lines are typically based on one of two types of digital technology: T-carriers and SONET, discussed in the following sections.

T-Carriers T-carriers originated when the dominant phone company at the time, AT&T, wanted to be able to carry multiple phone conversations over the same pair of wires and over longer distances. By digitizing voice data and organizing phone conversations into time slots on the media, AT&T could do just that, and the T-carrier technology was born.

Typical T-carrier lines are T1 and T3 that operate at 1.544 Mbps and 44 Mbps, respectively. **T-carrier lines** are derived from multiple 64 Kbps channels, making a T1 connection a grouping of 24 64-Kbps channels. A T3 line groups 672 64-Kbps channels. The channels aren’t carried over separate wires, although that might be the best way to think of how they work. A T-carrier line uses a single pair of wires for transmitting data and another pair for receiving data. It uses a

signaling method called **time division multiplexing (TDM)**, which allocates a time slot for each channel, making it possible to extract any number of the channels for a particular purpose. If a portion of a T-carrier line is used for one purpose and a different portion for another purpose, the line has been **fractionalized**. So you hear the term “fractional T1” or “fractional T3” to mean that some but not all of the line is being used for a specific purpose. For example, a T1 line can be fractionalized to use half the channels for a teleconferencing application and the other half for other network traffic traveling the WAN.



Each 64 Kbps channel is called a DS0, which means digital signal level 0. The next DS level is DS1, which is 24 64-Kbps channels, equivalent to a T1 carrier.

Multiplexing, or “muxing,” enables several communication streams to travel simultaneously over the same cable segment. Through multiplexing, a T-carrier network supports simultaneous communication links over the same set of cables. T1 uses multiplexing to combine data transmissions from several sources and deliver them over a single cable. After a transmission is received, it’s demultiplexed (demuxed) into its constituent channels, as shown in Figure 12-6 where three 64 Kbps channels are multiplexed into a single 192 Kbps channel and then demuxed into the original three channels.




Figure 12-6 Multiplexing/demultiplexing

Courtesy of Course Technology/Cengage Learning

The capability to mux/demux multiple channels make T-carriers enormously flexible. Customers can pay for a fractional T1 or T3 link if their bandwidth needs are modest and later upgrade to additional channels if their needs increase.

T-carrier lines require a CSU/DSU at each end of the link to convert the signals used by the T-carrier line into signals used by the LAN and to mux/demux the DS0 signals as needed. As Figure 12-3 showed previously, a CSU/DSU usually connects to the customer’s router on one side and the T-carrier line (WAN connection) on the other side. T-carriers can use standard twisted-pair cabling for carrying up to T1 speeds. Multiple T1s or T3s require coaxial or fiber-optic cabling and can also be carried via microwave and satellite transmitters.

T1 lines are the most common WAN connection method in the United States. The equivalent of T-carrier lines in Europe are E-carrier lines with E1 (30 64-Kbps channels) and E3 (480 64-Kbps channels) lines. Japan uses J1 and J3 connections with 24 and 480 64-Kbps channels, respectively.

SONET Networks Synchronous Optical Network (SONET) is a flexible, highly fault-tolerant technology that can carry signals of different capacities over a fiber-optic network. It defines optical carrier (OC) levels from OC-1 to OC-3072, in which each OC level is a multiple of the base OC-1 rate of 51.84 Mbps. Typical SONET data rates are OC-3, OC-12, OC-48, OC-192, and OC-768. An OC-3 SONET connection (155 Mbps) is often

used by networks implementing Asynchronous Transfer Mode (ATM) networks, discussed later in “Packet-Switched WANs.”

Because of the tremendous speeds SONET offers and its consequently high cost, it’s rarely used above OC-3 levels by companies other than ISPs and telcos. For example, OC-12 (622 Mbps) is a common SONET speed used by regional or local ISPs and Web hosting providers for their Internet backbone connections. OC-48 (almost 2.5 Gbps) and higher speeds are typical of large regional ISPs, and large ISPs use OC-192 and even OC-768 SONET networks.

One of SONET’s greatest benefits is its flexibility. SONET networks can carry traffic from a variety of other network types, such as T-carrier, ATM, and so forth, by using multiplexing techniques. SONET can also be arranged in a variety of physical topologies, including point-to-point, point-to-multipoint, star, and ring. The ring topology is the most popular, as it provides fault tolerance. SONET uses a dual-ring topology, like FDDI, that’s self-healing if a portion of one ring fails.



SONET is a North American term. In the rest of the world, the same technology is referred to as Synchronous Digital Hierarchy (SDH).

Although SONET can be used as a leased-line technology in which a customer has dedicated use of the bandwidth between sites or as an Internet connection, it’s often used as the underlying technology of other connection types, such as packet-switched WANs.

Packet-Switched WANs

A **packet-switched WAN** doesn’t create a dedicated connection between sender and receiver; instead, data is transmitted in frames or packets, and each packet is transmitted through the provider’s network independently. This process works much like frame transmission in a LAN that uses switches to deliver each frame from its source to its destination. Changing network conditions could cause frames to be transmitted over different paths in a large network, and instead of having a dedicated circuit over which data travels, data shares bandwidth with your provider’s other customers. At first glance, this sharing might not seem to be a good thing, but most network customers don’t use all the bandwidth a dedicated circuit provides, which means they’re paying for more bandwidth than they use.

Packet-switched networks should be nothing new to you because all LANs and internetworks use a similar procedure to get packets from a source computer to a destination computer. The biggest difference is the technologies that are used. Remember that WANs are defined at the Data Link and Physical layers of the OSI model, so these layers are where the differences lie. The most common packet-switched networks are X.25, frame relay, ATM, and MPLS, discussed in the following sections.

Virtual Circuits Before learning about the specific technologies, it’s important to understand that packet-switched WANs use a virtual circuit to ensure that packets are delivered reliably and at the agreed-on bandwidth level. A **virtual circuit** is a logical connection created between two devices in a shared network. No single cable exists between the two endpoints; instead, a virtual circuit maps out the path through the network of switches between the two devices and allocates the specified bandwidth throughout the circuit. This pathway between sender and receiver is created after devices at both ends of the connection agree on bandwidth requirements and request a pathway.

There are two types of virtual circuits: switched and permanent. **Switched virtual circuits (SVCs)** are established when needed and then terminated when the transmission is completed. The path between two communication points is maintained only as long as it's in active use. SVCs are best when communication between the two points is somewhat intermittent, allowing the circuit's bandwidth to be released for use by other SVCs. **Permanent virtual circuits (PVCs)** are similar to leased lines, in that the pathway between two communication points is established as a permanent logical connection; therefore, the pathway exists even when data isn't being transferred. PVCs are more expensive than SVCs because the circuit's bandwidth stays allocated even if data isn't being transferred, making PVCs a good choice only if communication between endpoints is fairly constant.

X.25 Networks X.25 is a packet-switching technology developed in the mid-1970s that has the advantage of running effectively over older copper phone lines. The X.25 specification provides an interface between public packet-switching networks and their customers. X.25 networks offer both SVCs and PVCs, although not all X.25 providers offer PVCs.

Early X.25 networks used standard phone lines as communication links, which resulted in numerous errors and lost data. Adding error checking and retransmission schemes improved the success of X.25 transmissions but severely reduced speed. With its extensive level of error control, X.25 could deliver only 64 Kbps transmission rates. A 1992 specification revision improved the maximum throughput of X.25 to 2 Mbps per connection, but this new version wasn't widely deployed.

X.25 is usually associated with public data networks (PDNs), such as CompuServe and Euronet, and was used to access these online services in the 1970s and 1980s. It remains popular in developing countries, where digital communication is less available and more expensive than in the North America and Europe. Using data terminal equipment (DTE) and data circuit-terminating equipment (DCE), explained later in “WAN Implementation Basics,” connecting to an X.25 network can be done in one of these three ways:

- An X.25 NIC in a computer
- A packet assembler/disassembler (PAD) that supports X.25 communication for low-speed, character-based terminals
- A LAN/WAN X.25 gateway

Even though X.25 networks offer reliable and error-free communication, this technology has been largely replaced by other higher-speed technologies, such as frame relay and ATM.

Frame Relay Networks Frame relay is a PVC packet-switching technology that offers WAN communication over a fast, reliable, digital link. It was developed from X.25 and ISDN technology. Error checking isn't required on the digital fiber-optic links most frame relay connections use, so overall throughput is improved. Instead, the devices on each end of the communication perform error checking.

Frame relay uses a PVC between communication points, so the same pathway carries all communications, which ensures correct delivery and higher bandwidth rates. A PVC is similar to a dedicated line, in that communication devices aren't concerned with route management and error checking. Instead, all the resources of devices are dedicated to moving data. This is why frame relay technology can maintain transmission rates from 64 Kbps to 44 Mbps (T3 speed). It fills the bandwidth gap between ISDN, which operates at a maximum of 128 Kbps, and ATM, which operates at 155 Mbps.

Frame relay services have grown in popularity. They're inexpensive (compared with leased lines) and allow customers to specify the bandwidth needed. Charges depend on the PVC's bandwidth allocation, also known as its **Committed Information Rate (CIR)**. CIR is the guaranteed minimum transmission rate the service provider offers. Customers can purchase frame relay services in CIR increments of 64 Kbps. Because customers can pay for a customized bandwidth solution, frame relay is sometimes preferred to T1 because it's generally less expensive.

A frame relay connection is established by using a pair of CSU/DSU devices—as with T1 lines—with a router or bridge at each end to direct traffic on and off the WAN link. An important difference between a frame relay connection and a T1 connection is that T1 is a point-to-point link, which means a T1 customer gets full-time bandwidth to the destination. However, frame relay connections are virtual circuits that go through a switch. This arrangement makes it possible to reach multiple destinations with a single connection. Therefore, a corporate customer can, for example, have a frame relay link to each of its branch offices as well as one to the Internet yet require only a single frame relay connection.

Figure 12-7 shows a WAN using leased lines, and Figure 12-8 shows the same WAN using frame relay connections.

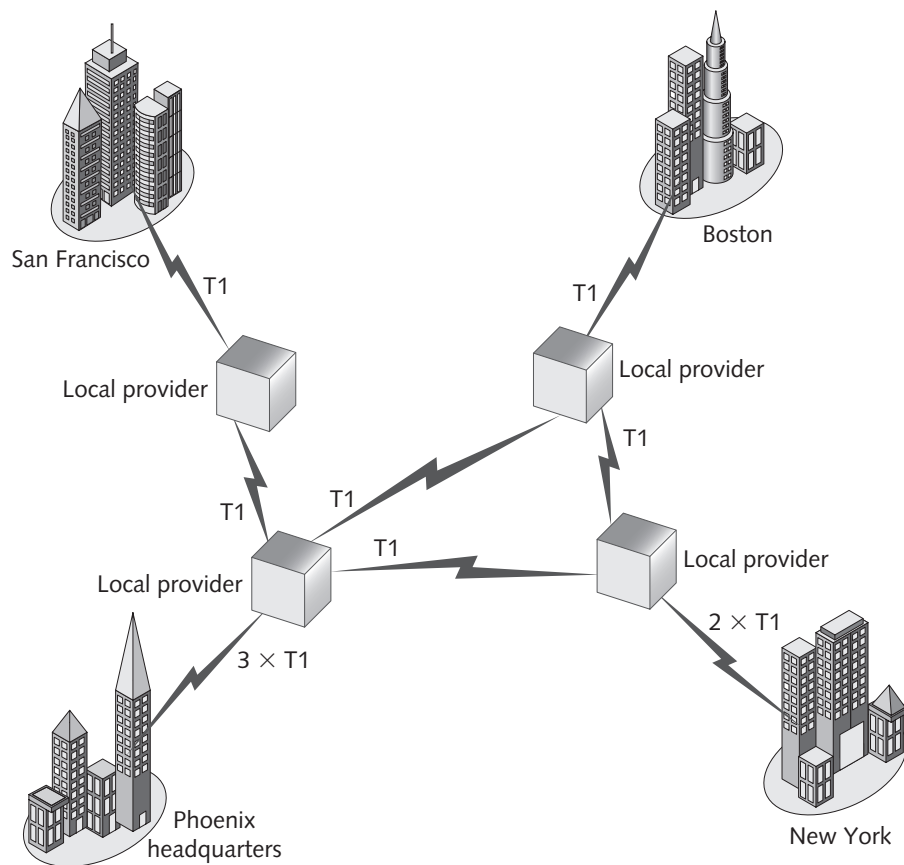


Figure 12-7 A leased-line network

Courtesy of Course Technology/Cengage Learning

The leased-line network has dedicated bandwidth extending from the corporate site to the local provider to the other sites. An additional connection between New York and Boston provides redundancy. The frame relay network allows the company to pay for only the bandwidth it needs between sites and can be configured to create redundant PVCs if necessary, precluding the need for customers to pay for the additional connection.




Figure 12-8 A frame relay network

Courtesy of Course Technology/Cengage Learning

ATM Networks Asynchronous Transfer Mode (ATM) is a high-speed network technology designed for both LAN and WAN use. ATM is a Data Link layer technology, so it specifies the format of data frames and the media access method. ATM can run over fiber-optic SONET networks and T-carrier connections in WANs and Cat-5 or higher UTP and STP cable in LANs. Because it specifies only the Data Link layer, the Physical layer doesn't matter, as long as it supports the bandwidth the ATM application requires. Fiber-optic media is

used most commonly in WAN applications. ATM bandwidth can be as low as a few Mbps up to 622 Mbps (OC-12), but the most common speed is 155 Mbps, an OC-3 SONET connection.

ATM is a cell-based packet-switching technology, in which the cells are a fixed length rather than the variable-length packets used in typical packet-based systems. In an ATM environment, data travels in short, 53-byte cells: 5 bytes of header information and 48 bytes of data. If a transmission (or the last cell in a transmission) is fewer than 53 bytes, the cell is padded to make up the difference. Fixed-length cells can be switched more efficiently than variable-length packets and are more predictable for time-sensitive data, such as voice and video. The predictable nature of fixed-length cells means a higher level of QoS for video and audio streaming and teleconferencing applications, among others.

Like frame relay, ATM uses virtual circuits created between the source and destination before cells are transferred, making ATM a connection-oriented technology that has the reliability of a circuit-switched technology and the flexibility and cost-effectiveness of a packet-switching technology. Conceptually, an ATM WAN looks much like the frame relay WAN shown previously in Figure 12-8.

ATM originated as a telephone company technology and is used quite heavily for the backbone and infrastructure in large communication companies. The wide availability and low cost of Gigabit and faster Ethernet speeds have diminished ATM's role as a LAN technology. When it's used in LANs, a service called LAN Emulation (LANE) must be used that encapsulates Ethernet frames into the correct cell format for transmission on an ATM network.

Because of Gigabit Ethernet's popularity, ATM is best suited for LAN applications in which voice, data, and time-sensitive information travel on the same media. ATM is also a solid choice, along with frame relay, for any type of high-bandwidth WAN application.

MPLS Networks Multiprotocol Label Switching (MPLS) has become a popular WAN technology because of its scalability and flexibility. It works with any Network-layer protocol and is independent of the Data Link layer technology. Consequently, MPLS runs over ATM, frame relay, SONET, and even Ethernet, among other Layer 2 technologies. Because MPLS works as a sort of go-between for Layers 2 and 3, it's often referred to as a Layer 2.5 technology.

MPLS creates a connection-oriented virtual circuit using labels assigned to each packet. A label is used to make packet-forwarding decisions in the MPLS network, making it unnecessary to view packet contents. Other technologies, such as frame relay and ATM, have used similar techniques, but MPLS has the advantage of using the best of these technologies.

Although MPLS is capable of supporting different Layer 3 protocols, currently it's used exclusively in IP networks, supporting both IPv4 and IPv6. It was first envisioned as a technology targeted at improving routing speed because of the simpler function of routing based on labels rather than IP addresses. However, faster routing technologies have made this reason a moot point. Nonetheless, MPLS has evolved into a crucial technology for implementing large-scale IP WANs. Some applications include network traffic engineering



(because it enables network engineers to better control the path traffic takes through a network) and high-speed VPNs. MPLS simplifies VPN deployment in large multisite VPN applications.



For a good overview of MPLS, see www.protocols.com/papers/mpls.htm.

TIP

WANs over the Internet

Using VPN connections over inexpensive Internet connections is a popular WAN alternative to the methods already discussed. The VPN discussion in Chapters 10 and 11 focused on using broadband cable and DSL connections to the Internet and centered on mobile users and telecommuters using VPNs for secure, convenient access. However, the same types of connections can be used for secure communication from branch offices to the corporate office. In addition, VPNs aren't limited to broadband connections; a branch office that needs the dedicated bandwidth of a T1 or multiple T1 connections can use a leased line as its gateway to the Internet and run a VPN over it to establish a secure connection with the central office. VPNs offer the following advantages over other WAN methods:

- *Inexpensive*—The cost of Internet access is much lower than leased lines or packet-switched WAN connections, particularly if broadband Internet is used.
- *Convenience*—A VPN can be configured as soon as Internet access is established, which can usually be done much faster than the sometimes complex installation and configuration of a traditional WAN connection.
- *Security*—Advanced authentication and encryption protocols protect the integrity and privacy of VPN traffic. Leased lines and packet-switched WANs, although usually considered secure, are still vulnerable to prying eyes while moving through the provider's network.
- *Flexibility*—After a VPN infrastructure is in place, it's available for WAN connections from branch offices as well as mobile users and telecommuters.

Although there are plenty of advantages in using a VPN for your WAN connections, there are some drawbacks, most notably the unpredictable nature of the Internet. A VPN probably isn't not a good solution for mission-critical applications because unlike a traditional WAN, you're paying only for the connection between your site and your ISP. After the data goes outside your ISP's network, the connection's speed and reliability are out of the ISP's hands and out of your control. With a traditional WAN, you're paying for an end-to-end connection, and your provider is responsible for your data throughout its entire journey.

VPNs are made secure through encryption and authentication protocols, but the data nonetheless travels through the public Internet. Given enough time and resources, almost any encryption protocol can be broken, and some companies dealing with highly sensitive data might prefer that their data travel on a private dedicated network.

Table 12-1 summarizes the WAN connection methods discussed in this section.

Table 12-1 WAN connection methods

WAN connection method	Description	Technologies	Advantages	Disadvantages
Circuit switched	Creates a temporary dedicated connection between sender and receiver	POTS, ISDN	Inexpensive, available everywhere, dedicated bandwidth	Low bandwidth; nonpermanent connect requires setup time; suspect reliability
Leased line	Provides a dedicated, permanent point-to-point link from end to end	T-carriers, SONET	Secure, reliable, dedicated bandwidth for mission-critical applications	Expensive
Packet switched	Packets travel through virtual circuits created in the provider's switches	X.25, frame relay, ATM	Pay only for bandwidth you need; less expensive than leased lines	Bandwidth shared with other customers; potential security concerns
VPN over Internet	Uses an Internet connection with authentication and encryption to provide a secure connection to the LAN	Broadband cable modem and DSL, dial-up Internet, FiOS, leased lines	Convenient and inexpensive; uses existing Internet connections; available anywhere an Internet connection is available	Unpredictable service level because of the nature of the Internet



As Internet access speeds become faster, with technologies such as Fiber Optic Service (FiOS) offered by Verizon bringing them up to 50 Mbps speeds, VPN over Internet becomes a more attractive WAN option than ever.

WAN Equipment

You have already learned some terms for the technologies that make WANs work, such as ISDN, ATM, and frame relay. The following sections explain how WANs are implemented.

Customer Equipment When an organization must build a WAN to connect geographically dispersed resources, some equipment is the organization's responsibility and some is the provider's responsibility. The organization requiring a WAN's services is always referred to as the customer, and the equipment at the customer site that's usually the customer's responsibility is called the **customer premises equipment (CPE)**. The customer might own or lease the equipment from the provider. CPE usually includes devices such as routers, modems, and CSUs/DSUs. Modems are needed when some type of analog connectivity is involved, and CSUs/DSUs are required for digital circuits.

Every WAN has a connection from the customer equipment (usually a cable from the CSU/DSU or modem) to a junction panel called the **demarcation point**: the point at which the CPE ends and the provider's responsibility begins. This junction is where the physical WAN connection is made from the customer to the telco or ISP (the provider).

Provider Equipment The provider location nearest the customer site is usually referred to as the central office (CO), and the network medium runs from the customer site demarcation point to the CO of the WAN service provider. The medium is usually coaxial copper or fiber-optic cable and is the provider's responsibility. For a wireless connection to the provider, a wireless transmitter is usually mounted on the customer's building. The connection between the demarcation point and the CO is called the **local loop** or **last mile**. The equipment specific to the WAN technology is usually placed at the CO. This equipment might be a frame relay switch, an ATM switch, a CSU/DSU, or another WAN device, depending on the type of WAN connection.

Going the Last Mile The CPE must be able to send data in the correct format to the connection that makes up the local loop and receive data coming from this connection. This requirement is where the CSU/DSU or modem comes in. The device that sends data to (and receives data from) the local loop is called **data circuit-terminating equipment (DCE)** or sometimes "data communications equipment," and the CSU/DSU or modem is called the DCE device. The device that passes data from the customer's LAN to the DCE is called **data terminal equipment (DTE)**. A typical DTE is a router or bridge that has one connection to the customer's LAN and another connection to the DCE that makes the WAN connection. Figure 12-9 illustrates this arrangement.




Figure 12-9 A WAN connection showing the CPE, demarcation point, and local loop

Courtesy of Course Technology/Cengage Learning

As you can see, getting all these definitions straight is half the struggle of understanding how to specify, design, and support these technologies.

Remote Access Networking

Large and small businesses alike are using fast, affordable remote access technologies that enable employees to access their office desktops and company resources from home and while on the road. VPNs are the favored method for remote access connections, but dial-up access is still widely supported by client and server OSs.

Windows server OSs include **Routing and Remote Access Service (RRAS)**, which supports both dial-up remote access and VPN remote access. Windows client OSs can create dial-up or VPN network connections to an RRAS server. Figure 12-10 shows how Windows RRAS might interact with a LAN to provide remote access. In this figure, RRAS provides both dial-up and VPN connections for remote clients. RRAS offers local area routing services as well as the capability to route between one or more remote or local connections.




Figure 12-10 Windows RRAS provides remote connectivity to clients

Courtesy of Course Technology/Cengage Learning

With an RRAS server configured on the LAN, users can dial in over POTS or use a VPN with any type of Internet connection. After the remote access connection is established, the remotely connected computer acts as though it were connected directly to the network, albeit more slowly.



The option for users to connect to a Windows remote access server is disabled by default for security reasons. This feature must be enabled in a user's account settings and/or set by configuring remote access policies in RRAS.

Making a VPN Connection in Windows

Configuring a VPN connection in Windows is a straightforward process. In Windows 7, you create a new connection in the Network and Sharing Center by selecting “Set up a new connection or network,” which starts the Set Up a Connection or Network Wizard (see Figure 12-11). To set up a VPN connection, click “Connect to a workplace.”

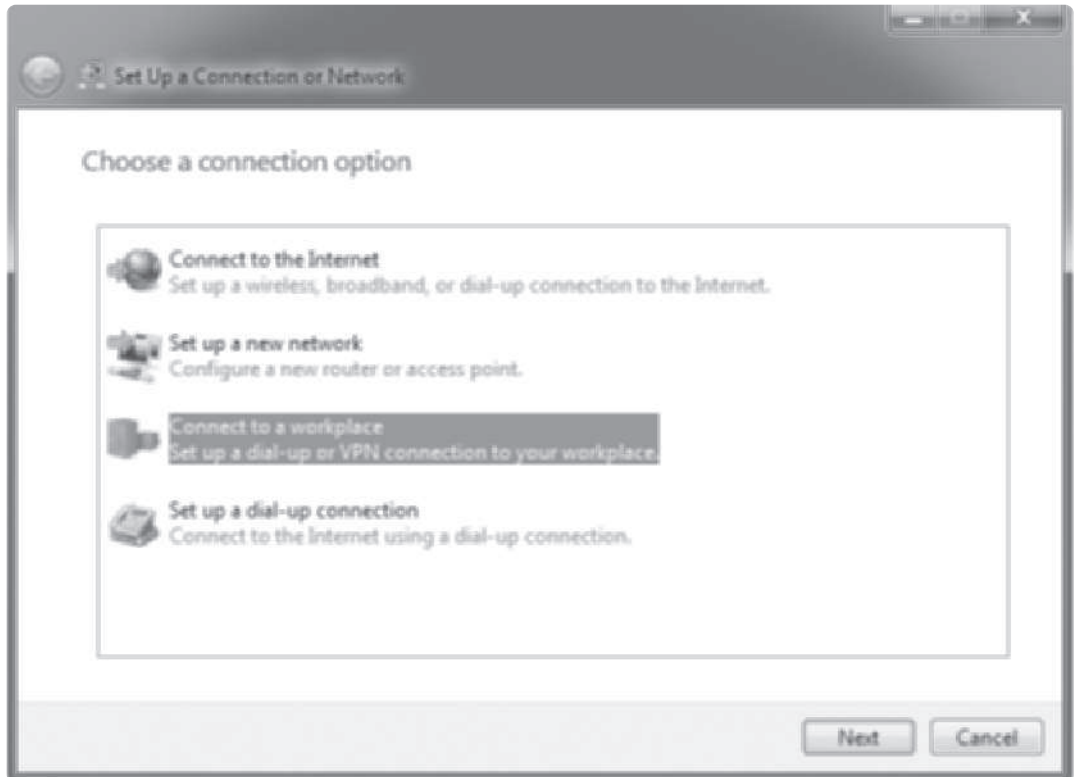


Figure 12-11 The Set Up a Connection or Network Wizard

Courtesy of Course Technology/Cengage Learning

In the next window, you specify using your Internet connection or dialing in directly. For a VPN, choose “Use my Internet connection.” From there, you simply enter the name or address of the VPN server along with your username and password. Windows 7 VPN connections attempt to make the connection by using the PPTP, L2TP, and SSTP protocols until a connection is made. Alternatively, you can configure using a specific VPN protocol in the VPN Connection Properties dialog box (see Figure 12-12).

As you can see, there are a number of options for authenticating to the VPN. The default settings usually work when connecting to an RRAS server, but other authentication options might be necessary to connect to third-party VPN servers. After the VPN connection is made, the client connection is assigned an IP address on the network.

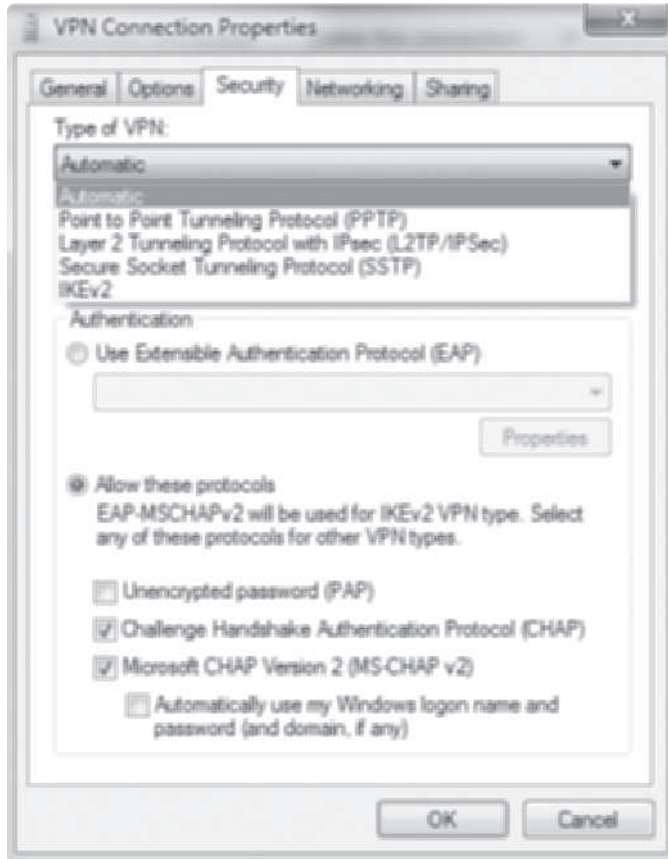


Figure 12-12 Selecting a VPN protocol

Courtesy of Course Technology/Cengage Learning



Making a Dial-Up Connection

All versions of Windows, starting with Windows 95, include **Dial-Up Networking (DUN)** software to make an RRAS connection. The DUN client also connects computers to ISPs for dial-up Internet access. The protocol that makes dial-up and most VPN connections possible is **Point-to-Point Protocol (PPP)**. It operates at the Data Link and Network layers and is used to carry a variety of protocols over different types of network connections. Although PPP is best known for use in dial-up networks, it's also used to provide advanced features to other types of WAN connections, such as leased lines and ISDN.

PPP allows WAN connections to use a variety of network protocols, including TCP/IP and IPX/SPX, and includes advanced authentication services. Its flexibility comes from the two protocols that are integral to PPP:

- *Link Control Protocol (LCP)*—LCP sets up the PPP connection, defines communication parameters and authentication protocols, and terminates the link when a call or connection is ended. It works at the Data Link layer of the OSI model.

- *Network Control Protocol (NCP)*—NCP encapsulates higher-layer protocols, such as IP, and provides services such as dynamic IP addressing. It can carry multiple protocols because it contains a field to indicate which protocol is encapsulated in each packet. NCP works at both the Data Link and Network layers.

In Windows 7, dial-up connections are made much like VPN connections. The main difference is that you select “Dial directly” instead of “Use my Internet connection” and enter the phone number of the RRAS server instead of its Internet address.



Serial Line Internet Protocol (SLIP) is an older protocol PCs use to connect to the Internet via a modem. It provides connectivity across phone lines with no error correction and no secure authentication and doesn't support dynamic IP address assignment. SLIP is no longer supported in Windows client OSs or Windows Server 2008 RRAS.

Remote Access Networking via the Web

Dial-up networking and VPN connect your computer to a network, making it behave as though it were located on the LAN. Another remote access model in common use is remote control of your desktop via a Web browser. Several online services are available for connecting your Web browser to your desktop, including LogMeIn and GoToMyPC. Using these services, you install a client component on your computer and then log on to the online service, which connects your Web browser to your computer's desktop. These services use authentication and encryption to maintain a secure connection. No-frills versions of these services are available free, and you can purchase services offering advanced features, such as remote file transfer and desktop sharing with other users, for a monthly or annual fee.

Using a third-party remote desktop solution is convenient and frees IT staff from having to support dial-up or VPN servers. However, not all networks permit this type of remote access for security reasons. Microsoft has a solution called Terminal Services Gateway (TSG) that allows remote desktop connections via SSL, the protocol that secures communication between Web browsers and Web servers. TSG doesn't use a Web browser to make the remote desktop connection. Instead, remote users use the standard Remote Desktop Protocol (RDP) client to access their desktops without having to establish a VPN or dial-up connection. After the TSG gateway is set up, remote users just specify the TSG server settings in the RDP client (see Figure 12-13).



Figure 12-13 TSG server settings on the RDP client

Courtesy of Course Technology/Cengage Learning



Hands-On Project 12-1: Creating a Dial-up Connection in Windows 7

Time Required: 10 minutes

Objective: Create a dial-up connection in Windows 7.

Required Tools/Equipment: Your classroom computer running Windows 7

Description: In this project, you configure a dial-up connection in Windows 7. You probably don't have a modem installed, so you simply go through the steps using an invalid phone number.

1. Log on to your computer as an administrator.
2. Open the Network and Sharing Center, and click **Set up a new connection or network**.
3. In the Choose a connection option window, click **Connect to a workplace**, and then click **Next**.

4. In the How do you want to connect? window (see Figure 12-14), click **Dial directly**. If your computer has no modem installed, you see the message “Windows could not detect a dial-up modem.” Click **Set up a connection anyway**.

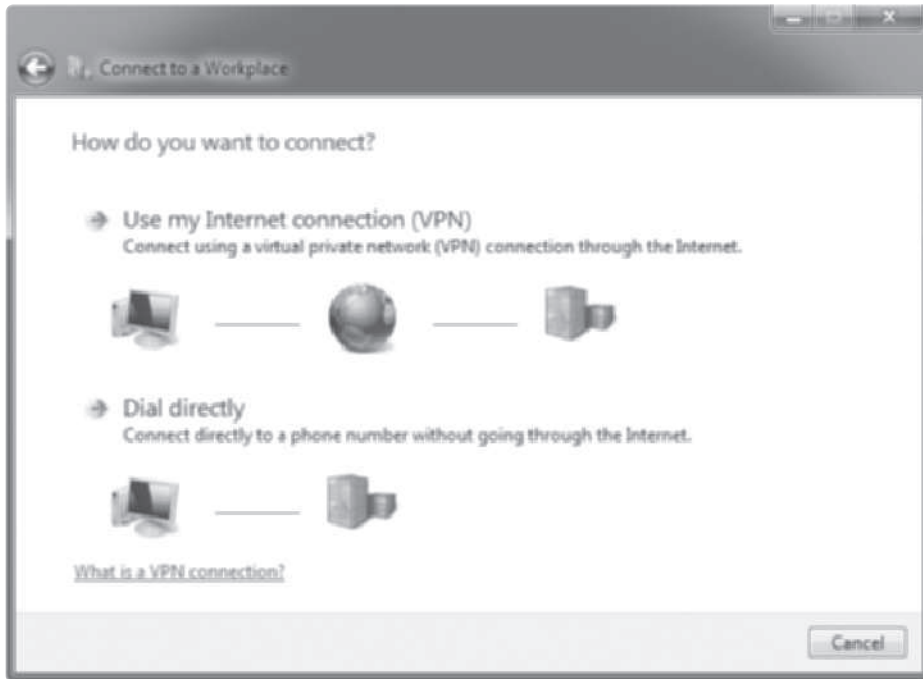


Figure 12-14 Selecting a connection type

Courtesy of Course Technology/Cengage Learning

5. In the Type the telephone number to connect to window, type **555-1234** in the Telephone number text box, and then click **Next**.
6. In the Type your user name and password window, type **RemoteUser** in the User name text box and **Password01** in the Password text box. If you're dialing into a Windows domain network, you can also enter the domain where your user account resides. Click **Create**.
7. In The connection is ready to use window, click **Close**.
8. Click **Change adapter settings** in the Network and Sharing Center. Right-click **Dial-up Connection** and click **Properties**.
9. Click the **Options** tab, and review the settings you can configure. Click **PPP Settings**.
10. In the PPP Settings dialog box, review the options you can configure for PPP, and then click **Cancel**.
11. In the Dial-up Connection Properties dialog box, click the **Security** tab, and review the encryption and authentication options.
12. In the Dial-up Connection Properties dialog box, click the **Networking** tab. Notice that all protocols except File and Printer Sharing for Microsoft Networks are selected. This

protocol is disabled for remote access connections to prevent remote users from accessing your shared files. Click **Cancel**.

13. Stay logged on for the next project.



Hands-On Project 12-2: Creating a VPN Connection in Windows 7

Time Required: 10 minutes

Objective: Create a VPN connection in Windows 7.

Required Tools/Equipment: Your classroom computer running Windows 7; if possible, the instructor can set up an RRAS VPN server

Description: In this project, you configure a VPN connection in Windows 7. If you have a VPN server to connect to, ask your instructor for the correct address, which you need in Step 5. You also need a username and password.

1. Log on to your computer as an administrator, if necessary.
2. Open the Network and Sharing Center, and click **Set up a new connection or network**.
3. In the Choose a connection option window, click **Connect to a workplace**, and then click **Next**. In the Do you want to use a connection that you already have? window, make sure **No, create a new connection** is selected, and then click **Next**.
4. In the How do you want to connect? window, click **Use my Internet connection (VPN)**.
5. In the Type the Internet address to connect to window, type **vpn.mydomain.com** in the Internet address text box if you don't have an actual VPN server to connect to. If your instructor gave you a VPN server's name or IP address, enter this information instead. Click **Next**.
6. In the Type your user name and password window, type **VPNUser** in the User name text box and **Password01** in the Password text box, unless your instructor gives you a different username and password. If you're connecting to a Windows domain network, you can also enter the domain where your user account resides.
7. If you're connecting to an actual VPN server, click **Connect**; otherwise, click **Cancel**.
8. If you connected to an actual VPN server, open a command prompt window, and then type **ipconfig /all** and press **Enter** to view the address settings assigned to your VPN connection. Close the command prompt window.
9. Close all open windows.



Cloud Computing

Many people rely on the Internet's services for communication, research, and entertainment. This trend has continued to the point that many functions once handled in a company's IT center are now handled by servers on the Internet—what's referred to as cloud computing. **Cloud computing** is a networking model in which data, applications, and processing power are managed by servers on the Internet, and users of these resources pay for what they use

rather than for the equipment and software needed to provide resources. It's like paying only for cell phone minutes you use instead of paying for the cell phone towers and switching equipment needed to make your phone work.

The word “cloud” is used to obscure the details of equipment and software used to actually provide resources. For the most part, customers don't care whether the equipment providing data consists of Windows or Linux servers, large tower computers, or rack-mounted computers, just as you don't care how your cell phone makes a call, as long as it works.

For some companies, cloud computing's allure is based on the following benefits:

- *Reduced physical plant costs*—Having fewer servers means less space is needed to house them, and less electricity and cooling are required to keep servers running.
- *Reduced upfront costs*—Paying only for services and software that are used means a company can avoid the startup costs of purchasing hardware and software.
- *Reduced personnel costs*—Having fewer servers and applications to support means fewer IT employees are needed to support hardware and applications.

Although cloud computing has seemingly limitless applications, three main categories of cloud computing are taking center stage:

- Hosted applications
- Hosted platforms
- Hosted infrastructure

The term “hosted” simply means the resource resides on another server or network than the one using the resource.

Hosted Applications

Hosted applications are also called “on-demand applications” or **software as a service (SaaS)** because the customer doesn't actually buy any software that's installed on its own equipment. Instead, the customer pays for the use of applications that run on a service provider's network. The most well-known example is Google Apps, which a business or a home user can use to run hosted applications, such as e-mail, calendar, word-processing, and spreadsheet programs. More complex applications involve large database systems, such as payroll services from ADP and customer relationship management software offered by companies such as Salesforce.com.

Hosted applications are usually offered as a subscription based on the number of users using the application. They take the burden of installation and maintenance off the customer so that companies can focus on maintaining their LANs and Internet access instead of maintaining hundreds of copies of an installed application. In addition, customers can take advantage of new software editions much faster than the standard deployment times of traditional application upgrades. Some application upgrades require client computer or OS upgrades, but with a hosted application, the vendor handles infrastructure upgrades when needed.

In addition, hosted applications are available anywhere the customer has a connection to the Internet. Mobile users and telecommuters have access to the same applications they use in the office without having to install the software on their laptops or home computers. Some applications can't even be installed on home computers, but with hosted applications, the software runs on remote servers, so local installations aren't necessary.

Hosted Platforms

Hosted platforms—also called **platform as a service (PaaS)**—are similar to hosted applications, but the customer develops applications with the service provider’s tools and infrastructure. After applications are developed, they can be delivered to the customer’s users from the provider’s servers. This setup differs from hosted applications in which the service provider owns the applications delivered to users; with hosted platforms, the customer develops and owns the application and then delivers it to a third party.

Developers who use PaaS can take advantage of many of the same benefits as users of SaaS. In addition, after an application is developed with PaaS, the developer can usually deploy the application immediately to customers who access it as a hosted application. The same operating environment used to develop the application is used to run it, which bypasses the sometimes complex and problem-prone process of migrating from a traditional development environment to a production environment.

The most common hosted platforms are Force.com’s Apex, Azure for Windows, Google’s AppEngine for Python and Java, WaveMaker for Ajax, and Engine Yard for Ruby on Rails. Others are available, but details on these development platforms are beyond the scope of this book. PaaS is still a new model for application development, and platforms will come and go as developers weed out what works and what doesn’t. Developing in the cloud is likely here to stay because it offers benefits that aren’t usually available in a locally managed environment. In addition, because small businesses and individual developers have access to expensive, full-featured development environments, entrepreneurs can be on an equal footing with the big boys, which increases competition and innovation—and that’s always a good thing.

Hosted Infrastructure

Hosted infrastructure, or **infrastructure as a service (IaaS)**, allows companies to use a vendor’s storage or even entire virtual servers as needed. Traditionally, if a company needs another 100 GB of storage to house a new database, it has to buy a new hard drive—assuming the server can accommodate a new hard drive. By using a hosted infrastructure, the company simply pays for another 100 GB of space without worrying about how that space is actually provided. In addition, if a customer needs another server to handle its application workload, it simply pays for the amount of processing and storage the additional server actually requires instead of the physical device. In most cases, IaaS servers are virtualized (as discussed in Chapter 8), meaning they run as virtual machines on more powerful physical servers.

IaaS differs from other hosted services because customers mostly rent the resources they’re using but are still responsible for application installation and upgrade. So although IT staff can be reduced because the IaaS vendor handles physical device upkeep, customers still need IT staff to configure and manage applications and server OSs.

IaaS isn’t just for server infrastructure. Companies can “upgrade” to the latest OS and desktop applications by using virtualized desktops through their IaaS providers. By accessing desktops remotely, IaaS customers can use thin clients (client computers with minimal hardware resources) or computers with older OSs to access the latest desktop OSs and applications.

Cloud computing isn’t for every company or situation, but it offers a flexible array of services that can complement an IT department’s existing resources and sometimes replace them. The

trend toward cloud computing is growing with no abatement in site. Only time will tell whether it's a flash-in-the-pan model or one that will become an integral part of daily computing.

Chapter Summary

- The most obvious difference between internetworks and WANs is the distance between the LANs being connected. WANs differ from internetworks in two other critical areas: the use of carriers or service providers and the use of serial communication technologies that can span long distances.
- Several types of devices are typically used in WANs for media access, signal transmission, and reception and to connect a WAN to a LAN: modems, channel service units/data service units, and routers.
- The methods for creating a WAN connection often dictate the technologies that can be used and the connection's properties. The most common connection methods are circuit-switched WANs, leased lines, packet-switched WANs, and VPN over the Internet.
- WAN equipment can be categorized as customer equipment, provider equipment, and the circuit that makes the connection between the demarcation point and the central office, called the last mile or local loop.
- Large and small businesses alike are using fast, affordable remote access technologies that enable employees to access their desktops and company resources from home and while on the road. VPNs are the favored method for remote access connections, but dial-up access is still widely supported by client and server OSs.
- Cloud computing is a networking model in which data, applications, and processing power are managed by servers on the Internet, and users of these resources pay for what they use instead of the equipment and software needed to provide resources.
- There are three main categories of cloud computing: hosted applications, hosted platforms, and hosted infrastructure.

Key Terms

analog signal A signal, represented by a sine wave, that varies over time continually and smoothly.

Asynchronous Transfer Mode (ATM) A high-speed, cell-based packet-switching technology designed for both LAN and WAN use; uses connection-oriented switches to allow senders and receivers to communicate over a network.

Basic Rate Interface (BRI) An ISDN format that consists of two 64-Kbps B-channels and a 16-Kbps D channel; generally used for remote connections. *See also* Integrated Services Digital Network (ISDN).

channel service unit/data service unit (CSU/DSU) A device that creates a digital connection between a LAN device, such as a router, and the WAN link from the service provider.

circuit-switched WAN A type of WAN connection in which a temporary dedicated connection is established between sender and receiver on demand.

cloud computing A networking model in which data, applications, and processing power are managed by servers on the Internet, and users of these resources pay for what they use rather than for the equipment and software needed to provide resources.

Committed Information Rate (CIR) A guaranteed minimum transmission rate offered by the service provider.

customer premises equipment (CPE) The equipment at the customer site that's usually the responsibility of the customer.

data circuit-terminating equipment (DCE) The device that sends data to (and receives data from) the last mile; usually a CSU/DSU or modem. *See also* channel service unit/data service unit (CSU/DSU) *and* last mile.

data terminal equipment (DTE) The device that passes data from the customer LAN to the DCE; usually a router. *See also* data circuit-terminating equipment (DCE).

demarcation point The point at which the CPE ends and the provider's responsibility begins. *See also* customer premises equipment (CPE).

Dial-Up Networking (DUN) The Windows software component (beginning with Windows 95) for setting up a connection to an RRAS server or connecting computers to ISPs for dial-up Internet access.

digital signal Represented as a square wave, a signal that uses binary 1s and 0s to represent two possible states.

fractionalized The term used to describe a T-carrier line in which portions are dedicated for different purposes. *See also* T-carrier lines.

frame relay A PVC packet-switching technology that offers WAN communication over a fast, reliable, digital link. Throughput is usually improved because error checking is done on endpoint devices instead of on the digital link. *See also* permanent virtual circuits (PVCs).

hosted applications A category of cloud computing in which a customer pays for the use of applications that run on a service provider's network; also called software as a service (SaaS).

hosted infrastructure A category of cloud computing in which a company can use a provider's storage or virtual servers as its needs demand; also called infrastructure as a service (IaaS).

hosted platforms A category of cloud computing in which a customer develops applications with the service provider's development tools and infrastructure; also called platform as a service (PaaS). After applications are developed, they can be delivered to the customer's users from the provider's servers.

infrastructure as a service (IaaS) *See* hosted infrastructure.

Integrated Services Digital Network (ISDN) A digital WAN technology developed to replace the analog phone system. It defines communication channels of 64 Kbps and is most often used by OSHO users.

last mile The connection between a WAN's demarcation point and the central office (CO); also called the local loop. *See also* demarcation point.

local loop *See* last mile.

modem A device that converts a sending computer's digital signals to analog signals for transmission over phone lines and then converts analog signals to digital signals for the receiving computer.

multiplexing A technology that supports simultaneous communication links over the same set of cables, so data transmissions from several sources can be combined and delivered over a single cable.

Multiprotocol Label Switching (MPLS) A highly scalable, flexible WAN technology that works with any Network-layer protocol and is independent of the Data Link layer technology; used exclusively in IP networks. It creates a connection-oriented virtual circuit, using labels assigned to each packet that make it unnecessary to view packet contents.

packet-switched WAN A type of WAN network in which data is transmitted in frames or packets, and each packet is transmitted through the provider's network independently. Instead of having a dedicated circuit over which data travels, a provider's customers share the bandwidth.

permanent virtual circuits (PVCs) Pathways between two communication points that are established as permanent logical connections; therefore, the pathway exists even when it's not in use. *See also* virtual circuit.

platform as a service (PaaS) *See* hosted platforms.

Point-to-Point Protocol (PPP) A remote access protocol that supports many protocols and is used to carry data over a variety of network connections.

Primary Rate Interface (PRI) An ISDN format that consists of 23 64-Kbps B-channels and one 64-Kbps D-channel. *See also* Integrated Services Digital Network (ISDN).

Routing and Remote Access Service (RRAS) A software component included in Windows Server 2008 that provides remote access through dial-up and VPN connections as well as routing and packet filtering.

software as a service (SaaS) *See* hosted applications.

switched virtual circuits (SVCs) A communication circuit that's established when needed and then terminated when the transmission is completed. *See also* virtual circuit.

Synchronous Optical Network (SONET) A flexible, highly fault-tolerant technology that can carry signals of different capacities over a fiber-optic network at high speeds. It defines optical carrier (OC) levels for incrementally increasing data rates, and SONET networks can be arranged in a variety of physical topologies.

T-carrier lines Communication lines that use one pair of wires for transmitting data and another pair for receiving data. They use the TDM signaling method, making it possible to extract any number of channels for a particular purpose. *See also* time division multiplexing (TDM).

time division multiplexing (TDM) A signaling method that allocates a time slot for each channel, making it possible to transmit multiple streams, or channels, of data on a single physical medium.

virtual circuit A logical connection created between two devices in a shared network, with bandwidth allocated for a specific transmission pathway through the network.

X.25 A packet-switching technology that provides an interface between public packet-switching networks and their customers; it has the advantage of running effectively over older copper phone lines. X.25 networks are SVC networks, meaning they create the best available pathway at the time of transmission. *See also* switched virtual circuits (SVCs).

Review Questions

1. In which of the following areas does a WAN differ from an internetwork? (Choose all that apply.)
 - a. WANs use service providers for the network connection.
 - b. WANs can't transport Network-layer protocols.
 - c. WANs use serial communication technologies that can span miles.
 - d. WANs don't use routers.
2. Which of the following is a device used to make WAN connections? (Choose all that apply.)
 - a. 10BaseT hub
 - b. CSU/DSU
 - c. Router
 - d. Ethernet repeater
3. Which of the following best describes a digital signal?
 - a. A signal that varies over time continually and smoothly
 - b. A signal whose states vary much like a lamp controlled by a dimmer switch
 - c. A channel service unit
 - d. A series of binary values
4. For what purpose is a CSU/DSU used?
 - a. Modulates a digital signal into an analog signal
 - b. Creates a digital connection between a LAN device and the WAN link
 - c. Routes packets from the LAN to the WAN
 - d. Creates a WAN connection over the public switched telephone network
5. Which of the following is a common WAN connection method? (Choose all that apply.)
 - a. Circuit switched
 - b. Packet leased
 - c. VPN over POTS
 - d. Packet switched
6. Which of the following is true about ISDN?
 - a. It uses a modem to modulate/demodulate the signal.
 - b. Its BRI format consists of two B-channels and one D-channel.
 - c. Its PRI format provides bandwidth up to 128 Kbps.
 - d. It uses a terminal adapter to connect to the network.
7. Which of the following is a typical situation in which leased lines should be considered? (Choose all that apply.)
 - a. Occasional use of the WAN link is needed.
 - b. 24/7 access is required.

- c. Fast upstream and downstream communications are needed.
 - d. You want to pay for only the bandwidth you use.
8. Which of the following combines several communication streams into a single faster communication stream?
- a. Multiplexing
 - b. Demultiplexing
 - c. CSU
 - d. DSU
9. Which physical topology does SONET readily support? (Choose all that apply.)
- a. Point-to-point
 - b. Star
 - c. Ring
 - d. Bus
10. Which of the following technologies uses packet-switching? (Choose all that apply.)
- a. ISDN
 - b. Frame relay
 - c. T1
 - d. X.25
11. Which of the following technologies provides permanent virtual circuits? (Choose all that apply.)
- a. PSTN
 - b. Frame relay
 - c. X.25
 - d. ISDN
12. Which technology uses 53-byte cells?
- a. POTS
 - b. ATM
 - c. Frame relay
 - d. X.25
13. To maintain security, WAN connections over the Internet should use which of the following?
- a. BRI
 - b. LANE
 - c. OC-3
 - d. VPN
14. What type of device is required to connect to a dedicated digital communication line?
- a. modem
 - b. NIC

- c. CSU/DSU
 - d. LANE
15. Which term best describes the place in a WAN connection where the customer's responsibility ends and the provider's responsibility begins?
- a. Data circuit-terminating equipment
 - b. Demarcation point
 - c. CPE
 - d. Central office
16. Which of the following places data on the local loop?
- a. DCE
 - b. DTE
 - c. router
 - d. CPE
17. Which of the following is the equipment at the customer site that's the responsibility of the customer?
- a. DTE
 - b. DCE
 - c. Demarcation point
 - d. CPE
18. Which of the following is a VPN protocol supported by RRAS? (Choose all that apply.)
- a. PPTP
 - b. L2TP
 - c. HTTP
 - d. SSTP
19. Which of the following is a component of the PPP protocol? (Choose all that apply.)
- a. Link Control Protocol
 - b. NCP
 - c. VPN
 - d. RDP
20. Which of the following can best be described as developing applications by using a service provider's development tools and infrastructure?
- a. Hosted applications
 - b. Hosted networking
 - c. Hosted platforms
 - d. Hosted infrastructure

Challenge Labs



Challenge Lab 12-1: Configuring an RRAS Server to Accept VPN Connections

Time Required: 1 hour or more

Objective: Configure a Windows Server 2008 server to accept VPN connections.

Required Tools/Equipment: A Windows client computer, such as Windows 7, and a Windows Server 2008 server

Description: In this lab, which can be done in groups, you configure an RRAS server in Windows Server 2008. To learn more about RRAS, go to <http://technet.microsoft.com/en-us/network/bb545655.aspx>. Configure the server to assign IP addresses dynamically to VPN clients. For all other settings, use the RRAS defaults. Then configure a Windows client to connect to the VPN server. Verify the connection by viewing connection details with the `ipconfig` command.

Case Projects



Case Project 12-1

As the network administrator for a growing company, you're asked to solve a remote access dilemma. The 12 employees who work from their homes complain about not being connected to the network except by e-mail. The company also has several employees who travel and would benefit from remote access connections. The director of marketing is responsible for part of the cost and wants only the best solution. Currently, you run a Windows Server 2008 network, and users want access to all systems. Develop a plan to connect your remote users. Your solution can involve more than one remote access type.

Case Project 12-2

CNT-Books wants an affordable way to establish remote connections for its salespeople, who log on from customer sites all over the country, and its three branch offices. The company's main office is in Phoenix, AZ, and its branch offices are in Los Angeles, CA, Chicago, IL, and Orlando, FL. Explain what kind of connections the salespeople and branch offices should use and what kinds of services should be installed on the main office's network to keep communication costs to a minimum.



chapter

13

Troubleshooting and Support

After reading this chapter and completing the exercises, you will be able to:

- Describe the benefits of documenting a network and list what elements should be documented
- Explain different approaches to network troubleshooting
- List the steps of the problem-solving process
- Explain how to make use of problem-solving resources
- Describe network troubleshooting tools
- Summarize common trouble situations
- Describe disaster recovery procedures, including backup plans and system recovery tools

The daunting task of supporting a complex internetwork can be both challenging and rewarding. Successful network administrators need to draw on a variety of skills, technologies, and techniques to meet the increasing demands of LANs and internetworks. The variety of technologies used in networks demand support personnel who are willing to continue learning long after their formal education is over. Having an excellent grasp of the fundamentals of network hardware and software enables you to adapt to the constantly changing computing environment you'll no doubt face.

In this chapter, you use the knowledge and skills you have gained throughout this book. Experience shows that a well-documented network is easier to troubleshoot and support. This chapter covers the importance of documentation and what should be documented in a typical network.

There are many different approaches to troubleshooting, and different situations call for different approaches. This chapter explains a variety of approaches to solving problems with examples of when a particular approach is and isn't recommended. No matter what approach you take, the problem-solving process has several steps. This chapter outlines eight steps; in some situations, you might use only four or five, but in others, all eight steps are likely to be needed.

A variety of problem-solving resources are available, some obvious and some not. This chapter discusses some of the most important resources for researching a problem. Your knowledge and experience might be the best resources you have and the best tools for solving problems. However, you often need software or hardware tools to help you gather information about a network and its devices. This chapter covers the most common tools for network monitoring and analysis that help you gather information and concludes with a discussion of disaster recovery, which includes backup and system repair and recovery.

Keep in mind that network support and troubleshooting is an enormous, complex topic, and entire books are written about it. This chapter is intended to give you an idea of where to start when facing a problem and some resources you can use as you gain experience as an IT professional.

Documenting Your Network

Believe it or not, there are network administrators who believe that failing to document the network means job security. It makes sense when you think of it from a narrow point of view. If your servers, cable plant, and internetworking devices are a mystery to everyone except you, it follows that you're indispensable. However, nobody can keep the details of a multitude of cable terminations, server configurations, and network devices straight in their head.

For example, the boss comes to you about the latest security vulnerability found in one of the OSs used in your network and asks whether the company's servers are vulnerable. He has read that servers without patch x using version y of the NIC driver are vulnerable. An administrator with up-to-date documentation can find the right documentation page quickly and respond reassuringly about the servers' status, thereby gaining the boss's confidence and trust. Administrators who believe documentation is unnecessary or too time consuming don't have this information at their fingertips and might respond with a confidence-draining "I don't know."

As another example, your boss tells you that 10 new offices on the third floor are going to be occupied by the end of the month. He asks what the cost in new switches, patch panels, and

cable runs will be. A well-prepared administrator can pull up a database or spreadsheet with information for the wiring closet on the third floor and promptly tell the boss how many ports are available on existing switches and patch panels. Perhaps the answer is that no extra equipment is required or only a new switch is needed, but your ability to respond quickly won't go unnoticed. Meanwhile, a documentation slacker tells the boss that some research is necessary before coming up with an answer.

There are many examples of why documenting your network is worth the trouble. The following is a list of some reasons network documentation is good for you *and* your network:

- Makes equipment and workstation moves, additions, and changes easier
- Provides information needed for troubleshooting
- Offers justification for more staff or equipment
- Helps determine compliance with standards
- Supplies proof that your installations meet hardware or software requirements
- Reduces training requirements
- Facilitates security management
- Improves compliance with software licensing agreements

The following sections describe many areas in which documentation can affect your network and the advantages of complete documentation, including how it can positively affect network changes, troubleshooting, IT staffing and training, standards compliance issues, technical support, and network security.

Documentation and Network Changes

In a network, a change is some procedure that requires modifications at workstations, such as changes of addresses, a NIC replacement, a software change, or a complete change in the workstation. Moves, additions, and changes in your network are much easier with accurate documentation. When a workstation is moved, the person doing the moving must know which patch panel and switch ports are being used so that they can be disconnected. After the workstation is at the new location, the mover must know which patch panel port to use and which switch port is correct for the workstation. Without good documentation, cables must be traced, questions must be asked, and time is wasted. Documenting the current configuration makes most changes proceed more smoothly, and documenting the change results facilitates future dealings with workstations.

Additions to your network can be done more quickly and with fewer chances of error if documentation is up to date. As noted, documentation of patch panel and switch port use, wiring diagrams, and the like can make estimating costs and scheduling more accurate and less time consuming.

Documentation and Troubleshooting

One of the first steps in troubleshooting is gathering information. If a user has connectivity problems, your network documentation can supply a wealth of information almost instantly. Physical and logical addressing, connectivity to devices, and even data about cabling can be useful pieces of information when trying to solve a problem. Accurate documentation of workstation MAC addresses helps you find issues such as IP address conflicts and the source of invalid or excessive frames.

In addition, when troubleshooting sessions requiring protocol analysis, documentation provides the names and addresses of devices that are likely to be involved in a particular type of packet exchange you're troubleshooting. Documentation also makes it easier to set up test networks so that you can duplicate conditions in your production environment to solve a complex problem.

Documentation and IT Staffing

Is network support running you and your staff ragged? Documenting the type and frequency of support calls can provide the justification for additions to staff or, at the very least, for more tools to make support more efficient. In addition, you can use statistics on network response time and bandwidth load as justification for upgrading servers or adding a switch.

Speaking of staff, the first thing you should hand a new network technician is a copy of the network documentation manual. Tell the new employee to read and learn it, and you'll have a trained technician (or at least a good start). A technician who knows where employees are located and which wiring closet their workstation cabling runs to can work more autonomously and confidently than one who must ask questions at every step.

Documentation and Standards Compliance

Compliance with standards is a necessity in today's standards-based networks for ensuring correct network operation and reducing the possibility of installation or configuration errors. For example, an Ethernet network has Cat 5e cable installed throughout. A network administrator observing a technician punch down an RJ-45 jack noticed that he punched down the green wires in the orange slots and the orange wires in the green slots. When the administrator pointed out this mistake, the technician explained that it was the only way to get the cabling to work, but he didn't know why he had to swap the orange and green wires. After a little investigation, it was discovered that the patch panels were wired according to the 568A wiring standard and the jacks were wired according to 568B! The technician didn't realize that two wiring standards existed, but if a network manual had been available that explained which standard was used at the patch panel (and the technician had been required to read it), the correct jacks could have been ordered to match the patch panels.

Documentation and Technical Support

If you call technical support to solve a network device problem, one of the first things the device manufacturer checks is whether your equipment, power supply, and cabling meet all applicable standards. For example, if you can't tell the manufacturer of a 100 Mbps switch that your cable installation passed the cable tests and possibly even provide the test results, technical support might just tell you to call back after you have confirmed that your cabling isn't the problem. In addition, if the new database server you installed is crashing, be prepared to tell the database vendor details about the server hardware, OS version, and patch installations. If you can't supply this information without walking over to the server and inspecting it, the database vendor will tell you to call back, which means another wait on hold in the support queue.

Documentation and Network Security

Physical and software security of devices, OSs with the latest security patches, and up-to-date virus protection are some factors in maintaining security. Documenting these items helps you

adhere to security policies and provides confirmation of your network's resistance to current threats or warns you of a vulnerability to these threats.

Several years ago, the Code Red virus unleashed on the Internet affected certain Microsoft IIS Web servers. When some administrators downloaded the free patch to install on infected servers, they discovered it required a newer service pack, which wasn't installed. On a network with dozens or hundreds of servers, knowing which servers had the service pack and patch installed and which servers didn't was critical. If this information is at hand, technicians can be given strict instructions on which patches and service packs should be installed on which servers. As you can see, a lot can be gained from good network documentation. The hardest part is determining what to document, and then establishing procedures and gathering tools to make documentation easier.

What Should Be Documented?

“What should be documented?” is the first question you should address before starting a documentation project or defining documentation policies. However, the answer isn't always straightforward. Networks of different size and complexity and with differing security and use policies often have different documentation needs.

This section discusses the elements of your network you should typically document. Keep in mind that the following list isn't exhaustive, and some environments might have other requirements:

- *Description of the network*—Your network description should be the section of your documentation manual that anyone could read and get a basic understanding of how your network works and what it consists of. This section should include information on the network topology, network technologies in use, OSs installed, and number of devices and users. This section should also provide contact information for the people responsible for various aspects of the network. You might also want to include key vendors and their contact information. This section is meant to be an overview document, with the details described in other sections, so keep tables and graphs to a minimum.
- *Cable plant*—This section, which will probably become worn from frequent use, describes the physical layout of network cabling, the terminations used, and the conventions for labeling cables and connectivity equipment. It also includes the results of tests done on the cable plant. Any time moves, additions, or changes are made, this section is usually consulted and possibly modified. Therefore, it must be kept current. Incorrect documentation of the cable plant is actually worse than none at all. Imagine moving a user workstation and unplugging the old patch cable from the switch only to find that you unplugged a critical server because the documentation was outdated. If you do this a few times, neither you nor anybody else will trust the documentation, and you'll go back to tracing cables every time a change must be made. In this case, you might as well discard the documentation.
- *Equipment rooms and telecommunication closets*—Equipment rooms and TCs house internetworking devices and servers and are the junction points for your work area and backbone cabling. An equipment room is also sometimes referred to as a computer room or an intermediate or main distribution facility. These rooms are often dedicated to network cabling and equipment but can also be shared with phone

equipment. Selecting features and locations of equipment rooms is important to maintaining an efficient and reliable network infrastructure. In addition, documenting the items in each room and their location is crucial to performing fast, effective changes or troubleshooting.

- *Internetworking devices*—Network changes and troubleshooting are made much easier by thorough documentation of internetworking devices. You need to know what devices are connected to other devices, the capabilities and limitations of each device, network management features available on each device, port use, and physical and logical addresses. Model numbers as well as hardware and software revision numbers might also be important when troubleshooting or considering upgrades. When you're finished with this documentation, you should be able to point to a switch, describe its capabilities, state which software version it's running, give the physical and logical addresses assigned to the switch, determine what other internetworking devices it's connected to, and list the critical resources attached to it. If this information isn't readily available, your documentation work isn't finished.
- *Servers*—All computers that provide shared resources or network services (including file and print servers, Web servers, DNS and DHCP servers, and other resources or services the company depends on) must have detailed documentation. Hardware configuration, OS and application version numbers, NIC information, and serial and model numbers are just a few items that should be available to you at a glance. Remember that a server isn't a static piece of equipment, so you can't simply install it and forget it. The interaction with internetworking devices, workstation client software, and new applications requires server hardware and software that's compatible and up to date. Installing a new device or application often requires a particular OS version or service pack, so you need to know how your server is configured before proceeding with a hardware or software upgrade.
- *Workstations*—Documentation of workstations is often the most difficult to maintain because there are usually a lot of them, and users, rather than the network administrator, are in control of them. However, don't let these factors prevent you from keeping accurate workstation documentation. Workstations, and the people who use them, are likely to be the source of most support and troubleshooting events. Knowing a workstation's hardware and software configuration and its physical and logical addresses can save you a lot of time and effort when solving a problem. Furthermore, network policies should limit how much users can change their workstations without the administrator's knowledge.

When you enforce these documentation policies, your records won't become hopelessly out of date. Maintaining accurate records on 20 or 30 workstations can be a major task, but maintaining records on thousands of workstations might seem too daunting a task to even attempt. Fortunately, many applications are available to help automate the process so that much of the documentation can be gathered only periodically and even done remotely.



For useful information on network documentation and tips on documenting your network, visit www.networkdocumentation.com. You can find a variety of mostly free tools for gathering information and diagramming networks at www.filetransit.com/files.php?name=Network_Documentation.

Approaches to Network Troubleshooting

Tackling different problems requires different approaches. Sometimes it makes sense to just try a solution and see whether it works. Sometimes you can use a similar system as a working model, or you might have to buckle down and research the problem thoroughly. In this section, you learn about different methods and circumstances in which some methods work and others don't. With this knowledge, you can try a variety of approaches for your environment.

Trial and Error

The trial-and-error approach to network problem solving isn't very scientific, and technical purists often frown on it. Nevertheless, few network specialists can deny having used it in everyday practice. There's a time and place for it, however, and you shouldn't rely on this method exclusively because you can do more harm than good in some situations.

As the name suggests, the trial-and-error method requires an assessment of the problem, an educated guess of the solution, an implementation of the solution, and a test of the results. You repeat the process until the problem is solved. This approach can be used under the following conditions:

- The system is being newly configured, so no data can be lost.
- The system isn't attached to a live network, so no other users are affected by changes.
- You can undo changes easily.
- Other approaches would take more time than a few trial-and-error attempts.
- There are few possible causes of the problem, which makes your educated guess of the solution a good bet.
- No documentation and other resources are available to draw on for arriving at a solution more scientifically.

As mentioned, it's not always wise to just try something and see whether it works. Changes made to one system on a network can affect other systems or make an existing problem worse. The trial-and-error method isn't advisable under these conditions:

- A server or internetworking device is live on the network.
- The problem is being discussed over the phone and you're instructing an untrained user.
- You aren't sure of the consequences of the solutions you propose.
- You have no sure way to undo the changes after they're made.
- Other approaches will take about the same amount of time as the trial-and-error approach.

If you determine that trial and error is the right approach for your problem, however, you should follow some guidelines:

- Make only one change at a time before testing the results. That way, if the problem is solved, you know which change is the solution. You can add this information to your network support documentation for future use.
- Avoid making changes that might affect the operation of a live network. For example, if you suspect an incorrect TCP/IP address, don't change the address without first

verifying that the new address is available. Using an address that's already in use could cause another device to stop working.

- Document the original settings of hardware and software before making changes so that you can put the system back to its original state.
- Avoid making a change that can destroy user data unless a recent backup exists.
- If possible, avoid making a change that you can't undo.

The following examples help you determine under what circumstances this troubleshooting method is suitable.

You're called to solve a problem on a client/server network of about 100 computers, and employees access the Internet. The problem is that a workstation running Windows 7 can't access the Internet. You sit down at the workstation and open Control Panel to check the settings. You check the IP address settings. TCP/IP is configured to use DHCP. You know that DHCP isn't used on this network, so it must be the problem. You recall from an earlier visit that the network address is 206.17.44.0. You decide to configure the computer with an address that you select from the network randomly to see whether this step solves the problem.

Should you use trial and error to solve this problem? Absolutely not. Although you might have happened on the cause of the problem, simply choosing an address without knowing whether it's already in use can cause a conflict with another machine. The correct course of action is to consult the network documentation that lists all IP addresses in use. This document should also have other settings, such as DNS server addresses and the default gateway address. You must consider the effect that changes you make have on the rest of the network. If you're unsure, play it safe and consult the documentation.

In the second example, you've been asked to troubleshoot a new PC running Windows Vista on a client's network. The network is small, has only seven PCs, has Internet access through a cable modem and router, and is set up in a workgroup environment. An employee has already done some of the work, such as assigning a computer and workgroup name, but the new PC still can't communicate with other PCs on the network or access the Internet. You find that a static address is assigned to the computer, but you know addresses are supposed to be assigned via DHCP by the router. You decide to configure TCP/IP to get an address via DHCP.

Would trial and error be a safe, effective troubleshooting method in this situation? If you said yes, that's correct. Configuring TCP/IP to use DHCP is a reversible action and is safe because there's no chance of conflicting with another station. If it works, you have solved the problem. If not, no harm done. The only caveat is that you should note the static IP address settings in case the machine is supposed to have a static address for a reason you don't yet know.

In the third example, you get a call from a client having intermittent problems with a subnet. The client tells you that when employees try to access a server on a different subnet, sometimes it works, but sometimes the connection times out. The network has four subnets connected through one router. There are no problems on the other subnets. You have seen a similar problem at other networks, and resetting the router seemed to solve the problem. You tell the client to power down the router, wait 10 seconds, and power it back up.

Is this approach a reasonable way to solve this problem? By powering down the router, you affect all four subnets. This action could cause loss of data, time, and possibly even money. Additionally, you don't know whether the router configuration has been saved, so powering

down the router could cause even worse problems after it restarts. Finally, you should never instruct a user to perform a procedure when you have no way to make sure it's being done correctly.

Sometimes using trial and error to solve problems is quicker and easier than other methods. In fact, this option might be your only way of solving some problems in a timely manner. However, you must be careful not to make matters worse if your proposed solution can affect other systems or cause data loss.

Solve by Example

Solving by example is the process of comparing something that doesn't work with something that does, and then making modifications to the nonfunctioning item until it performs like the functioning one. It's one of the easiest and fastest ways to solve a problem because it requires no special knowledge or problem-solving skills. When most organizations purchase new computers, they purchase similar models and configure them identically. You simply take advantage of this fact when confronted with a problem on a machine.

Some problems can be difficult to troubleshoot, particularly when they involve an OS configuration. In addition, hunting down the problem and fixing it could take considerable time. If you have a working example of a device that's nearly identical, however, you can copy the configuration, or parts of the configuration, from a working machine. This effort might involve checking Control Panel for installed components, copying system files such as device drivers, copying configuration files, or even making a copy of an entire disk.

To see how this process works, take a look at an example. Mike, a networking consultant, has been called into a client's office because 2 of the firm's 20 computers lock up periodically when accessing a Web-based database program on a Windows server. Mike checks Control Panel on the two computers, and everything seems to be in order. The two computers can access the network, but when they run Internet Explorer to access the database application, the machines invariably lock up within a half hour of use. Mike asks whether other computers run the same Web application and is told yes. He takes a look at one of the computers that's not locking up. After some investigation, Mike finds that the working computers are using the 64-bit version of Internet Explorer, and the two computers that lock up are using the 32-bit version. Mike uses the 64-bit version of Internet Explorer to run the Web application and after several hours of testing, no lockups occur, and Mike calls it a day.

What's your analysis of this situation? Mike certainly could have pulled out his network analyzer and captured network packets to try to determine what the problem was, or he could have played trial and error with different network settings to see whether he could correct the lockups. However, with several working examples nearby and one obvious difference between the working machines and the faulty ones, he did the smart thing by taking what works and applying it to what didn't work.

As with the trial-and-error method, there are some caveats to using the solve-by-example method. Here are some general rules to follow:

- Use the solve-by-example approach only when the working sample has a similar environment as the problem machine. For example, don't compare a machine having problems accessing Windows Server 2008 with one accessing Linux.



- Don't make configuration changes that will cause conflicts. For example, don't change the TCP/IP address of a nonworking machine to the same address as a working machine's.
- Don't make any changes that could destroy data that can't be restored.

In another example, Sophie is fairly new to networking but has been asked to connect some new computers to a stack of two switches. A similar stack of switches already exists, and she's supposed to make these connections in a similar fashion. Armed with a box of patch cables, Sophie starts plugging computers into the switches. When she gets to the last switch port, she realizes she must save this port to connect to the next switch. She connects the two switches with a patch cable but doesn't see a link light indicating that the connection is good. She tries another patch cable with the same result.

Not sure what to try next, Sophie examines the similar stack of switches and recognizes a button next to one of the ports. It's a two-position button with the positions marked as Normal or Uplink. The button is set to the Uplink position on one switch and Normal on the switch it's connected to. Because the switches she's setting up are the same, Sophie compares her switches with the working switches. The button is set to the Normal position on both her switches, so she changes it to the Uplink position on one switch. The link light indicator comes on, and Sophie finishes her job.

Is the solve-by-example method appropriate in this situation? Because Sophie had an example in an environment similar to the one she was having trouble with, she was able to make these changes with confidence. The switches weren't being used on the live network yet, so her changes wouldn't cause any problems.

The Replacement Method

The replacement method of problem solving is a favorite among PC technicians. It requires narrowing down possible sources of the problem and having known working replacement parts on hand so that they can be swapped out. Sounds simple, and it is—at least after the source of the problem has been identified. That's where the difficulty and the skill come in. The replacement method is effective only if the problem's source can be determined and the source is a defective part. A lot of time and money can be wasted in replacing parts that aren't defective, so you need to apply your troubleshooting skills before you show off your installation skills. Follow these rules in order when using the replacement method:

1. Narrow the list of potentially defective parts down to a few possibilities.
2. Make sure you have the correct replacement parts on hand.
3. Replace only one part at a time.
4. If your first replacement doesn't fix the problem, reinstall the original part before replacing another part.

Step by Step with the OSI Model

The step-by-step method of troubleshooting involves using the OSI model. In this approach, you test a problem starting at the Application layer and keep testing at each layer until you have a successful test or reach the Physical layer. Alternatively, depending on the problem, you might start at the Physical layer and work your way up the OSI model. This method of problem solving is what most people think of as network support. To use this approach, you must understand how networks work and where you should use troubleshooting tools.

Networks are complex, multilayered systems. When confronted by a problem for which there's no obvious fix, remembering the layered approach to network systems can be helpful. If you conceptualize the problem following the seven layers of the OSI model, you can take a step-by-step approach to solving the problem. To see how this approach works, start by reviewing the simple network diagram in Figure 13-1.

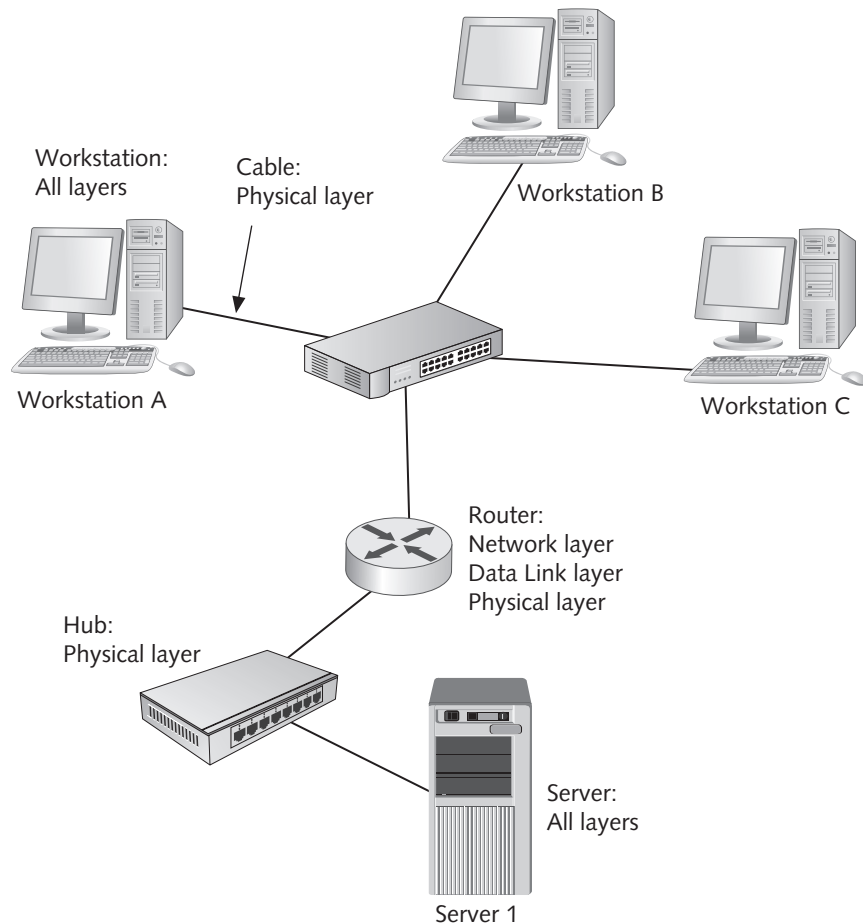


Figure 13-1 Troubleshooting with the OSI model

Courtesy of Course Technology/Cengage Learning

Suppose the user on Workstation A complains that an error occurs when she tries to access files on Server 1. Users at Workstations B and C aren't having similar problems. No more information is available. When you arrive on the scene, you see that Workstation A is running Windows XP. First, you open Network Places to view network resources. As expected, no resources are available. This step involves the upper layers of the OSI model. Now that you have determined these functions aren't working correctly, you can start looking at the lower layers. Your goal is to find the lowest layer at which there's functionality.

You check the network documentation and see that TCP/IP is the protocol being used. You then verify that TCP/IP is installed correctly on this computer, so you know you have functionality in

Layers 3 and 4. Time to keep moving down the OSI model. A common tool for troubleshooting TCP/IP is the Ping program. Referring to the network documentation, you try to ping Server 1. No success. Look at the diagram in Figure 13-1 again. If you can communicate with the router but not the server, you can narrow your search. Try pinging the router. Again, no success.

Which layers remain to be tested? The Data Link and Physical layers. Data Link–layer problems most often affect the entire network, or the problem lies with a single computer’s NIC or drivers. Unless you have reason to believe there’s a problem with the drivers, it’s best to leave them for later. Most network technicians would move on to the Physical layer because this layer is where problems restricted to one workstation are most likely to occur. After a brief investigation, you find that the patch cable from the jack to the workstation has gone bad, and you replace the cable. Problem solved.

Many network technicians approach a problem by starting with the Physical layer and then working their way up. The approach you take depends on your experience and information you have learned from interviewing users. What’s important is that you understand everything required for the network connection to work, which enables you to test and check all components involved with the tools available to you.

Here’s an example of starting with the Physical layer and working your way up: A user complains that she can’t access the Web site <http://books.tomsho.com>. You have no other information, and the user isn’t at her computer when you arrive. Starting with the Physical and Data Link layers, you check cable connections and link lights on the computer’s NIC. They look good, so you move to the Network layer. You check the computer’s IP address configuration and attempt to ping the default gateway. Again, you’re successful. Next, you try to ping a known address in another network to verify that the default gateway is routing correctly. You might check DNS next, which verifies the Transport and Session layers (because DNS is often described as a Session-layer protocol). To do this, you ping a Web site by name or use Nslookup to verify that DNS lookups work. You might also want to ping the address the user is having trouble with to determine whether the target server’s name can be resolved and it can communicate via pinging. If all your checks so far have been successful, checking the Application layer is next. You can use a protocol analyzer, such as Wireshark, to capture packets so that you can see exactly what’s happening with HTTP, the protocol used in Web communication. Application-layer troubleshooting is the most difficult because protocols in the upper layers are the most complex, but with experience, you’ll be able to solve problems occurring at all layers of the OSI model.

The Problem-Solving Process

One of the most difficult aspects of network problem solving is deciding where to begin. What’s described next is a general framework for approaching problems that you can apply in almost any situation. The specific actions you take depend on the situation. The process described in this section can be applied to a variety of problems, both in your networking environment and in everyday life. Here are the steps of the problem-solving process:

1. Determine the problem definition and scope.
2. Gather information.
3. Consider possible causes.
4. Devise a solution.

5. Implement the solution.
6. Test the solution.
7. Document the solution.
8. Devise preventive measures.

Several steps in this process might be repeated. For example, if Step 6 doesn't lead to a solution for the problem, you probably need to repeat Steps 2 through 6 until you do have a solution. Each step might also require several substeps (explained in the following sections) before you can move on to the next step. For example, Step 4 might require setting up a test environment to duplicate the problem and test possible solutions before implementing a solution on a live network. Figure 13-2 is a flowchart of the basic process.

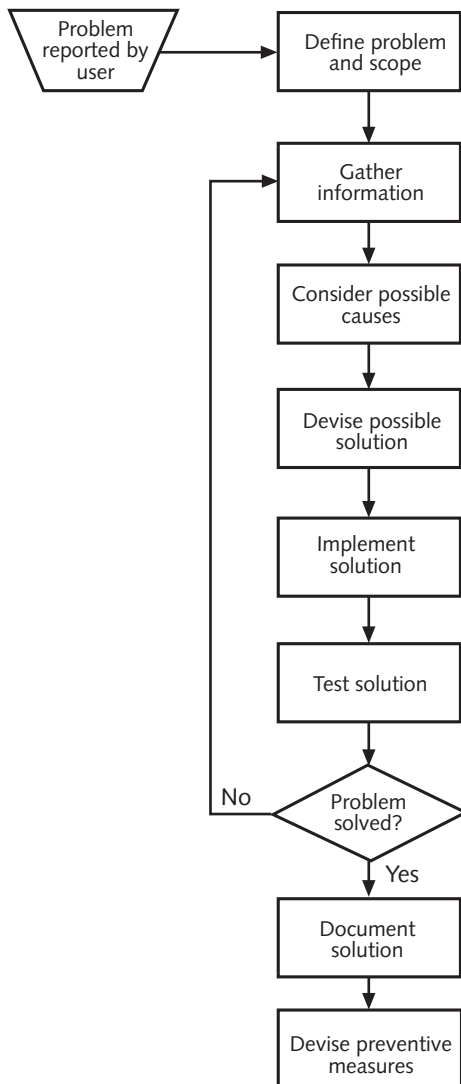


Figure 13-2 The problem-solving process

Courtesy of Course Technology/Cengage Learning

Step 1: Determine the Problem Definition and Scope

Before a problem can be solved, it must be defined. “Mary’s computer doesn’t work” does not define the problem well enough to create a plan of action. “Mary can’t run Word because an illegal operation occurs every time she tries to run it” is much better. A problem definition should also describe what does work and what doesn’t work. If Mary can’t run Word, stating whether she can run her e-mail program or other applications should be included in the problem definition.

You need to know who and what are affected by the problem. You take a vastly different approach if an entire floor rather than a single user is affected by a network problem. Is the problem related to a single application—for example, e-mail—or are all functions affected? If you’re working with routers, is only one router exhibiting problems, or are several routers affected? Determining the scope of a problem is important not only for deciding where to start your troubleshooting process, but also for deciding what priority to assign to the problem. The malfunction of a network switch or server demands a higher priority than a problem affecting one user. Determining the scope quickly and accurately, therefore, is not only part of the first step, but also an essential step of the troubleshooting process.

Most network problems come to the network administrator’s attention by way of a user phone call or e-mail. This communication is your first opportunity to learn more about the problem. Although this part of the troubleshooting process is more art than technical skill, there are some questions you can ask to start you on your way:

- Is anyone else near you having the same problem?
- What about other areas of the building?
- Is the problem occurring with all applications or just one?
- If you move to a different computer, does the problem occur there, too?

The goal of your questions is to determine a problem definition and scope. If a solution comes about as a result of this interview, all the better, but that’s not the goal of this step, nor is determining the cause of the problem. Rather, the goal is to define the problem in detail and determine the scope of the problem accurately. Examples of a problem definition and scope include the following:

- Jim can’t access the e-mail server. Other servers are available to Jim, and no one else reports the problem.
- Third-floor users can’t log on to the network, but users on all other floors can.
- Vanessa can’t print to the new LaserJet printer on the fourth floor and has tried several applications. No other users have tried to print to this printer.
- Jake reports that the network is slow while accessing his Documents folder. Access to the Internet and other resources seem to work with normal performance.

After you define the problem and understand its scope, you can assign a priority to the problem. Assigning priorities takes a little experience and some political savvy. You must have a clear understanding of what areas the organization deems most critical to its business functions. Creating a document to show management in which order problems are solved when there’s a backlog is helpful.

Most IT departments are understaffed, so backlogs are usually the norm. Besides prioritizing according to business functions, prioritizing according to who reports the problem or who it

affects is common. If the president of the company can't check e-mail, giving this problem top priority can be a double-edged sword. On one hand, if you solve the problem right away, you curry favor with the boss. On the other hand, if you solve the problem right away, you must not have much else to do, so why are you always complaining that you're understaffed? In any event, after you have prioritized the problem, you can assign the support person who's best equipped to solve the problem. After the problem is ready to be tackled, you can move to the next step: gathering information that can help you solve the problem.

Step 2: Gather Information

This step is where your user interview skills can really shine. Most of the initial information you get about a problem comes from users. Knowing what questions to ask and how to ask them can mean the difference between a quick fix and an all-nighter.

Did It Ever Work? Strangely enough, this question is often overlooked. There's a big difference between something that worked once and then stopped working and something that never worked at all. Users often don't volunteer this information, so it pays to ask. If something worked once and now doesn't, you can assume something has broken the process. If it never worked at all, there's a good chance it wasn't set up correctly in the first place. If it worked at one time, you go into troubleshooting mode and continue with the interview. If it never worked, you go into installation mode and look at it as another task to put on the to-do list.

To understand this principle, consider this example: Alexis gets a call from Matt, a network user. Matt tells Alexis that he can't print to the printer down the hall. After determining that Matt is the only one having the problem as far as he knows, Alexis goes to the information-gathering step. She asks Matt whether he could ever print to that printer, and he replies that he couldn't. Alexis can go into installation mode at this point. This problem has just become a simple printer installation that isn't really a problem at all. Had Alexis not asked that question, she might have gone into troubleshooting mode, continuing to ask questions, checking printer queues, determining printer permissions, and investing time in a host of other time-wasting activities.

When Did It Stop Working? Assuming that the problem involves a function that used to work and has stopped or changed in some way, you need to find out when the change occurred. The purpose of this question is to find out the problem's time and date of occurrence and determine what else might be going on at that time to cause the problem. For example, is another application running when the problem occurs, or does the air conditioner kick on about the time the problem occurs? This line of questioning can also give you information about its urgency. If the user has lived with the problem for two weeks and is just now reporting it, and you have bigger and hotter fires to put out, you might be able to put this one on the back burner. A good support technician must know how to listen to customers and understand their sense of urgency or frustration. You might also want to ask the following:

- Does the problem occur all the time or only intermittently?
- Are there particular times of the day when the problem occurs?
- Are other applications running when the problem occurs?

Has Anything Changed? You have to be careful with this question when you're talking to users about workstation problems. If users think you're implying they might have caused the problem, they're likely to clam up. Their answers are essential pieces of information. New applications on workstations, new hardware devices, and updates to existing applications or drivers can all cause problems. While you're asking users this question, you need to ask yourself, too. Were any changes made to the network that could cause the problem? Were any upgrades made to servers, or were new router configurations implemented?

Never Ignore the Obvious Sometimes it's easy to get caught up in a problem, pull out the network analyzer, and start some serious troubleshooting. One of the most common problems, which thankfully has one of the easiest solutions, is an unplugged cable. Don't assume that your users will have checked this possibility. Experience suggests that a substantial percentage of network problems involve an unplugged cable. Maybe the culprit was the cleaning crew, or maybe one of your technicians did it while installing a new sound card or working on a server or router.

Sometimes you can discover the obvious when you realize that people have their own unique perceptions of a problem. For instance, descriptions such as "slow network response" are subjective; what seems normal to one person might be considered a problem by another. Suppose an employee has been on the night shift for the past year and has recently taken a shift during the day. This employee reports that server response is very slow, but you have had no complaints from other users. Because the night shift works half-staffed, this employee has probably become accustomed to a server operating with a lighter load during the night, which results in quick response times. His idea of a slow response might simply be a normal response time for the day shift.

Define How It's Supposed to Work Gathering solid facts about a problem is difficult if you don't have a good definition of how things are supposed to work. Having good documentation and a clear baseline of your network pays off. Periodic baselines are compared with previous baselines to spot trends that indicate problems ahead. For example, if average network utilization increases 2% to 3% per month for several months, you can prepare for a performance upgrade that will no doubt be required before too long.

A baseline of your network should include network utilization statistics; utilization statistics on server CPUs, memory, hard drives, and other resources; and normal traffic patterns. This information can be compared with statistics you gather during the troubleshooting process. It can help you determine whether reports of slow response time are valid and point you in the direction of the problem's source if they are indeed valid. It can also help you know when it's time to upgrade the network infrastructure or servers.

Step 3: Consider Possible Causes

In this step, based on symptoms and other information you have gathered, you consider what could be the cause of the problem. Experience is invaluable in this step, as the more problems you have seen, the more likely you are to recognize symptoms of a particular problem. As you proceed through this step, you'll probably gather more information.

Your goal in this step is to create a checklist of possible things that could have gone wrong to cause the problem. For example, an entire area of a building has lost connection with the

network, but no other areas are affected. Without knowing anything else, you could construct the following list of possible causes:

- The connection in the main wiring closet to the rest of the network has failed.
- The switch all workstations are connected to has lost power or failed in some way.
- All workstations have acquired a virus through the network, and the virus affects their network connection.
- A major upgrade has been made recently on all workstations in that area, and incorrect network addresses were configured.

You could create quite a long list if you put your mind to it. Of course, during Step 2, you would probably have eliminated all but a few of the possible causes. If you find yourself with a long list of possible causes, you probably need to go back to Step 2 and gather more information. After creating a list of possible causes, you can investigate each one and rule out or confirm it. In the previous example, you would probably check the wiring closet to see the status of devices there, or if you had a network management program, you could verify the health of wiring closet devices remotely.

Step 4: Devise a Solution

After determining a likely cause, you can devise a solution. In the example discussed in Step 3, assuming the cause of the problem is a failed switch, devising a solution is easy: Replace the switch. Suppose, however, that many users have reported a periodic loss of connection to certain resources. After you've gathered information and considered possible causes, you find that several routers become overutilized periodically and start dropping packets. This problem isn't so simple. You don't want to rush in and replace the problem routers with bigger, stronger, faster routers because this solution could affect other routers or other network components. Is the problem with the routers, or are the dropped packets simply another symptom of the problem? You don't really know. Before devising a solution, it's important to consider the following:

- Is the identified cause of the problem truly the cause, or is it just another symptom of the problem's true cause?
- Is there a way to test the proposed solution adequately?
- What results should the proposed solution produce?
- What are the ramifications of the proposed solution for the rest of the network?
- Do you need additional help to answer some of these questions?

The last question is a sore point for many network professionals. However, being an expert in everything is impossible, and some network problems can be too complex or the equipment needed to answer questions is too expensive for many IT departments. A broken network that results in reduced productivity and, therefore, lost money can be more expensive than calling in experts occasionally.

After you have the solution, it's time to carry it out immediately, right? Wrong. Before you implement the solution, you must be prepared for the possibility that it could make things worse than the existing problem. Whether your problem and proposed solution affect an entire network or just a few users, you must devise a **rollback plan** so that you can return

things to their original state if the solution doesn't work. Depending on the scope of the problem and solution, you might need to do the following:

- Save all network device configuration files.
- Document and back up workstation configurations.
- Document wiring closet configurations, including device locations and patch cable connections.
- Conduct a final baseline to compare new and old results if a rollback becomes necessary.

Step 5: Implement the Solution

If you have done a good job with the first four steps, the implementation step should go fairly smoothly. During this step, you create opportunities for intermediate testing and inform users of your intentions. Then you put the plan into action.

Create Intermediate Testing Opportunities You need to design the implementation so that you can stop and test it at critical points, instead of testing the completed solution only to find that something doesn't work. Testing small steps in which a limited number of things could go wrong is far easier than testing a complex solution with dozens or hundreds of problem areas.

Suppose your solution is to add a network segment to your internetwork to alleviate broadcast problems. You have purchased a new router and a switch to accommodate the workstations that will form the new segment. One way to carry out this solution is to hook up all the equipment, configure the router and switch, assign new addresses to workstations, plug in all the cables, and then hope for the best. When this method doesn't work, however, where do you start looking? Is the problem the router configuration or the switches? Is your addressing scheme incorrect?

A better way to tackle this solution is to have a step-by-step plan that allows intermediate testing. For example, you can use the following steps that alternate between implementing and testing to test the new router and switch:

1. Configure the router.
2. Verify its stand-alone operation by pinging each interface.
3. Attach the router to the rest of the network.
4. Verify that all parts of the network can be reached by pinging.
5. Use the Trace Route program to verify the path selection.
6. Install and configure the switch.
7. Configure workstation addresses for the new network.
8. Cable workstations to the new switch.
9. Verify connectivity in the network.
10. Connect the router to the switch.
11. Verify that you can ping the router interface from workstations.
12. Verify that you can reach other networks from workstations.
13. Create a baseline of the new network segment.

A carefully planned implementation of your solution with testing along the way allows you to catch unforeseen results at a stage when they're easy to see and easy to fix.

Inform Your Users When your action plan affects other parts of the network and, therefore, other users, you need to inform users of the possible disruption to some network services while work is progressing. Give users plenty of time to schedule downtime of the network. Nothing can take the wind out of your sails more quickly than getting a frantic call from a boss who needs the network for a big presentation just as you're halfway through a day-long network upgrade.

Put the Plan into Action After you have your checklist of actions and intermediate testing ready and have informed users, it's time to take action. Provided you have done everything correctly up to now, this step is the easy part. You have your list of actions; now is the time to carry them out.



Remember that making only one testable change at a time is crucial.

TIP

Take notes about every change you make to the network or servers. For example, document a driver upgrade or an IP address change. This way, you know the network's current state when your changes are finished. A well-documented network is easier to troubleshoot and upgrade in the future.

Step 6: Test the Solution

It's 3:00 a.m. and you're finished with the upgrade. Time to go home, right? Wrong. It's time to test your solution as a whole. If the issue is a simple workstation connectivity problem, you verify that the station can access the resources assigned to it. If it's a major network or server overhaul, however, the testing is more involved. In either case, if you have done intermediate testing during the implementation step, the testing step should be fairly straightforward.

Your testing should attempt to emulate a real-world situation as closely as possible. If you're testing a workstation problem, verifying that the workstation can ping a server isn't enough. If possible, you should attempt to log on to the network as a user with similar privileges as the workstation's main user. Next, attempt to access applications that would likely run from the workstation. Take notes about what you learned and saw.

If you're testing a major network upgrade, you have probably already tested end-to-end connectivity during implementation. Now you need to put some stress on the network. Start some workstations on the upgraded part of the network, if possible, with the help of some assistants, and run some network-intensive applications. Access the Internet, if Internet access is included in your network. All the while, you should be gathering information about how the network behaves while you're working it. Compare your results with the results you saw before the changes were made. Again, take notes about the results of your testing. When you have tested everything possible, go home and get some sleep; tomorrow will be the real test, when users begin using your new solution.

Step 7: Document the Solution

If you have made it this far, congratulations—you have solved the problem! It's time to take all the handwritten notes made during the implementation and testing steps and turn them into a cohesive document. This step is as important as any of the previous steps. No matter how big or small the problem was, a similar problem will likely happen in the future. If you took notes about the problem and the solution, you have this documentation available as a valuable resource for solving the next problem of its kind. Your documentation should include everything pertinent to the problem, such as the problem definition, the solution, the implementation, and the testing. If necessary, you should be able to reproduce both the problem and the solution from your documentation. If the problem and its solution have implications for the entire network, including this information in your overall network plan is advisable.

Step 8: Devise Preventive Measures

After solving and documenting a problem, you should do everything you can to prevent this problem or similar problems from recurring. For example, if the problem was the result of a virus that spread throughout the network and caused considerable damage before it was found, you can install virus protection programs on the network and tighten policies for software and e-mail downloads. This preventive measure is obvious and reasonably easy to do.

Suppose, however, that the problem is a degenerative one, in which the network gradually becomes slower and less responsive. Preventing this problem isn't as simple as installing software and sending a policy memo. There are some measures you can take, however. For instance, you can devise certain rules for your network's operation. For example, you can specify that no more than 50 workstations be installed on a network segment, or stipulate that your Linux servers can have no more than 200 simultaneous logins before adding a server or adding a CPU to the server. These types of rules help prevent performance problems in the future. In addition, if those in charge of the budget approve these rules, you have instant justification for an upgrade when the time comes.

Devising preventive measures is proactive rather than reactive network management. If you let the problem come to you, it's always far more serious than nipping it in the bud before it causes serious productivity issues. You might be tempted to pat yourself on the back and rest on your laurels after solving a difficult problem, but coming up with methods to prevent problems in the first place is wiser.

Making Use of Problem-Solving Resources

This section covers some resources available to you during troubleshooting. Each resource has its place, and experience will tell you what's appropriate for different situations.

Experience

Your most effective weapon in supporting the network and diagnosing and solving problems is your own experience. Unfortunately, people often don't make effective use of their experience. Whether you have been limited to working on computers in the classroom and at home or have been working on a large multiplatform network, there are plenty of opportunities to expand and enhance your experience.

Make the Most of Your Experience Few people have photographic memories. They see something, say they're going to remember it for future use, and then promptly forget it. Sometimes people remember generalities but forget the details, which is easy to do with networks because so much is changing constantly.

Take notes about what you see and learn. This advice applies even if you've been in the computing world for years, but it's particularly pertinent when you're first starting out and your experiences are limited. Keep a journal of your experiences. Even if you never read it again, the act of writing information down helps preserve it in your memory for future use. Say you're upgrading a network with a VPN server. After several configuration changes, you finally get everything to work. If you write down the details of what worked and what didn't work, you have a reference for the next time you have to perform a similar installation.

**TIP**

An electronic journal is helpful because you can file your entries alphabetically and search for them when needed. Of course, a print-out is also useful when your network crashes and electronic documentation is unavailable.

If It Happened Once, It Will Happen Again One mistake technicians make is thinking that a problem is so obscure that it's not worth the time and effort to make a note of it. However, hardware and software are standardized now, and millions of people use the same or similar components in their computers and networks. So if you're seeing a problem now, you'll probably see it again. Make a note of it, and the next time the problem occurs, you can be the hero by already having the solution at hand.

Colleagues' Experience One of the most overlooked resources for solving problems is your colleagues and classmates. Use the people you know as a resource. They'll appreciate you coming to them for possible answers and, in turn, they'll come to you in the future. Some people build up a network of colleagues and put them on an e-mail distribution list. When facing a difficult question or problem, they can send an e-mail to several knowledgeable people. There's a good chance one of them has had a similar problem in the past and can steer the problem-solving process in the right direction.

Experience from Manufacturers' Technical Support Sometimes there's nothing left to do but call for help. Every time you install a new piece of hardware or a new application, one of the first things you should do (besides reading the installation manual) is enter the manufacturer's technical support number in your database of important phone numbers.

The best time to call technical support is when you have a specific error number or message that you can report to the manufacturer. Be prepared to have a lot of other information ready, too. The more prepared you are, the more responsive the support person is likely to be. Typically, information you need includes the software's version number or the hardware's serial number, the OS and version, whether it's an application problem, and, for a router or switch problem, the firmware revision number. You need to be as detailed as possible about the problem or error's circumstances so that the manufacturer can reproduce it if needed. Gather all pertinent information before you call technical support; if you don't have the necessary information, you'll have to call back a second time.

In addition, use some of the troubleshooting methods discussed earlier to rule out obvious problems, such as a defective part. If you have another part handy, use the replacement method so that you can tell technical support you have already tried swapping parts. You can also try the suspect part or application on a different system so that you can report this information to technical support. Again, the more prepared you are, the better results you'll get. In addition, if you've tried all the obvious troubleshooting techniques and can report this fact to technical support, it's likely that your problem will be transferred to a more knowledgeable person or tech support will be prepared to send a replacement part.

The World Wide Web

If you can describe the problem with a few words or an error message or number, the Web is the first place to look for answers. Most manufacturers put time and effort into building databases of problems and solutions so that their customers can research the problem themselves without calling the technical support line.

The Web is one of the best resources for computer and networking professionals. What used to take days or weeks to accomplish via phone calls and driver updates on floppy disks sent by mail can be done in minutes by using the Web. You can't install a new NIC on the new version of Windows you just installed? Get on the Web and download the latest driver. Every time you try to send an e-mail, you get error number 3744? Go to the software developer's Web site and enter the error number in a search, and you might get a response explaining how to solve the problem.

Most manufacturers store their technical support problems and solutions in a database called a knowledge base or a frequently asked questions document. A knowledge base is a searchable database containing descriptions of problems and errors along with known solutions, if any. It can also contain installation notes and compatibility information. A **frequently asked questions (FAQ)** document is usually a text document with two parts to each entry. The first part is a question the manufacturer has anticipated or actually received from customers; the second part is an answer to the question. A FAQ is more helpful for general installation and configuration help, although it can have information about error messages, solutions, and compatibility issues.

Using a Knowledge Base or Search Engine The old adage of “garbage in, garbage out” applies perfectly to using a knowledge base or search engine. With a search engine, you have to enter the right words, phrases, or error numbers to find the information you want. Even then, finding what you're looking for can take several attempts, and you might have to sift through several entries before you find the information that will help with your particular problem.

When you're researching a problem, be as specific as possible. If you have error numbers or messages, enter them. With error messages, you get the best results if you enclose them in quotation marks. For example, if the error message says “Too many open files,” enter this exact phrase enclosed by quotation marks to get the best and fewest search results. Enter as many keywords or phrases as possible to limit the number of results returned; you can get hundreds or thousands of results if the keywords you enter are too general. If your first search returns no results, cut back on the specificity of the search and try again. After a while, you'll get a feel for the type and amount of information to enter.

Finding Drivers and Updates When installing a new piece of hardware, OS, or networking device, one of the first things you should do is check whether bug fixes, driver updates, or new firmware revisions are available. Before you call a manufacturer's technical support line, make sure you have the latest versions, or the support person will probably tell you to call back if the problem persists after you have installed the new version.

Most manufacturers devote a section of their Web sites to the latest fixes and drivers you can download. A word of caution: Read the installation guide or Readme.txt file before installing OS updates because you might need to be aware of special preinstallation items before you start the update.

**TIP**

Many drivers can be updated directly in Windows. In Device Manager, right-click the device and click Update Driver Software. Windows attempts to find the latest driver available and install it or informs you that the driver is up to date.

Consulting Online Support Services and Newsgroups Many online support services are dedicated to technical subjects such as networking. You can use these services to tap into the knowledge of experienced network professionals by posting questions. One excellent source is Experts Exchange (www.experts-exchange.com). This Web site is a subscription pay service, but you can earn points toward your subscription by answering questions posted by members. In addition, many companies use user communities to their best advantage by creating newsgroups or support forums that users of their products can go to exchange experiences and help one another.

Researching Online Periodicals Given rapid industry change, periodicals dealing specifically with computers and networking can be the best sources of information on new products, trends, and techniques. Many periodicals are available on the Internet, and some offer free subscriptions to networking professionals. Some of the most popular networking journals include *Network Computing*, *Information Week*, and *Network World*. Several publications focus on Windows (such as *Windows IT Pro Magazine*; www.windowsitpro.com) or Linux (for example, *Linux Journal*; www.linuxjournal.com).

Network Documentation

Many network administrators dislike the task of network documentation, but it's one of the best resources for knowing what's happening with a network and what needs to be done to fix a problem. Good network documentation can mean the difference between a five-minute fix and hours, or even days, of troubleshooting.

As mentioned, you should document everything that's important to installing, maintaining, and troubleshooting the network. Your documentation should read like a user's manual for network administrators. You know it's complete when you feel as though you could leave your network in a stranger's hands for a month, and everything would still be working fine when you come back. Good documentation should include information in at least two categories: network topology and internetworking devices. This guideline is general; your network might have many more categories and subcategories. If your documentation is weak in either area, you should set aside time to improve it.

Network Diagrams A picture is worth a thousand words, and this statement is certainly true for a network. Your documentation should include network diagrams showing a logical picture of the network and another diagram showing the network’s physical aspects, such as rooms, devices, and connections. Complete documentation shows a level of detail down to the floor plan and location of jacks. Figure 13-3 shows a logical diagram of a network, and Figure 13-4 shows a physical diagram.




Figure 13-3 A logical network diagram

Courtesy of Course Technology/Cengage Learning

Internetworking Devices Internetworking devices require different levels of documentation, depending on the type of equipment. Simple, unmanaged switches require the least information, for example, whereas routers normally require the most. Besides depicting internetworking devices in your network diagrams, you should list them in table form, as shown in Table 13-1’s example of a list of managed switches. A similar table should be created for all types of devices so that they can be located and identified easily when needed. This information also helps with expansion plans because it includes the number of free ports available, where you can add workstations and other devices.

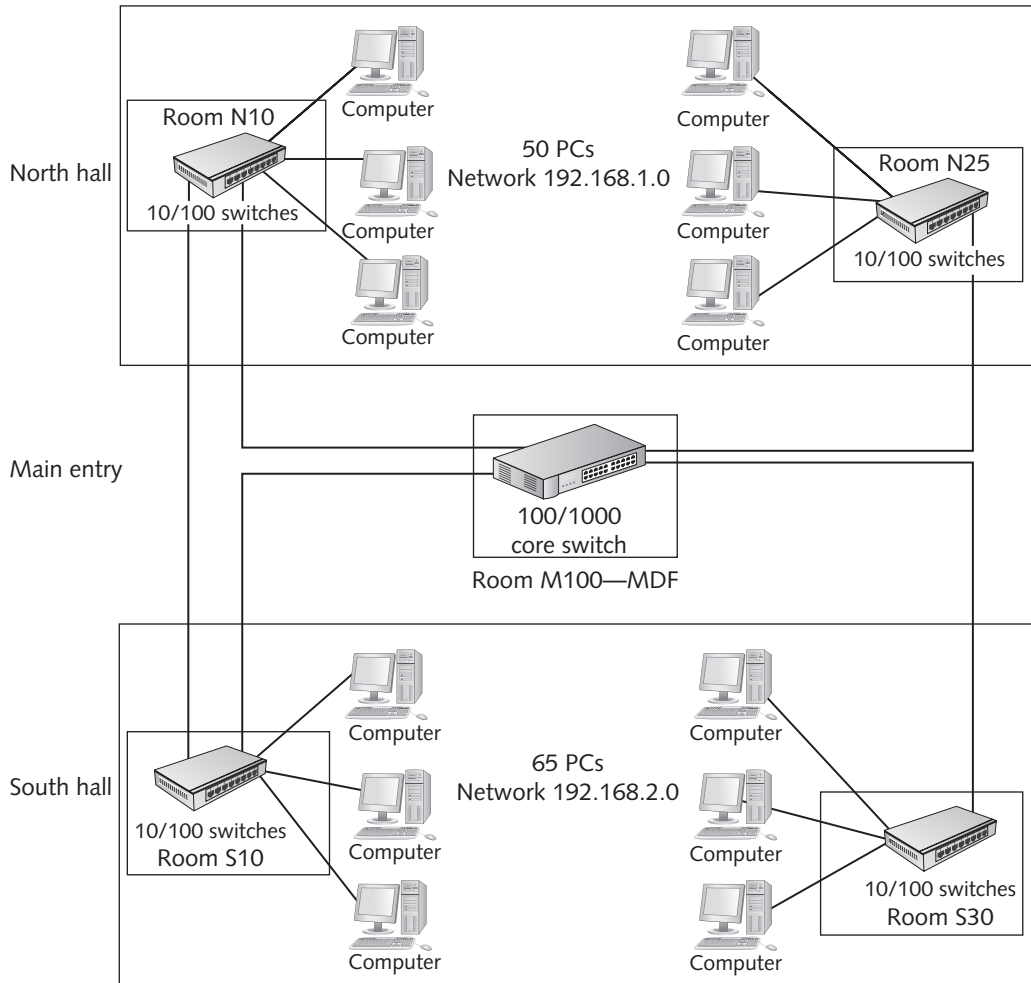


Figure 13-4 A physical network diagram

Courtesy of Course Technology/Cengage Learning



Table 13-1 Network equipment list: switches

Switch model/serial #	Location	IP address	MAC address	# Ports/# free
Cisco 2950/2117760	Room N10	192.168.1.240/24	00000cab3546	24/0
Cisco 2950/2117761	Room N25	192.168.1.241/24	00000cab3547	24/0
Bay 28115/1A74215	Room S10	192.168.2.240/24	000003f25567	24/4

Network Troubleshooting Tools

Experience, colleagues, the Web, phone support, and documentation are helpful resources for network support and troubleshooting. Sometimes, however, the only place you can get the information you need is from your own network. Many networking problems occur at lower

layers of the OSI model, where they're often difficult to troubleshoot. Fortunately, there are tools for diagnosing these problems. The next sections discuss some of the most common tools and their possible uses on a network.

Ping and Trace Route

You have already worked with Ping and ICMP quite a bit, but take a closer look at the output of the ping command. Figure 13-5 shows the results of three ping commands sent by a computer with IP address 172.31.210.1: one successful and two not.

```
C:\Users\gtonsho>ping 172.31.210.2

Pinging 172.31.210.2 with 32 bytes of data:
Reply from 172.31.210.2: bytes=32 time=1ms TTL=128
Reply from 172.31.210.2: bytes=32 time<1ms TTL=128
Reply from 172.31.210.2: bytes=32 time<1ms TTL=128
Reply from 172.31.210.2: bytes=32 time<1ms TTL=128

Ping statistics for 172.31.210.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Users\gtonsho>ping 172.31.210.3

Pinging 172.31.210.3 with 32 bytes of data:
Reply from 172.31.210.1: Destination host unreachable.
Reply from 172.31.210.1: Destination host unreachable.
Reply from 172.31.210.1: Destination host unreachable.
Reply from 172.31.210.1: Destination host unreachable.

Ping statistics for 172.31.210.3:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

C:\Users\gtonsho>ping 172.221.1.5

Pinging 172.221.1.5 with 32 bytes of data:
Reply from 67.135.198.249: Destination net unreachable.
Reply from 67.135.198.249: Destination net unreachable.
Reply from 67.135.198.249: Destination net unreachable.
Reply from 67.135.198.249: Destination net unreachable.

Ping statistics for 172.221.1.5:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),

C:\Users\gtonsho>
```

Figure 13-5 Successful and unsuccessful pings

Courtesy of Course Technology/Cengage Learning

The first ping command has four replies from the address that was the target of the command: 172.31.210.2. You see four replies because the ping command sends four ICMP messages by default. The output for each successful reply contains the following pieces of information:

- The responding computer's IP address
- The number of bytes in the ICMP message
- The time receiving a reply takes
- The reply packet's TTL

The ping command tells you whether the computer can communicate with another computer via IP or, more specifically, ICMP. With a successful reply, you know the target machine is up and

running and there's a path between your computer and the target. This command also tells you the amount of time elapsed before receiving a reply, measured in milliseconds (ms). On a single LAN, the time should be very short—typically less than 10 ms. Even addresses on the Internet should normally reply in less than 100 ms. Time values that exceed 100 ms substantially can indicate a problem with the destination computer or the path between the source and destination computers.

Slow response times can be transient. A single Ping message or even several messages in a row with long response times don't necessarily indicate a problem; the computer or a network link might have been busy for a short time. Try the `ping` command several times to see whether response times are consistently high before assuming there's a problem. What if response times are consistently high? If you're the administrator of both the source and destination computers as well as path between them, you need to check CPU use on the computers and network use on the links between them. A network monitor program (discussed later in "Networks Monitors") can help determine whether network traffic is part of the problem. Performance Monitor, discussed in Chapter 9, can help determine whether the computer is responsible for the poor response times.

Review Figure 13-5 again. The second `ping` command receives a "Destination host unreachable" message. This message is generated by the source computer, indicating that an ARP request was sent to retrieve the MAC address 172.31.210.3, but an ARP reply was never received. This message is generated when the destination host is on the same network as the source, but the source can't retrieve the destination's MAC address. If the MAC address can be retrieved but the ping request times out, the destination computer's firewall might be blocking the ping.

In Figure 13-5, the third `ping` command receives a "Destination net unreachable" message. This message is generated by a router in the path between the source and destination computers when the router determines that the destination network doesn't exist. Not all routers return this message; if they don't, the Ping program generates a "Request timed out" message.

Using Ping for Connectivity Troubleshooting When troubleshooting a suspected connectivity problem with a computer, you should verify Physical-layer connectivity first. If there are link lights on the switch and/or NIC, the Physical layer is probably okay. You can verify the Network layer by trying the following checks:



These checks aren't necessarily in the order in which you should perform them. The order depends somewhat on the problem's symptoms and your troubleshooting experiences.

- *Run `ipconfig /all`*—Displays all pertinent information for your IP address configuration, including the IP address, default gateway, and DNS servers.
- *Ping the loopback address (127.0.0.1)*—A successful response verifies that the IP protocol is functioning correctly. It doesn't mean the IP address configuration is correct, however. The following checks do that.
- *Ping the local IP address*—Verifies the computer's capability to receive ICMP packets. If you can ping the loopback address but not the local computer's IP address, it's likely the firewall is blocking ICMP packets.
- *Ping the default gateway*—The default gateway is the address of the router the computer sends packets to when the destination host is on another network. If you can't

ping the default gateway, you won't be able to send packets outside your local LAN. If the host you're trying to communicate with is on the same LAN, you can skip this check.

- *Ping the host's IP address*—Verifies whether you can communicate with the target computer by using ICMP.
- *Ping the hostname*—Verifies that you can resolve the hostname to the correct IP address. If this check is unsuccessful, try the next two checks.
- *Ping DNS servers*—A response from one or more DNS servers indicates that the computer can communicate with a server that can resolve names to IP addresses, but it doesn't indicate that DNS lookups are working. If you can ping the DNS server, the next check verifies whether the DNS server can perform DNS lookups.
- *Use Nslookup*—Determines whether the DNS server can resolve the name of the host you're trying to communicate with. If it can't resolve the hostname to an address, try a well-known Internet name to see whether the problem happens only when looking up the target host's address.

Using Trace Route Trace Route, discussed in Chapter 5, is used in Windows as the `tracert` command to determine the path between two devices. Each router between the source and destination replies with a message, and the command output shows how long receiving the reply took. For each router in the path, three ICMP packets are sent, so three replies are received to indicate an average response time (see Figure 13-6).

```
C:\Users\gtonsho>tracert books.tonsho.com
Tracing route to books.tonsho.com [67.218.126.125]
over a maximum of 30 hops:
  0  1  <1 ns    <1 ns    1 ns    172.31.1.250
  1  2  1 ns     1 ns     1 ns     172.16.8.1
  2  3  1 ns     1 ns     1 ns     ycexpress.video.yc.edu [198.68.121.28]
  3  4  18 ns    21 ns    12 ns    pho-edge-06.inet.quest.net [67.135.198.249]
  4  5  5 ns     3 ns     4 ns     pho-core-01.inet.quest.net [205.171.12.77]
  5  6  75 ns    17 ns    14 ns    lap-brdr-03.inet.quest.net [67.14.22.78]
  6  7  14 ns    15 ns    21 ns    63.146.27.34
  7  8  26 ns    17 ns    17 ns    ae-61-68.ebr1.losangeles1.level3.net [4.69.144.50]
  8  9  18 ns    18 ns    17 ns    ae-5-5.car1.sandiego1.level3.net [4.69.133.205]
  9 10  17 ns    17 ns    18 ns    ae-11-11.car2.sandiego1.level3.net [4.69.133.218]
 10 11  17 ns    17 ns    18 ns    add2net-inc.car2.sandiego1.level3.net [4.53.122.70]
 11 12  18 ns    21 ns    17 ns    quasar.lunarpages.com [67.218.126.125]
Trace complete.
C:\Users\gtonsho>
```

Figure 13-6 Output of the `tracert` command

Courtesy of Course Technology/Cengage Learning

Trace Route does a reverse DNS lookup on each router's IP address and displays the router name, if possible. The router name can help you determine where the router is physically located, and the response times can help you determine whether there's a bottleneck between the source and destination. A consistently high response time by one or more routers can indicate that the

path is congested with excessive traffic. If a router in the path is administered locally, you can monitor the network link and the router itself to see where the problem lies.

In addition to showing you where bottlenecks might exist in an internetwork, Trace Route can confirm your network design. If you have a complex internetwork with multiple routes to some destinations, this command can show you which path your packets are taking. Most large internetworks with multiple routes for fault tolerance or load sharing have a preferred path. Router configuration determines the path packets should take, and Trace Route can verify whether your network configuration is operating as expected.



Hands-On Project 13-1: Troubleshooting with the OSI Model

Time Required: 10 minutes.

Objective: Troubleshoot a problem with the OSI model approach.

Required Tools/Equipment: Your classroom computer running Windows 7.

Description: You have been called to troubleshoot a problem with the connection to a Web server. The user states that when she starts a Web browser and tries to access a page at *www.tomsho.net*, the Web browser displays an error message after a short time. You try using the step-by-step OSI model approach to solving this problem.

1. Log on to your computer as an administrator.
2. Start a Web browser and go to **www.tomsho.net**. After a short time, you should get an error message. Write down what layers of the OSI model you tested, and then exit the Web browser.

3. Open a command prompt window. You want to see whether the Web server's name can be resolved to an IP address by using DNS. Type **nslookup www.tomsho.net** and press **Enter**. You should get a response giving you the address for *www.tomsho.net*. Some networking diagrams place DNS at the Session layer. Because the name can be resolved to an address, you move on. Write the response you received:

4. Next, you want to see whether you can communicate with the Web server by pinging it. Type **ping www.tomsho.net** and press **Enter**. You should get a successful reply. Write the reply you received and what layers of the OSI model you tested:

5. You now know that you can communicate with the Web server by using the Ping program. Start a Web browser again and go to **www.course.com**. If you're successful, you can conclude that the Web server software on *www.tomsho.net* isn't working correctly because the Web client software appears to be functioning. Write which layers of the OSI model are most likely the problem:

6. Close all open windows, but stay logged on if you're going on to the next project.

Network Monitors

Network monitors are software packages that can track all or part of the network traffic. By examining the packets sent across the network, a network monitor can track information such as packet type, errors, and traffic to and from each computer; collect this data; and generate reports and graphs. A number of full-featured products for network monitoring and analysis are available from companies such as WildPackets and Fluke Networks. IPSwitch's WhatsUp Gold Standard software is an inexpensive but capable network monitoring tool for small to medium businesses that monitors devices and includes an alert management system to notify an administrator of possible problems.

Sometimes your network monitoring needs are more modest. For example, if you run several servers, you might want to be notified if a server stops responding. Inexpensive applications such as Simple Server Monitor (shown in Hands-On Project 13-3) can monitor servers by using Ping and other protocols. If a server stops responding, the program can send an e-mail to an administrator.



Hands-On Project 13-2: Using Wireshark to Monitor Traffic

Time Required: 10 minutes

Objective: Use Wireshark to view network statistics.

Required Tools/Equipment: Your classroom computer running Windows 7, with Internet access and Wireshark installed

Description: You want to view statistical information about your network traffic by using Wireshark, specifically which protocols are in use and a list of network conversations.

1. Log on to your computer as an administrator, if necessary.
 2. Open Wireshark and start a capture session. Start a Web browser, go to several Web sites, and then close the Web browser. Open a command prompt window, and then type `ping -t www.yahoo.com` and press **Enter**. This command pings the target computer until you stop it. Minimize the command prompt window while the command continues to send pings.
 3. Click **Statistics, Protocol Hierarchy** from the Wireshark menu. You see a list of protocols that have been captured and the number of packets and bytes for each. Click **Close**.
 4. Click **Statistics, Conversations** from the Wireshark menu. You see a list of conversations that have been captured. Click the **IPv4** tab to limit the display to IPv4 conversations. The display is updated periodically, so you should see packets and bytes increase in the conversation between your computer and *www.yahoo.com* as long as the ping continues. Click **Close**. Write the IP addresses of the computers exchanging the most packets.
-
5. Click **Statistics, Summary** from the Wireshark menu. You see a summary of the number of packets captured, average size, total bytes, and so forth. Click **Close**.

- Open the command prompt window you minimized earlier. Press **Ctrl+C** to stop the ping, and then close the command prompt window.
- Close Wireshark and any other open windows. Stay logged on if you're going on to the next project.



Hands-On Project 13-3: Using Simple Server Monitor

Time Required: 20 minutes

Objective: Download and install a trial version of Simple Server Monitor.

Required Tools/Equipment: Your classroom computer running Windows 7 with Internet access

Description: You want to monitor the status of servers on your network. You research several options and decide that Simple Server Monitor meets your requirements.

- Log on to your computer as an administrator, if necessary.
- Start a Web browser and go to <http://simpleservermonitor.com>.
- Click **Download a 30-Day Free Trial**. Click the link indicating the current version, and then download and install Simple Server Monitor.
- After Simple Server Monitor is installed, run the application (see Figure 13-7).

Add Server icon

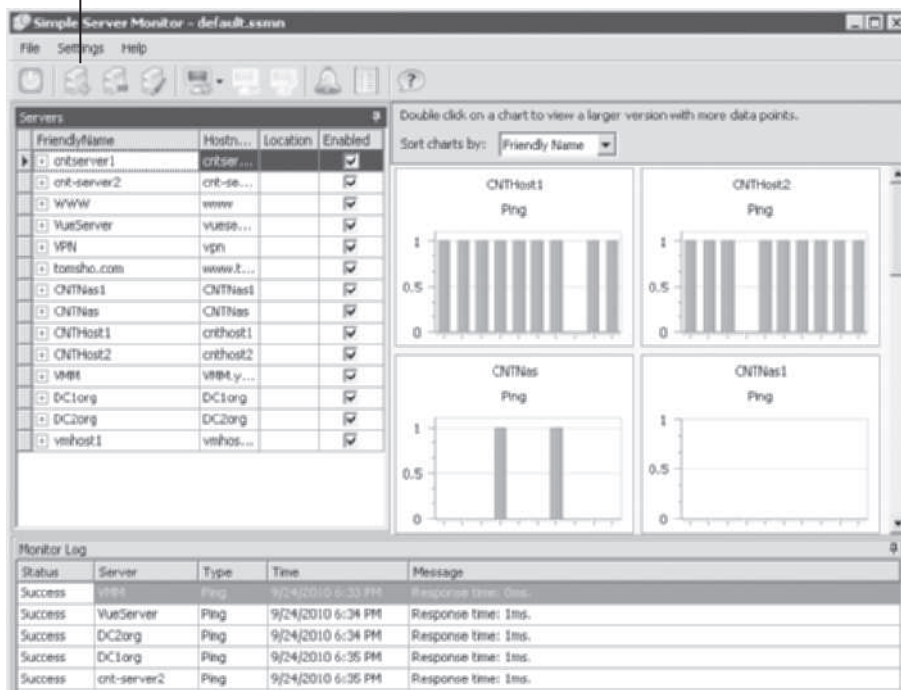


Figure 13-7 The Simple Server Monitor interface

Courtesy of Course Technology/Cengage Learning

5. Click the **Add Server** icon (tower with a green plus sign) to add a device to monitor. In the Hostname text box, type the IP address of your default gateway (see Figure 13-8). If you don't know the address, use the `ipconfig` command to retrieve it.

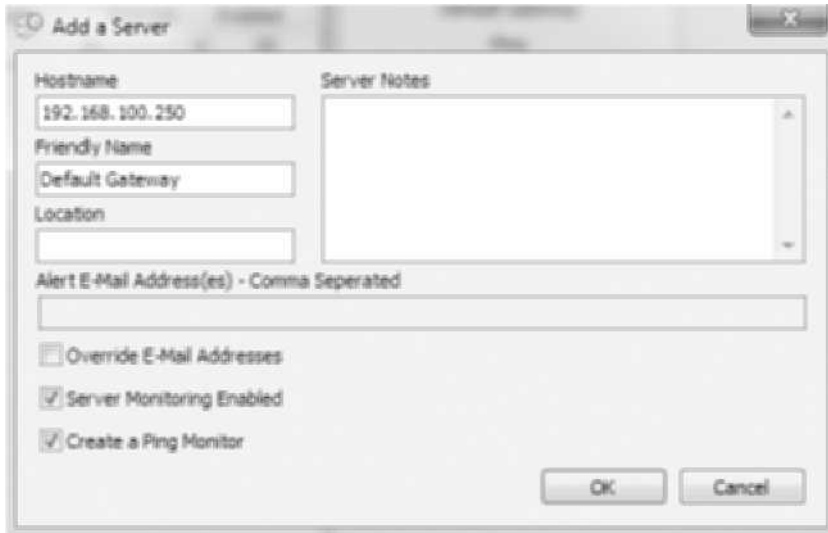


Figure 13-8 The Add a Server dialog box

Courtesy of Course Technology/Cengage Learning

6. In the Friendly Name text box, type **Default Gateway**. Note that you can include an e-mail address in this window. Simple Server Monitor e-mails you when a device fails to respond if you include your e-mail address. You need to configure Simple Server Monitor with the address of an SMTP server if you want to receive e-mail alerts. Click **OK**.
7. Add your DNS server address and add the address of some Internet servers, such as *www.yahoo.com* and *www.google.com*.
8. In the left pane, right-click **Ping** under the Default Gateway monitor and click **Edit Monitor**. In the Ping Monitor window, you can adjust the interval to change the frequency at which Simple Server Monitor attempts to contact the device. You can also adjust the timeout and number of retries.
9. Aside from monitoring with Ping, Simple Server Monitor can verify that particular services are running. Right-click one of the Internet servers you added (such as *www.yahoo.com*), point to **Add Monitor**, and click **Add HTTP Monitor**. Click **OK** in the HTTP Monitor window. This monitor verifies that the Web server service is running. Ping tests up to the Network layer; the HTTP monitor tests through the Application layer.

10. Add another device that you know can't reply (for example, 10.255.250.200, if it's an invalid address in your network) so that you can see what Simple Server Monitor displays when the monitor fails to get a reply.
11. Simple Server Monitor runs when Windows starts by default. To prevent it from doing so, click **Settings** and click to clear the **Run on Windows Startup** check box. Using Simple Server Monitor, you can monitor all your servers, routers, and other important devices so that you're informed quickly if a device stops responding. Close all open windows, but stay logged on for the next project.

Protocol Analyzers

A protocol analyzer enables you to capture packets and analyze the network traffic generated by different protocols. You have used Wireshark to capture different packet types. Microsoft Network Monitor, an application similar to Wireshark, can be installed from Add/Remove Programs in Windows Server 2003 and earlier, but you must download it in Windows Server 2008 or Windows 7.

Protocol analyzers can help solve problems from the Data Link layer up to the Application layer because they decode the Data Link-, Network-, and Transport-layer headers and display data processed by the Application layer. With a protocol analyzer, you can troubleshoot problems with DNS, authentication, DHCP, IP addressing, remote access, and many other processes. It's also a great teaching and learning tool because you can capture entire conversations created by a certain protocol to see exactly how it works.

Many protocol analyzers, including Wireshark, have an expert mode that attempts to detect problems with the stream of frames it decodes. For example, it can detect problems such as TCP retries, in which TCP has to resend segments because no acknowledgement was received, or duplicate acknowledgements, which can occur when segments are sent again before the acknowledgement is received, indicating a slowdown in communication. Protocol analyzers can detect a host of other potential problems, too, that just might help you solve your original problem.

The most advanced protocol analyzers combine hardware and software in a self-contained unit. These analyzers sometimes include a built-in cable analyzer for solving Physical-layer problems, too. Some examples of protocol analyzers include the following:

- *Microsoft Network Monitor*—An older version is available as an installation option in Windows Server 2003, but you can download the latest version for Windows Server 2008, Windows 7, Windows Vista, Windows XP, and Windows Server 2003.
- *WildPackets OmniPeek*—This software-only protocol analyzer handles all major networking protocols. For more information, visit www.wildpackets.com.
- *Fluke Network OptiView Network Analyzer*—This portable hardware-based network analyzer can be used with wired or wireless networks. For more information, visit www.flukenetworks.com.



- *Wireshark*—This popular free protocol analyzer is available for both Windows and Linux/UNIX environments. It supports all major protocols and a number of lesser-known protocols. For more information and to download it, visit www.wireshark.com.

Most experienced network administrators rely on protocol analyzers to establish baselines for network performance and to troubleshoot their networks, especially when they suspect software problems or when Network-layer (Layer 3) devices appear to be responsible for network problems.



Hands-On Project 13-4: Installing and Using Microsoft Network Monitor

Time Required: 10 minutes

Objective: Install Windows Network Monitor.

Required Tools/Equipment: Your classroom computer running Windows 7 with Internet access

Description: You have used Wireshark for network monitoring and protocol analysis, but you want to try Microsoft's latest version of Network Monitor to compare its features.

1. Log on to your computer as an administrator.
2. Start a Web browser and go to download.microsoft.com. In the Search text box, type **Network Monitor 3.4** and press **Enter**. Click **Microsoft Network Monitor 3.4** in the search results.
3. If your Windows OS is 64-bit, click **Download** next to the file with x64 in its name. If your Windows OS is 32-bit, click **Download** next to the file with x86 in its name. (If you aren't sure, click **Download** next to the file with x86 in its name.) Click **Run** when prompted to run or save the file.
4. When asked whether you want to run this software, click **Run**. Click **Yes** in the Microsoft Network Monitor message box. In the Welcome to the Microsoft Network Monitor 3.4 Setup Wizard window, click **Next**.
5. In the End-User License Agreement window, click **I accept**, and then click **Next**. In the Choose Setup Type window, click **Typical**. In the Ready to Install window, click **Install**. If you see a User Account Control message, click **Yes**.
6. In the Completing the Setup Wizard window, click **Finish**. If you see a User Account Control message, click **Yes**.
7. When the installation is finished, log off and log back on. Double-click the **Microsoft Network Monitor** desktop icon to start Network Monitor. You see a window similar to Figure 13-9. Read the message in the right pane to become familiar with this tool's features.



Figure 13-9 The Microsoft Network Monitor Welcome window

Courtesy of Course Technology/Cengage Learning

8. Click **New Capture** on the toolbar to start a packet capture session, and then click the **Start** button.
9. Start a Web browser, if necessary, and go to a Web site. Then exit the Web browser. In the Network Conversations pane, you can select an application to view the packets generated by that application. Click the name of the Web browser you used to see the packets generated by your Web browsing session (see Figure 13-10).
10. Click **Experts**, **Download Experts** from the menu. In the Web page that opens, click **TCP Analyzer** under Other Community Experts. Under Downloads, click the file (x64 or x86) for your environment, and then click **Download**. Install the file.
11. After the TCP Analyzer expert is installed, click **Experts** from the Network Monitor menu, point to **TCP Analyzer**, and click **Run Expert**. Type a filename, such as **HTTP Traffic**, and click **Save**.

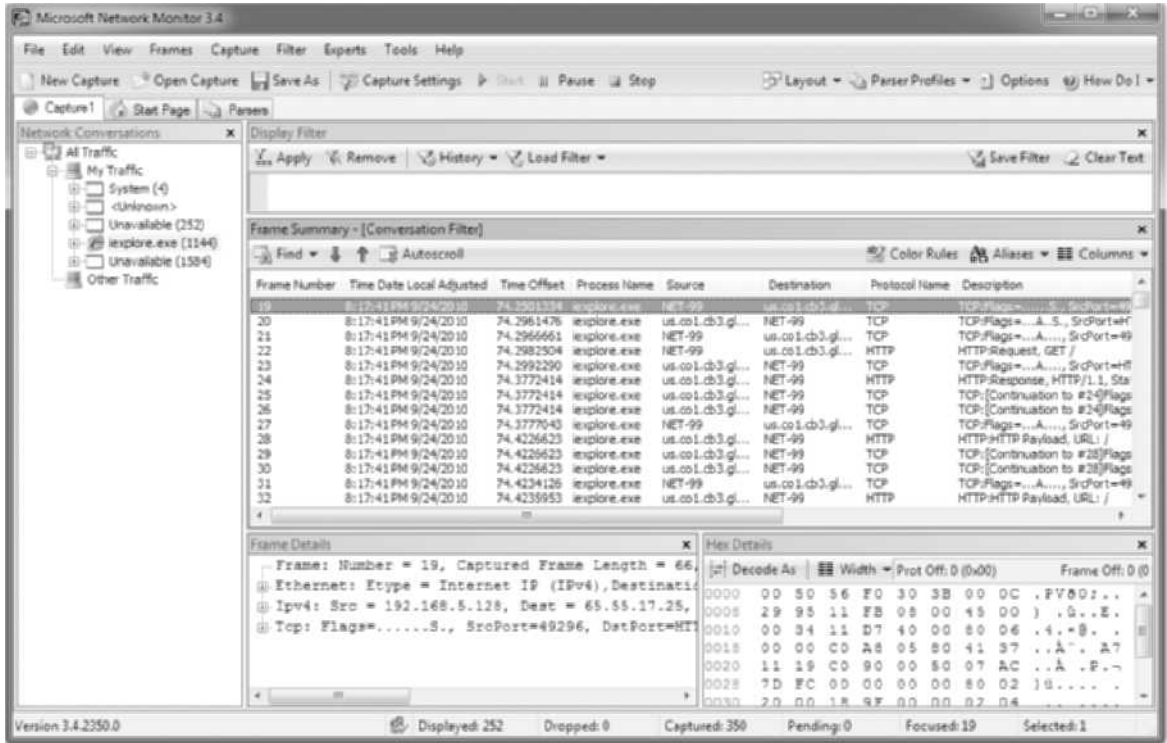


Figure 13-10 A Web browser session captured in Microsoft Network Monitor

Courtesy of Course Technology/Cengage Learning

12. The TCP Analyzer runs and displays an analysis of the TCP packets you captured. Close TCP Analyzer after you have reviewed the information. Don't be concerned if you don't understand all the information in the analysis.
13. Explore other features of Network Monitor, and then close it when you're finished. Close all open windows, but stay logged for the next project.

Time-Domain Reflectometer

You can use a **time-domain reflectometer (TDR)** to determine whether there's a break or short in a cable, and then measure the cable's length. A TDR can pinpoint how far from the device the break is located by sending an electrical pulse down the cable that reflects back when it encounters a break or short. It measures the time it takes for the signal to return and, based on the type of cable tested, estimates how far down the cable the fault is located. A high-quality TDR can determine the location of a break within a few inches. TDRs are available for fiber-optic as well as electrical cables.

Although cable installers use them most often, TDRs can be invaluable diagnostic tools for network administrators, too. When you're having new cables installed, ask the installer to use a TDR to document actual lengths of all cables. Rent a TDR (or hire someone who owns one) to measure any cables on the network whose lengths aren't documented already. The TDR function is standard in most advanced cable testers. Remember that each medium has distance

limitations, so running a TDR scan on each cable segment is critical in documenting your Physical-layer installation.

Basic Cable Testers

You can purchase basic cable testers for less than \$100. Typically, they test only correct termination of a twisted-pair cable or continuity of a coaxial cable. They're excellent tools for checking patch cables and testing for correct termination of a cable at the patch panel and jack. However, these testers can verify only that cable wires are terminated in the correct order (referred to as the "wiremap") or that there are no breaks in the cable. These low-priced testers can't check a cable for attenuation, noise, or other possible performance problems in the cable run.

Advanced Cable Testers

More expensive than TDRs or basic cable testers, advanced cable testers perform a battery of tests on a cable for the purpose of certifying the cable for a particular application. For example, you can set the cable tester for Cat 6 cable, and it performs the tests needed to certify the connection for Cat 6 operation. In addition to length and wiremap tests, these cable testers perform several tests for crosstalk, attenuation, EMI, and impedance mismatches. Some advanced cable testers function at both the Physical and Data Link layers to measure frame counts, collisions, CRC errors, and broadcast storms. Anyone who plans to be in the business of installing cable needs an advanced cable tester, which can cost from around \$1000 to several thousand dollars, depending on the features.

Advanced Monitoring Tools

Simple Network Management Protocol (SNMP), part of the TCP/IP protocol suite used for network management, is an industry-standard protocol that most networking equipment manufacturers support. In a Windows environment, SNMP management can be installed through the Programs and Features applet in Control Panel.

To use SNMP, SNMP **software agents** are loaded on network devices you want to manage and monitor. Each agent monitors network traffic and device status and stores information in a **management information base (MIB)**. To use the information gathered by software agents, a computer with an SNMP management program must be on the network. This management station communicates with software agents and collects data stored in the MIBs on network devices. Then it combines information from all networking devices and generates statistics or charts of current network conditions. With most SNMP managers, you can set thresholds for sending alert messages to network administrators when these thresholds are exceeded.

In addition, you can manage many network components with SNMP. With software agents, you can configure networking devices and, in some cases, reset them from the management station. SNMP can manage network devices, such as switches and routers, and important network resources, such as servers. An SNMP management program can query these devices and even make configuration changes remotely to help managers control networks from a single application.

Remote Monitoring (RMON) is an advanced network-monitoring protocol that extends SNMP's capabilities. It comes in two versions: RMON1 and RMON2. SNMP defines a single MIB type to collect network data, but RMON1 defines nine other MIB types, called RMON groups, to provide a more comprehensive set of data about network use. RMON-capable

devices, such as hubs, routers, and switches, contain software agents called probes that collect data and communicate with a management station by using SNMP.

RMON1 is designed to capture data and collect statistics at the Data Link and Physical layers. RMON2 can collect and analyze traffic at the Network layer and higher layers, which makes detailed analysis of enterprise network and application software operation possible. RMON-capable devices aren't inexpensive, but being able to monitor networks and solve network or application problems before they become serious is well worth the expense, considering the benefit of increased productivity for organizations.



For more information, you can find an excellent paper on RMON1 and RMON2 at www.cisco.com/en/US/docs/internetworking/technology/handbook/RMON.html.

Common Troubleshooting Situations

Say you have the following problem on a network: 10 Mbps hubs are being replaced with 100 Mbps switches. It's been confirmed that all NICs can support 100 Mbps, but certain stations simply can't communicate on the network. After some investigation, you've determined that the NIC driver software was set to force the NIC to communicate at 10 Mbps in half-duplex mode, but the switch was set to communicate only at 100 Mbps in full-duplex mode. A quick change in the NIC configuration solved the problem.

Using the structured problem-solving approach to network troubleshooting described earlier, you can eventually solve networking problems such as this one. To help get you started with this sometimes arduous task, this section outlines some common network problems and possible solutions.

Cabling and Related Components

Many networking problems occur at the Physical layer and include problems with cables, connectors, and NICs. The first step in troubleshooting these problems is to determine whether the problem lies with the cable or the computer. One easy way to do this is to connect another computer to the cable. If it functions normally, you can conclude that the problem is with the original computer. If it exhibits the same symptoms, check the cable first, and then check the device it connects to, and so forth.

After you determine that the cable is the likeliest culprit, make sure it's connected to the computer correctly, and verify that it's the right type of cable for the connection. Make sure you use the same type of UTP cable throughout the network. Double-check cable lengths to make sure you don't exceed the maximum length limitation for the network medium. By using a TDR, you can identify and correct these types of problems quickly.

If you suspect a faulty or misconfigured NIC, check the back of the card. As discussed, the NIC might have indicator lights to show whether it's functioning and its network connection is active. If the NIC lacks these indicators, you must replace the suspect NIC with a known working NIC—in much the same way you replace a suspect computer with a known working one to determine whether the network or the computer is the cause of the problem.

If the NIC seems functional and you're using TCP/IP, try using the Ping program to check connectivity to other computers. If the NIC works but the computer still can't access the network, you might have more serious hardware problems (for example, a faulty bus slot), or NIC configuration settings might be invalid. Either way, you must conduct further troubleshooting.

Power Fluctuations

Power fluctuations in a building—caused by an electrical storm or a power failure, for example—can affect computers adversely. First, verify that servers are functioning. When possible, remind users that it takes a few minutes for servers to come back online after a power outage.

One way to eliminate the effects of power fluctuations, especially on servers, is to connect them to uninterruptible power supplies (UPSs). UPS systems provide battery power to computers so that they can be shut down without data loss. Some perform shutdowns automatically, thereby eliminating the need for human intervention when power failures or severe power fluctuations occur.

Upgrades

Because networking technology changes constantly, frequent upgrades of equipment and software, such as the file server's OS, are necessary. During these upgrades, it's common for some equipment to run on an old OS and some to run on a new one. When you perform network upgrades, remember three important points:

- Ignoring upgrades to new software releases and new hardware can lead to a situation in which a complete network overhaul is necessary because many upgrades build on top of others. If administrators don't keep current, they might need to do an overwhelming amount of research and endure a lack of technical support for older software or hardware. Keep current and do one upgrade at a time to make your life easier.
- Test any upgrade before deploying it on your production network. Ideally, use a test laboratory where you can try all upgrades and work out any problems. If a test lab isn't an option, select a small part of your network—one department or a few users—and perform the upgrade. This method gives you an opportunity to work through possible issues before imposing changes (and the problems that sometimes go with them) on the entire network.
- Don't forget to tell users about upgrades. A well-informed user is an understanding user. Everyone who might be affected by an upgrade must be informed when it will occur, what's involved, and what to expect.

Poor Network Performance

If all goes well, the network monitoring and planning you do will ensure that the network performs optimally. However, you might notice that your network slows down; this problem can happen quickly, or it might be a gradual deterioration. Whether performance problems are exhibited slowly or suddenly, answering the following questions should help pinpoint the causes:

- What has changed since the last time the network functioned normally?
- Has new equipment been added to the network?

- Have new applications been added to computers on the network?
- Is someone playing electronic games across the network? (You'd be surprised at the amount of traffic networked games can generate.)
- Are there new users on the network? How many?
- Could any other new equipment, such as a generator, cause interference near the network?

If new users, added equipment, or newly installed applications seem to degrade network performance, it might be time to expand your network and add equipment to limit or contain network traffic. Higher-speed backbones, network partitions, and additional servers and routers are alternatives worth considering when you must increase capacity to accommodate usage levels that have grown beyond your network's current capabilities.

Disaster Recovery

If your network is well documented, recovering from a disaster will be much easier. Disasters can come in many forms, from a simple disk crash that disables a key server to a fire or flood that devastates your entire workplace. The procedures for recovery from total devastation are beyond the scope of this book; instead, the following sections focus on backup procedures and recovery from system failure.

Backing Up Network Data

A comprehensive backup program can prevent major data loss. A backup plan is an important part of an overall disaster recovery plan and should be revised as your needs—and data and applications—change. To formulate your backup plan, review the following guidelines:

- Determine what data should be backed up and how often. Some files, such as program executables and OS files, seldom change and might require backup only weekly or monthly.
- Develop a schedule for backing up data that includes the type of backup to perform, how often, and at what time of day. The next section reviews the most common backup types.
- Identify the people responsible for performing backups.
- Test your backup system regularly. The person responsible for backups should perform these tests, which include backing up data and restoring it. After a backup system is in place, conduct periodic tests to ensure data integrity. A data backup is no use to you if the restore process doesn't work.
- Maintain a backup log listing what data was backed up, when the backup took place, who performed the backup, and what medium was used. Windows Backup maintains a summary of the backup history, as do third-party backup programs, but these logs can't be accessed if the backup server fails.
- Develop a plan for storing data after it's been backed up. This plan should include on-site storage, perhaps in a fireproof safe, and off-site storage in the event of a catastrophe. For both on-site and off-site storage, make sure only authorized personnel have access to the backup medium.

Backup Types

Chapter 9 discussed using Windows built-in backup programs to perform basic backup of Windows computers, but most backup plans require more flexible backup software when more than a few computers are involved. Third-party backup programs can provide a comprehensive backup solution with centralized management. Before using some of these products, you need to understand some of the terminology for backup types:

- *Full*—A **full backup** copies all selected files to the selected medium and marks files as backed up; also called a normal backup.
- *Incremental*—An **incremental backup** copies all files changed since the last full or incremental backup and marks files as backed up.
- *Differential*—A **differential backup** copies all files changed since the last full backup and doesn't mark files as backed up.
- *Copy*—A **copy backup** copies selected files to a selected medium without marking files as backed up.
- *Daily*—A **daily backup** copies all files changed the day the backup is made and doesn't mark files as backed up.

Of these five types, full, incremental, and differential backups are most useful as part of a regular backup schedule. A copy backup is good for copying files to a new location or making a secondary backup for off-site storage without affecting the regular backup routine. A daily backup is good for identifying files that were changed on a particular day so that you can, for example, collect all the files you were working on at home and bring the changed files back to the office the next day.

A good model for creating a backup schedule combines a weekly full backup with daily differential backups so that you can do backups quickly on a daily basis and restore data easily by restoring the contents of two backups: the full backup first and the differential backup next. This method uses more storage space for differential backups because the backup gets larger each day until the next full backup is done. However, it has the advantage of requiring only the full backup and the last differential backup.

You can also use incremental backups for daily backup, but restoring data is more difficult because you need the full backup plus each incremental backup done since the last full backup. The advantage of an incremental backup is that each backup takes up less space because incremental backups copy only files that were changed since the previous incremental or full backup.

When creating a backup schedule, posting the schedule and assigning one person to do backups and sign off on them each day is a good idea. That way, you can see at a glance when the last backup was done, and train one person to perform backups and care for the backup medium.

Windows systems have an additional backup type called a system state backup that copies the boot files, the Registry, Active Directory on domain controllers, and other critical information. If a Registry file or boot file becomes corrupted, sometimes all that's needed to get the system running again is the latest system state backup.

If you back up to tape media, make sure you can restore data. Use the “verify data” option that comes with your backup software to ensure that data copied to tape matches data on the drive. Create some test files, back up these files, and then practice restoring them to the

server to check that the restore operation works correctly. In addition, make sure tapes are stored in a cool, dry, dark place to minimize the risk of damage by heat, moisture, or light. Take a tape off the shelf periodically, and make sure it's readable and its data can be restored *after* the tape has been removed from the machine. For example, a miscalibrated tape drive might accept tapes for backup but fail to restore their contents—a condition usually discovered only when you need to restore data. Have a policy to rotate tapes so that no single tape set is reused in the same week. In addition, have a policy to remove tapes from the set after a predetermined time to avoid worn tapes that might affect performance.

System Repair and Recovery in Windows

Windows systems occasionally fail to boot or exhibit problems or errors after booting that indicate the system is damaged or corrupted and possibly needs repair. Windows OSs include repair utilities to correct these problems. The repair and recovery tools for Windows systems discussed in this section are Startup Repair, Last Known Good Configuration, System Restore, and Driver Rollback.

Startup Repair Startup Repair is a feature of Windows Recovery Environment (WinRE) that's accessed by booting to the installation DVD and selecting “Repair your computer” instead of the “Install now” option. Startup Repair attempts to repair your system automatically if it finds problems with the boot environment. If repairs can't be made automatically, you see the System Recovery Options dialog box shown in Figure 13-11.



Figure 13-11 The System Recovery Options dialog box

Courtesy of Course Technology/Cengage Learning

The System Recovery Options dialog box has the following options:

- *Startup Repair*—Attempts to fix problems automatically that keep Windows from starting. Examples are corrupted boot configuration files and a damaged Master Boot Record.
- *System Restore*—Restores Windows system state settings to an earlier point in time.
- *System Image Recovery*—Restores one or more disks from a backup image created by Windows Backup.
- *Windows Memory Diagnostic*—Performs a complete memory test.
- *Command Prompt*—Opens a command prompt window, where you have access to a variety of command-line troubleshooting tools.

Last Known Good Configuration The Last Known Good Configuration repair utility can fix a Windows system that doesn't boot because a new device driver was installed. It restores specific Registry information that existed after the last successful boot. If a problem occurs after a successful boot and user logon, the Last Known Good Configuration boot option doesn't help. It's very specific in what it restores—the Registry key HKLM\System\CurrentControlSet. The best time to use this utility is when you install a new device driver and restart the system, and the OS doesn't boot. To use it, start your Windows computer, and press F8 after the system BIOS window is displayed. You should see a window similar to Figure 13-12.

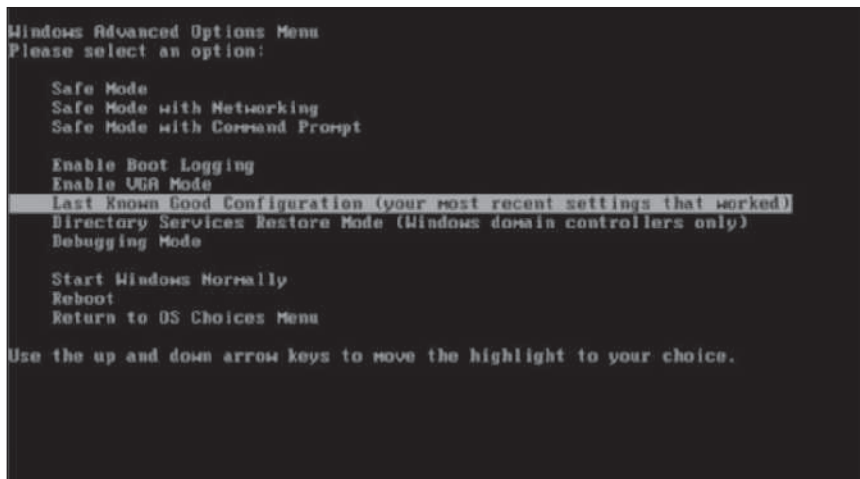


Figure 13-12 Selecting the Last Known Good Configuration option

Courtesy of Course Technology/Cengage Learning

System Restore Windows client OSs include System Restore (which isn't included in Windows server OSs), a handy utility that restores a system to a previous state. It monitors all drives and partitions and records changes made to system files, such as the Registry, and

some applications. When major changes are made, System Restore creates a restore point so that the computer can be restored to its operating parameters before the changes, in case something goes wrong with a driver installation, application install, or Registry change.

Users can also create their own restore point. For example, if you're about to install a hardware device and several associated applications, you can create a restore point before the install. If the system becomes unstable after the installation, you can restore Windows to the point before the installation. System Restore doesn't delete user files or application files that have been installed; you must still use Add/Remove Programs to uninstall applications. However, any system files (such as the Registry) changed by the application setup program are undone. System Restore can run from a regular boot or from the System Recovery Options dialog box discussed previously.

Driver Rollback Included in Windows client and server OSs, the Driver Rollback feature is used when a new driver installed for an existing device causes a problem with the system. To use this feature, open Device Manager. Double-click the device for which you want to roll the driver back to a previous version, click the Driver tab, and then click the Roll Back Driver button. Driver Rollback is useful only when a previously working device stops working correctly after a driver update.



Hands-On Project 13-5: Creating a System Restore Point

Time Required: 10 minutes

Objective: Create a system restore point in Windows 7 and revert your system to this restore point.

Required Tools/Equipment: Your classroom computer running Windows 7

Description: In this project, you create a system restore point in Windows 7 and then restore your machine by using it.

1. Log on to your computer as an administrator.
2. Click **Start**, right-click **Computer**, and click **Properties** to open the System control panel.
3. Click **System Protection** to open the System Properties dialog box to the System Protection tab. In the Protection Settings section, make sure protection is set to **On** for at least the drive marked as the system drive (which is drive C in Figure 13-13).
4. Click the **Configure** button to view current system protection settings. By default, Windows 7 is set to restore system settings and previous versions of files. You can change the option to restore only previous versions of files or turn off system protection. For the system drive (usually C), you should leave the default setting. You can also adjust the maximum amount of disk space reserved for system protection. The more you reserve, the more restore points Windows keeps before deleting them automatically. Click **OK**.



Figure 13-13 The System Protection tab

Courtesy of Course Technology/Cengage Learning

5. Click the **Create** button to create a new restore point. When prompted to enter a name for the restore point, type **TestRestore**. The name should be descriptive, such as **BeforeInstallingNewDriver**. The date and time are added automatically to the name you enter. Click **Create**.
6. After the restore point has been created, click **Close**.
7. Restart Windows 7. When Windows 7 begins to boot, press **F8** until you see the Advanced Boot Options menu (see Figure 13-14). If you don't press **F8** in time, restart Windows 7 and try again. Press **Enter** to select the Repair Your Computer option.

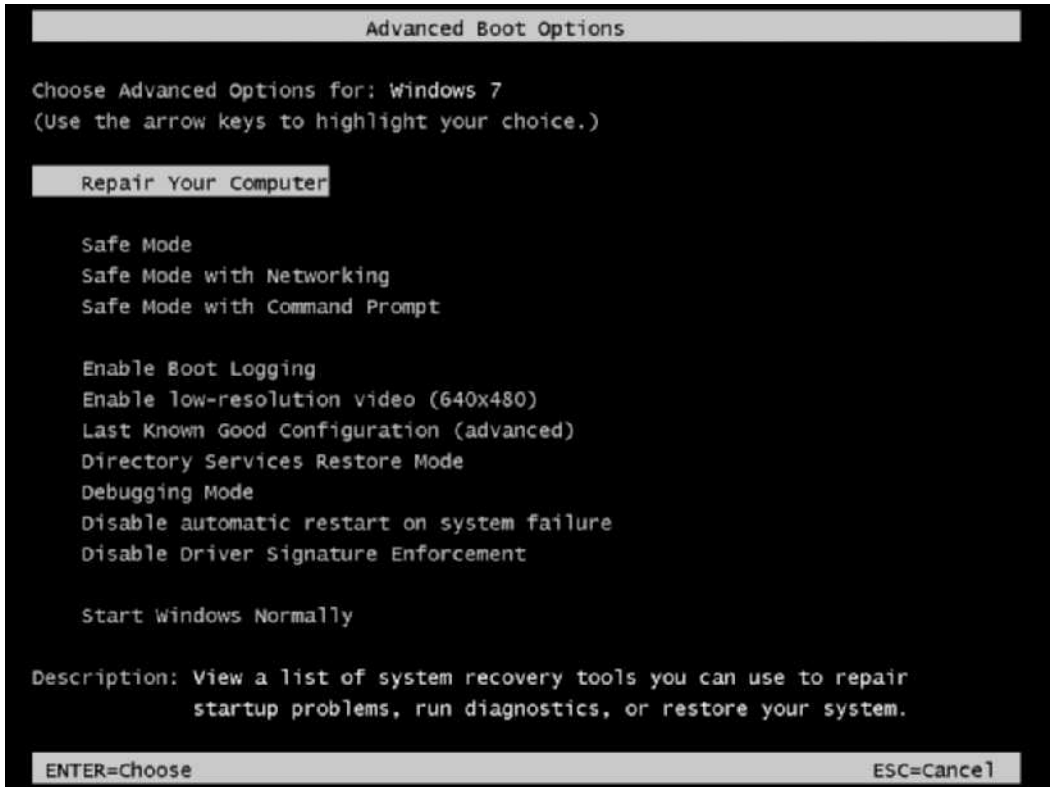


Figure 13-14 The Advanced Boot Options menu

Courtesy of Course Technology/Cengage Learning

8. In the System Recovery Options window, select a keyboard input method, if necessary, and then click **Next**.
9. When prompted, enter the administrator password, and then click **OK**.
10. In the System Recovery Options dialog box (shown previously in Figure 13-11), click **System Restore**.
11. In the Restore system files and settings window, click **Next**. You see a list of system restore points with your most recent one at the top and selected (see Figure 13-15). Click **Next**.
12. Click **Finish** in the Confirm your restore point window, and then click **Yes** to confirm you want to continue.
13. After the restore is finished, click **Restart**. List two to three situations in which you can put System Restore to good use:

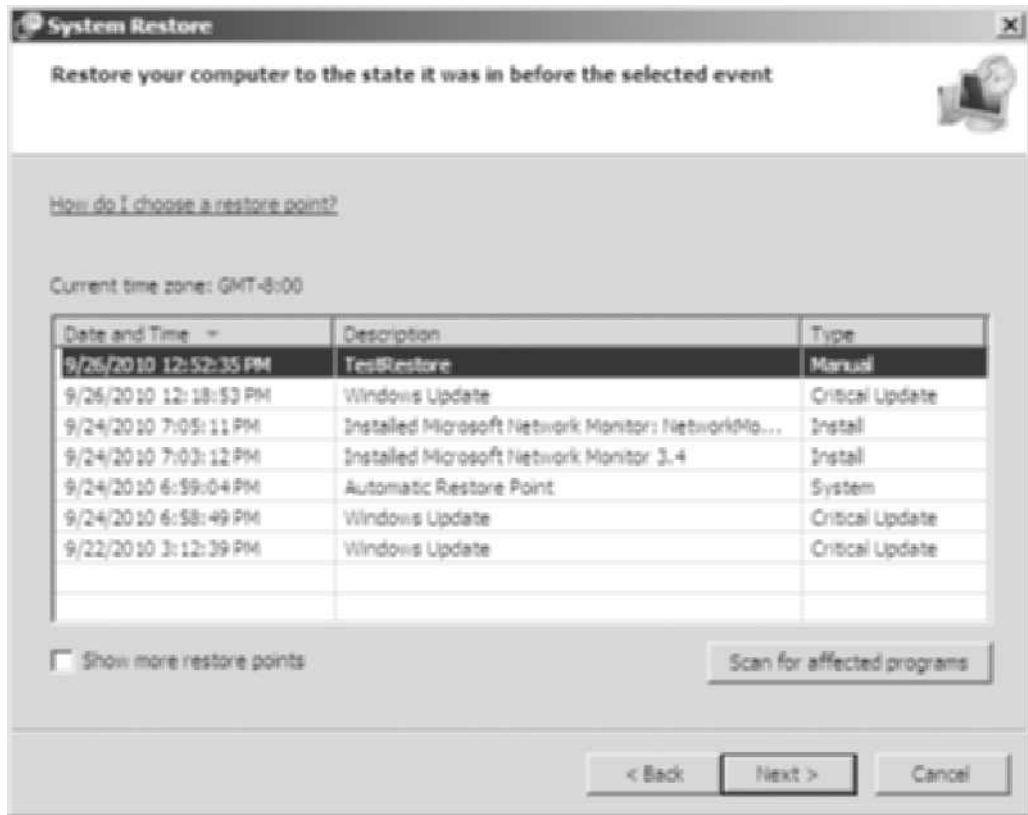


Figure 13-15 Choosing a restore point

Courtesy of Course Technology/Cengage Learning

- System Restore is a good tool to use when your system won't boot or if certain processes and applications don't work correctly but did at one time. You can run it from the Advanced Boot Options menu if Windows won't boot; if Windows does boot, you can run it from the System control panel. Close all open windows.

Chapter Summary

- Documenting a network thoroughly offers the following advantages: makes moves, adds, and changes easier, provides information needed for troubleshooting, offers justification for adding staff or equipment, helps determine compliance with standards, supplies proof that your installations meet hardware or software requirements, reduces training requirements, facilitates security management, and improves compliance with software licensing agreements.

- Elements of the network that should be documented include a network description, the cable plant, equipment rooms and telecommunication closets, internetworking devices, servers, and workstations.
- There are many approaches to network troubleshooting, including trial and error, solve by example, the replacement method, and step by step with the OSI model. Different methods are suitable for different circumstances and the technician's level of experience.
- The problem solving process has eight steps: Determine the problem definition and scope, gather information, consider possible causes, devise a solution, implement the solution, test the solution, document the solution, and devise preventive measures.
- Many resources are available to help troubleshoot problems, including an administrator's experience, the Web, and network documentation. Sometimes a combination of these methods can be used to solve problems.
- You can use several tools to get information from your network and its devices to help solve problems, including the Ping and Trace Route programs, network monitor programs, protocol analyzers, and basic and advanced cable testers.
- Some common troubleshooting situations are related to cable plant components, electrical power, and network and software upgrades. Poor network performance can be a result of many factors, and you must know what questions to ask to narrow down the problem's scope and possible causes.
- If your network is well documented, recovering from network disaster will be easier. Disaster can come in the form of a simple disk crash that disables a key server or a fire or flood that devastates your entire workplace. Some tools to prevent disasters and expedite recovery include data backups and system repair and recovery features in Windows.

Key Terms

copy backup A backup that copies selected files to the selected medium without marking files as backed up.

daily backup A backup that copies all files changed the day the backup is made; doesn't mark files as backed up.

differential backup A backup that copies all files changed since the last full backup; doesn't mark files as backed up.

frequently asked questions (FAQ) A document with two parts to each entry. The first part is a question the manufacturer has anticipated or received from customers; the second part is an answer to the question.

full backup A backup that copies all selected files to the selected medium and marks files as backed up. Also called a normal backup.

incremental backup A backup that copies all files changed since the last full or incremental backup and marks files as backed up.

management information base (MIB) A collection of network data stored by Simple Network Management Protocol software agents. *See also* software agents.

network monitors Programs that monitor network traffic and gather information about packet types, errors, and packet traffic to and from each computer.

Remote Monitoring (RMON) An advanced network-monitoring protocol that extends Simple Network Management Protocol's capabilities; contains software agents called probes that collect data and communicate with a management station by using SNMP.

rollback plan The part of an upgrade plan with instructions on how to undo the upgrade if problems occur during or after the upgrade.

software agents Simple Network Management Protocol components that are loaded on network devices; they monitor network traffic and device status information and send it to a management station.

time-domain reflectometer (TDR) A network troubleshooting device that can determine whether there's a break or short in the cable and, if so, approximately how far down the cable it's located. Also shows the total cable length.

Review Questions

1. A(n) _____ can determine whether a cable break or short exists and approximately how far down the cable it's located.
 - a. Network monitor
 - b. RMON
 - c. Time-domain reflectometer
 - d. Protocol analyzer
2. At what layer of the OSI model is a CRC error detected?
 - a. Data Link
 - b. Network
 - c. Physical
 - d. Session
3. Which type of plan makes returning a network to its state before an upgrade easier?
 - a. Backup plan
 - b. Rollback plan
 - c. Upgrade plan
 - d. Downgrade plan
4. Based on the eight-step problem-solving process discussed in this chapter, what's the first step in network troubleshooting?
 - a. Devise a solution.
 - b. Consider possible causes.
 - c. Test the cables.
 - d. Determine the problem definition and scope.
5. Which problem-solving approach requires a solid understanding of how networks work?
 - a. Trial and error
 - b. Step by step with the OSI model

- c. Solve by example
 - d. Replacement method
6. Under which condition is using the trial-and-error approach not advisable? (Choose all that apply.)
- a. You can undo changes easily.
 - b. A server or network device is live on the network.
 - c. You can't undo changes easily.
 - d. No data can be lost.
7. Which layer is where problems restricted to one workstation are likely to occur?
- a. Data Link
 - b. Transport
 - c. Physical
 - d. Presentation
8. After implementing a solution, what's the next step?
- a. Document the solution.
 - b. Test the solution.
 - c. Devise preventive measures.
 - d. Consider possible causes.
9. Which of the following can be used to prevent data loss during a power fluctuation or failure?
- a. TDR
 - b. UPS
 - c. MIB
 - d. SNMP
10. Which of the following is an element of the information-gathering step in the problem-solving process? (Choose all that apply.)
- a. Find out whether the function ever worked.
 - b. Determine whether something has changed.
 - c. Compare current operation with a baseline.
 - d. Consider possible causes.
11. When measuring network performance, what do you need as a point of reference?
- a. SNMP
 - b. Baseline
 - c. Message information base
 - d. Protocol analysis

12. Which of the following is the TCP/IP protocol used to configure and monitor network resources?
 - a. ICMP
 - b. SNMP
 - c. DHCP
 - d. SMTP
13. A logical topology includes jack locations and room numbers. True or False?
14. What tool does its job by measuring the amount of time it takes for a signal to travel the length of the cable and back?
15. What is an advanced network-monitoring protocol with more capabilities than SNMP?
 - a. IMAP
 - b. RMON
 - c. DNS
 - d. SMTP
16. For which of the following network conditions must you use a protocol analyzer or network monitor for further diagnosis? (Choose all that apply.)
 - a. Cable break
 - b. Cable short
 - c. Slow network performance
 - d. High rate of transmission errors
17. At what layers of the OSI model does a software protocol analyzer operate?
 - a. Layers 1, 2, and 3
 - b. Layers 1 to 4
 - c. Layers 2 and 3
 - d. Layers 2 to 7
18. Which of the following is *not* found in the output of a successful Ping reply?
 - a. IP address of the responding computer
 - b. The number of bytes in the ICMP message
 - c. The sequence number
 - d. The TTL of the reply packet
19. Which documentation section describes the physical layout of network media?
 - a. Internetworking devices
 - b. Cable plant
 - c. Telecommunication closets
 - d. Servers

20. When using the Ping program to solve a network connection problem, which of these steps can you skip if the target host is on the same LAN?
- Ping the 127.0.0.1 address.
 - Ping the host's IP address
 - Ping the default gateway.
 - Ping DNS servers.

Challenge Labs



Challenge Lab 13-1: Troubleshooting Network Problems

Time Required: 1 hour or more

Objective: Troubleshoot a variety of network problems created by a partner.

Required Tools/Equipment: Your Windows client computer (such as one running Windows 7); a network share network that client computers should be able to open when their computers are in good working order

Description: This lab can be done in pairs, with each student creating a problem on his or her partner's computer. A working system should be able to access the Internet and a file share on the local network (set up by the instructor). While troubleshooting, students should keep a log stating the problem, symptoms, and the layers of the OSI model they test and the final solution to the problem. The following are some problems students can create (but they shouldn't be done in order):

- Set an incorrect static IP address.
- Set an incorrect subnet mask.
- Set an incorrect DNS server.
- Set an incorrect default gateway.
- Disable Client for Microsoft Networks in the network connection's properties.
- Disable the IPv4 protocol in the network connection's properties.
- Disconnect a cable and/or replace a patch cable with a crossover cable.
- Disconnect the switch from the router (if possible).
- Power off the switch (if possible).



Challenge Lab 13-2: Using a Protocol Analyzer to Troubleshoot DNS

Time Required: 30 minutes or more

Objective: Troubleshoot DNS by capturing the packets in a DNS query.

Required Tools/Equipment: Your Windows client computer (such as one running Windows 7) with Wireshark installed

Description: This lab can be done in groups. Capture packets involved in DNS queries, and determine the queries generated by your computer and the responses from the server. Next, capture only DNS packets, and use Nslookup to query *www.yahoo.com* and *www.yahoo-xyz.com*. Answer the following questions:

- What packet-capture filter did you use?

For each lookup, answer these questions:

- Did your computer generate more than one DNS query for each lookup? If so, why?

- What was the response from the DNS server for each query?

- Was more than one IP address returned? If so, why?



Challenge Lab 13-3: Using NetInfo to Gather Documentation

Time Required: 30 minutes or more

Objective: Use NetInfo to gather information about your network.

Required Tools/Equipment: Your Windows client computer with NetInfo installed

Description: This lab can be done in groups. Use NetInfo, which you installed in Chapter 10, to scan your network for a variety of information. Answer the following questions:

- What tools did you use in NetInfo to gather information?

- What type of information can you gather with NetInfo tools?

- How can the information you gathered be useful in documenting and supporting your network?



Challenge Lab 13-4: Installing an Advanced Network Monitor

Time Required: 45 minutes or more

Objective: Use WhatsUp Gold Network Monitor to gather network-monitoring information.

Required Tools/Equipment: Your Windows client computer (such as one running Windows 7) with Internet access

Description: This lab can be done in groups. Download WhatsUpGold from *www.ipswitch.com*. You need to enter some information about yourself before you can complete the download. While the file is downloading, read about some of the program's features and capabilities. After you install the program, start it and explore its many features. Then answer the following questions.

- What are three major features of WhatsUp Gold?

- How can these features help you to maintain and support a network?

Case Projects



Case Project 13-1

A user calls to report that she's unable to log on to e-mail, and you ask a couple of questions. Because you know that no one else is using the network right now, you can't determine whether the problem is unique to her machine or affects the entire network. Probing further, you learn that she's also unable to print. You decide this problem is probably easier to troubleshoot from the user's computer.

Using the eight-step problem-solving process covered in this chapter, outline the items you must check and the questions you must ask when you arrive at the user's office. Based on the possible responses to your questions, describe the actions you will take to correct potential causes of this problem.

Case Project 13-2

Document the computers, servers, and network equipment in your classroom. Design a form for gathering this information, including space for items such as model number, serial number, NIC type, MAC address, logical address, location, patch panel port connections, and hub/switch port connections. What other information might be important?

Case Project 13-3

Describe two network problems that can be solved by replacing 10 Mbps hubs and NICs with 100 Mbps switches and NICs. What must you verify about the existing network configuration before you perform the upgrade?



Common Networking Standards

The discussions and examples in this book and your own experiences in assembling new networks or working with existing ones have shown that a network is built of many parts. These components include the networking medium (cabled or wireless), network interfaces and supporting equipment, computers, connections, and many types of hardware. Other network parts include software drivers, networking protocols, networking services, application interfaces, and network-related programs. Given the complexity of assembling a network, it might seem miraculous that networks function at all.

What ensures their functioning is networking standards. At every level of the OSI model, network standards enable NIC developers to build NICs that attach to standard cables with standard connectors. Standards also permit e-mail vendors to count on basic delivery services built on the TCP/IP-based Simple Mail Transport Protocol (SMTP) and to get support for attachments of many kinds from Multipurpose Internet Mail Extensions (MIME).

If not for standards, networks couldn't work together effectively because vendors would have to work out the details of managing communication at many levels each time they try to solve a networking problem. Standards enable vendors to simplify assumptions about the way components behave, connect, and communicate on networks. This appendix examines the major networking standards, from the Physical layer all the way up to the Application layer, and gives pointers for additional information.

Standards-Making Process

Committees create most standards because this process involves many groups, each with its own special interests and agendas. Therefore, standards invariably involve compromises and alternatives that exceed practicality and technology. Nevertheless, it's possible to describe the general standards-making process. Most standards setting occurs within the framework of standards-setting bodies, industry associations, trade groups, or other organizations that consist mainly of unpaid volunteers, with a small core of paid professionals. As general members propose "hot topics" or specific networking needs germane to the umbrella group, or if ideas come from other possible channels, special interest groups (SIGs) form. SIGs include representatives from governments, vendors, academia, the consulting community, and user groups (especially large and well-funded ones that can afford staff to participate in this endeavor). Sometimes representatives from particular factions play pivotal roles in these groups.

In a SIG, working groups coalesce around certain topics. Each group selects a chairperson and appoints members of the working group, who then address the problems related to that group's focus and discuss ideas related to the topic. In a working group, constituencies usually propose ideas, which invariably start out based on proprietary technologies or idiosyncratic viewpoints. As a proposal takes shape, it broadens as members of different constituencies work to ensure that their viewpoints are addressed. Over time, these groups work hard to achieve consensus, which emerges from a long series of rough draft proposals that are amended until the SIG is ready to submit a rough draft for outside review. This process can—and sometimes does—take years, but 3 to 9 months is typical. Even so, many rough drafts never go beyond this step because the groups can't reach consensus or because newer technologies emerge and supplant their proposals.

The rough draft is submitted to the SIG for further discussion and approval, followed by another series of drafts and rewrites. Perhaps the entire SIG reaches consensus that the proposal is worthy of draft status; otherwise, the proposal is abandoned. Again, this process usually takes 3 to 9 months. The SIG then submits the draft to its parent group or the body of the entire organization for more discussion and another approval process, which results in acceptance or rejection. This process takes another 3 to 6 months. If the entire membership accepts a proposal, it's published as an official standard for that group after it's submitted in a final, approved form. This step can take from several weeks to several months, depending on the proposal document's size and the remaining work it requires.

Official standards must be reviewed on a regular cycle and amended as needed. Champions or key proponents from the original working group usually take stewardship of standards and maintain their currency and accuracy. Typically, reviews occur yearly or twice yearly (if the membership doesn't call for earlier review) and can take from 1 to 6 months.

A standard becomes obsolete when the organization designates it as such. This designation usually means the organization approved a newer standard to take its place or a subsequent revision involves so much change that the preceding version becomes obsolete, and its replacement is designated as the new official standard.

Clearly, this process is complex and labor intensive. The delay inherent in any consensual process explains why proprietary technologies and approaches to networking (among other fields) continue to play an important role in business and industry. New and improved proprietary technologies keep the pressure on standards makers to deliver usable results as quickly as possible and provide a never-ending stream of alternatives. Among the hundreds of industry consortia, trade groups, professional associations and societies, and SIGs in the networking community, only a small number manage the standards with the most influence on networking hardware and software. The following sections describe the most important standards makers in networking.

Important Standards Bodies

Standards come from many sources, some more influential than others. Most standards bodies discussed in this appendix exert considerable influence around the world. Some focus more on hardware and signaling issues; others are concerned more with software. Be aware that many more standards bodies exist than are covered in this appendix. Familiarize yourself, at a minimum, with the main groups and their networking standards and technologies.



TIP

For an outstanding online reference on networking standards, visit www.cmpcmm.com/cc/standards.html.

Here's a list of the most important standards-setting bodies, described in more detail in the following sections:

- American National Standards Institute (ANSI)
- Comité Consultatif International Téléphonique et Télégraphique (CCITT)
- Electronic Industries Alliance (EIA)
- Internet Architecture Board (IAB)
- Institute of Electrical and Electronics Engineers, Inc. (IEEE)
- International Organization for Standardization (ISO)
- Object Management Group (OMG)
- The Open Group (TOG)
- World Wide Web Consortium (W3C)
- Internet Corporation for Assigned Names and Numbers (ICANN)

American National Standards Institute

ANSI creates and publishes standards for programming languages, communication methods and techniques, and networking technologies. It's also the U.S. representative to ISO, the main international standards-setting body for networking and wireless communications, and to the CCITT, the main international standards-setting body for telephony and long-haul digital communication.

ANSI programming languages include C, COBOL, and FORTRAN as well as a version of Structured Query Language (SQL) commonly used in database access and programming. ANSI standards also cover the small computer systems interface (SCSI) used for high-speed, high-capacity disk drives and other microcomputer peripheral devices. A standard PC device driver, *Ansi.sys*, used to drive character-mode screen displays in DOS (and DOS emulation modes) is commonly found on PCs.

Major ANSI specifications include the following:

- *ANSI 802.1-1985/IEEE 802.5*—Token ring access, protocols, wiring, and interfaces
- *ANSI/IEEE 802.3*—Coaxial cable standards, CSMA/CD definition for Ethernet
- *ANSI X3.135*—SQL database query methods for client/server database access
- *ANSI X3.92*—Privacy/security encryption algorithm for network use
- *ANSI X3T9.5*—FDDI specification for voice and data transmission
- *SONET*—Fiber-optic specification for transmitting time-sensitive data (such as real-time video) across a global network



For more information on ANSI standards, visit ANSI's Web site at www.ansi.org.

Comité Consultatif International Téléphonique et Télégraphique

CCITT (in English, the Consultative Committee for International Telegraphy and Telephony) is a permanent subcommittee of the International Telecommunications Union (ITU), an organization operating under the auspices of the United Nations. This committee's parent body includes representatives from 160 countries, mostly from national Postal, Telephone, and Telegraph (PTT) services. CCITT is responsible for many standards that apply to communication, telecommunication, and networking, including X.25 packet-switched networks, X.400 electronic messaging systems, X.500 directory services, encryption and security, the V.*nnn* standards for modems, and the L.*nnnn* standards for ISDN. (In these generic standards designators, *nn* and *nnn* stand for sequences of two or three digits.)

Because CCITT works closely with ISO, many standards carry designations from both groups. CCITT recommendations appear once every 4 years, most recently in 2004. In March 1993, CCITT was officially renamed the International Telecommunication Standardization Sector (ITU-T, sometimes called ITU-TS or ITU-TSS), but nearly all resources still refer to this organization by its original name. CCITT includes the following study groups:

- A, B—Working terms, definitions, and procedures
- D, E—Tariffs
- F—Telegraph, telemetric, and mobile services
- G, H—Transmissions
- I—ISDN
- J—Television transmission
- K, L—Facilities protection
- M, N—Maintenance
- P—Telephone transmission
- R–U—Terminal and telegraph services
- V—Telephone-based data communication
- X—Data communication networks

The V-series modem and teledata communication standards are as follows:

- V.22—1200 bps full-duplex modem
- V.22*bis*—2400 bps full-duplex modem
- V.27—Fax/modem communication
- V.28—RS-232 interface circuits
- V.32—Asynchronous and synchronous 4800/9600 bps

- *V.32bis*—Asynchronous and synchronous up to 14.4 Kbps
- *V.35*—High data-rate communication across combined circuits
- *V.42*—Error checking
- *V.42bis*—Lempel-Ziv data compression for modems
- *V.90*—Modem standard for 56 Kbps downstream, 33.6 Kbps upstream

The X-series standards, which overlap with OSI standards, include the following:

- *X.200 (ISO 7498)*—OSI reference model
- *X.25 (ISO 7776)*—Packet-switching network interface
- *X.400 (ISO 10021)*—Message handling
- *X.500 (ISO 9594)*—Directory services, security, and encryption
- *X.700 (ISO 9595)*—Common Management Information Protocol (CMIP)



For more information on CCITT standards, visit www.itu.int/home/index.html.

Electronic Industries Alliance

The EIA (www.eia.org), founded in the 1920s, is an industry trade organization of U.S. manufacturers of electronic components, parts, and equipment. It supports a large library of technical documents (many available online), including standards for interfaces between computers and communications equipment. The EIA works closely with other standards organizations, including ANSI and CCITT. Many EIA standards have CCITT counterparts, so EIA RS-232 is the same as CCITT V.24. The EIA's best-known standards are those for serial interface connections, particularly connections between computers and modems:

- *RS-232*—Defines serial connections for modems, including DB-9 and DB-25 connectors
- *RS-422*—Defines a balanced multipoint interface, commonly used for data acquisition
- *RS-423*—Defines an unbalanced digital interface, also used for data acquisition
- *RS-449*—Defines a serial data interface with DB-37 connectors that specifies RS-422 and RS-423 as subsets of its capabilities

Internet Architecture Board

The IAB is the board governing the Internet and the parent body for the many other boards overseeing Internet protocols, technologies, research, and development. It can be considered the primary controlling authority over Internet standards, but no single body controls the Internet. The IAB is part of the Internet Society, a general membership organization for people interested in Internet technologies and related social issues. (Visit www.isoc.org for information on joining.) The following are some important IAB constituent bodies:

- *Internet Engineering Task Force (IETF; www.ietf.org)*—The group under the IAB that develops, approves, and maintains standards defining valid Internet protocols, services,

and related information. It manages RFC documents defining draft, experimental, proposed, historical, and official Internet standards.

- *Internet Network Information Center (InterNIC; www.internic.net)*—Responsible for providing information on Internet domain registration services. InterNIC currently contracts this function to third parties worldwide. The InterNIC Web site has a database you can check to see whether another party already has the name you want to register, a form to report a problem with a registrar, and a FAQ with answers about domain registration.
- *Internet Corporation for Assigned Names and Numbers (ICANN)*—Responsible for managing the Internet’s IP address space as well as related domains and domain names. It’s also responsible for doling out IP addresses—typically to ISPs, who then allocate them to customers. ICANN took over responsibility for the Internet Assigned Numbers Authority (IANA). Its Web site is at www.icann.org, but you can still find information at IANA’s Web site (www.iana.org).
- *Internet Engineering Steering Group (IESG)*—Executive group that guides activities of the IETF’s many constituent elements.
- *Internet Research Task Force (IRTF)*—Works on long-term research proposals, new technologies, privacy and security issues, and other aspects of proposed Internet technologies with social as well as technical implications.

The number and nature of Internet standards is too vast a subject for this appendix. One RFC provides a map to all other current, valid RFCs. At this writing, it’s RFC 2500. Titled “Internet Official Protocol Standards,” it summarizes all the current official Internet standards. You can also search for the most recent version with the RFC-Full Text Search engine at www.faqs.org/faqs/. If you search for RFC 2500, you should be able to find any older versions because new RFCs always list obsolete versions of the RFCs they replaced.

Institute of Electrical and Electronics Engineers, Inc.

The IEEE (www.ieee.org) is a U.S.-based professional society that publishes many technical standards, including networking-related standards. The IEEE’s 802 Committee developed some of the most important LAN standards in use today. After the IEEE finishes work on a standard, it usually shares the work with ANSI, which might then forward it to the ISO. This process explains why several elements of the IEEE 802 standards family are also ANSI and ISO standards.

Several working committees were formed as part of the IEEE 802 project because no single group was capable of handling the many topics involved in this mammoth undertaking. The following committees were created to cover the full range of topics in the 802 project:

- 802.1—Internetworking
- 802.2—Logical Link Control (LLC)
- 802.3—CSMA/CD Network (Ethernet)
- 802.4—Token Bus Network
- 802.5—Token Ring Network
- 802.6—Metropolitan Area Network (MAN)

- 802.7—Broadband Technical Advisory Group
- 802.8—Fiber-optic Technical Advisory Group
- 802.9—Integrated Voice/Data Networks
- 802.10—Network Security
- 802.11—Wireless Networks
- 802.12—High-Speed Networking
- 802.13—Unused
- 802.14—A now-defunct working group that specified data transports over cable TV
- 802.15—Wireless PAN (covers emerging standards for wireless personal area networks)
- 802.16—Wireless MAN
- 802.17—Resilient Packet Ring (covers emerging standards for very high-speed, ring-based LANs and MANs)
- 802.18—Wireless Advisory Group (monitors radio-based wireless standards)
- 802.19—Coexistence Advisory Group (addresses issues of coexistence with current and developing standards)
- 802.20—Mobile Broadband Wireless (addresses always-on multivendor mobile broadband wireless access)

International Organization for Standardization

The ISO focuses on defining global-level standards. Member countries are represented by government bodies or their main standards-setting bodies. (For example, ANSI represents the United States, and the British Standards Institute represents Great Britain.) The ISO also includes representatives from businesses, educational institutions, research and development organizations, and other international standards bodies, such as CCITT. ISO's overall charter is broad—establishing international standards for all services and manufactured goods and products.

For computing, ISO seeks to establish global standards for data communication and information exchange. These standards are intended to promote interoperability across networking environments worldwide and to allow mixing vendor systems and products without regard to system type or country of origin. The ISO's primary efforts in interoperability are directed at the Open Systems Interconnection initiative (OSI or ISO/OSI). You can find an overview of important OSI standards at www.iso.org/iso/home.htm.

Object Management Group

The OMG (www.omg.org) represents a federation of more than 700 member organizations from business, industry, government, and academia involved in devising tools that enable system vendors to create platform- and OS-neutral applications. The OMG's efforts extend to programming and scripting languages, application and data-conversion interfaces, and protocols. For a fee, the OMG offers certification services to verify that products conform to standards and specifications agreed on by OMG member organizations.

The cornerstone of the OMG's efforts is its Object Management Architecture (OMA), which defines a common model for object-oriented applications and runtime environments. A key element of the OMG's efforts focuses on Common Object Request Broker Architecture (CORBA), a set of standard interfaces and access methods that allow interchanging objects and data across a wide variety of platforms and OSs. In addition, The Open Group (described in the next section) incorporates the OMG's architecture in its Distributed Computing Environment and Distributed Management Environment.

The Open Group

The Open Group (www.opengroup.org) formed in February 1996 by consolidating two leading open systems consortia—the X/Open Company Limited (X/Open) and the Open Software Foundation (OSF). Under the TOG umbrella, OSF and X/Open work together to deliver technology innovations and promote wide-scale adoption of open systems specifications. Founded in 1988, the OSF hosts industry-wide, collaborative software research and development for distributed computing. Founded in 1984, X/Open's brand is recognized worldwide as a guarantee of compliance with open systems specifications and now includes ownership of the UNIX trade name.

TOG focuses on defining vendor-neutral computing and development environments with a special emphasis on user interfaces. Its legacy from the OSF includes the following main elements:

- *Distributed Computing Environment (DCE)*—Simplifies development of software for use in heterogeneous networked environments.
- *Distributed Management Environment (DME)*—Defines tools to manage systems in distributed, heterogeneous computing environments.
- *Single UNIX Specification*—Defines a common reference model for an advanced UNIX implementation, with support for SMP, enhanced security, and dynamic configuration.
- *X-Window System*—Provides a well-recognized industry standard model for a platform-neutral GUI.
- *Motif Toolkit API*—A well-recognized industry standard for a common user interface definition that recognizes IBM's Common User Access (CUA) model.
- *Network File System (NFS)*—Defines a well-accepted standard model for a UPD/IP-based distributed file system.
- *Common Desktop Environment (CDE)*—Offers tools for building client-side application front ends. Its current release integrates the Motif 2.0 GUI, the X-Window System, and common application interfaces for standardizing application displays across distributed multiplatform environments.
- *Baseline Security Services (XBSS) and Secure Communication Services (GSS-API)*—XBSS defines basic security-related functions to be provided by open systems with recommended default settings for security parameters; GSS-API is an application programming interface that ensures secure communication when interacting with peer applications across a network.
- *Structured Query Language (SQL) Definitions and Services*—Defines application access to relational databases, using SQL embedded in C and/or COBOL. TOG's

XPG4 SQL includes dynamic SQL, which corresponds to ISO/IEC 9075:1992. XPG4 SQL also includes specifications developed with the SQL Access Group that allow application portability to distributed environments.

The World Wide Web Consortium

Founded in the early 1990s in the wake of CERN's decision to release its work on HTML and HTTP to the public, the W3C is the standards-setting body for Web markup languages, specifications, accessibility guidelines, and more. Organizations such as Massachusetts Institute of Technology (MIT) in the United States and INRIA (Institut National de Recherche en Informatique et en Automatique, the French National Institute for Research in Computer Science and Control) are involved in staffing and housing this organization.

At first, a Web-oriented standards group might not seem to have much to do with networking, but the importance of the Web in finding networking information can't be overstated. Web-based services are also essential for in-house networks (intranets) and the public Internet. Savvy network administrators must know how to use the Web and manage Web servers on their networks. Key W3C standards include the following:

- *HTML and Extensible HTML (XHTML)*—Used to create many Web pages
- *Extensible Markup Language (XML)*—The basis for XHTML and other XML-based applications
- *HTTP*—Transports Web page requests from clients to servers and responses from servers to clients
- *Accessibility guidelines*—Developed to make the Web equally available to all users, regardless of visual or reading disabilities
- *Cascading Style Sheets (CSS)*—Provide detailed instructions on how to display content from HTML and XML documents
- *XML-based applications*—More than 30 XML-based applications for everything from mathematical notation to wireless phone access to Web data



For more information on these and other standards, tools, and best practices, visit the W3C's Web site at www.w3.org.

Internet Corporation for Assigned Names and Numbers

ICANN is more of a controlling and organizing body for the names (domain names) and numbers (IP addresses and port numbers) used to access Internet resources. It oversees distribution of top-level domains (such as .com, .org, and .net) and helps manage the distributed DNS system, which works to ensure that Internet resources can be located by using names rather than numbers. ICANN is also responsible for distribution and management of the IP address space and assignment of TCP and UDP port numbers used to identify Application-layer protocols. As mentioned, it's taken over IANA's responsibilities. You can read more about ICANN and IANA at www.icann.org and www.iana.org.

This page intentionally left blank



Older and Obsolete Technologies

This appendix contains information on older or obsolete technologies that were included in previous editions of this book. Details of these technologies were removed in chapter coverage to make room for current topics but might still be important for some readers.

Thinwire Ethernet (Thinnet)

Thinwire Ethernet is a thin, flexible cable approximately 0.25 inches (0.64 cm) in diameter. Thinwire cabling is easy to work with and fairly inexpensive to build or buy. (Prefabricated cables in many lengths are widely available.) It's especially well suited for small or constantly changing networks. With BNC T-connectors, thinwire cables attach more or less directly to networking devices and to each computer's NIC.

Working with the U.S. military, cable manufacturers designated Radio Government (RG) specifications for various types of cable, including many varieties of coaxial. Thinnet belongs to the RG-58 family and has a characteristic impedance of 50 ohms. (Impedance, measured in ohms, is the electrical resistance to current flowing in this type of cable.) The main differences between members of the RG-58 family lie in the center conductor. For some members, this conductor is solid wire; in others, it has a braided core. Table B-1 compares some types of RG cable.

Table B-1 Well-known types of RG cable

Designation	Type	Impedance	Description
RG-58/U	Thinwire	50 ohms	Solid copper core (U stands for utility grade; not recognized as valid thinwire cable by IEEE 802.3)
RG-58 A/U	Thinwire	50 ohms	Stranded copper core (A/U indicates a tinned copper braid as the center conductor with foam dielectric insulator)
RG-58 C/U	Thinwire	50 ohms	Military version of RG-58 A/U (uses a solid dielectric insulator)
RG-59	CATV	75 ohms	Broadband cable; used for cable TV and sometimes ARCnet

Table B-1 Well-known types of RG cable (*continued*)

Designation	Type	Impedance	Description
RG-6	Broadband	75 ohms	Larger diameter and higher bandwidth than RG-59; used as CATV drop cable
RG-62	Baseband	93 ohms	Used for ARCnet and IBM 3270 terminals
RG-8	Thickwire	50 ohms	Solid core; approximately 0.4 inches in diameter
RG-11	Thickwire	75 ohms	Stranded core, approximately 0.4 inches in diameter; used for CATV trunk lines

Table B-2 summarizes the key characteristics of thinwire Ethernet cable.

Table B-2 Thinwire Ethernet characteristics

Characteristic	Value
Maximum cable length	185 m (607 ft.)
Bandwidth	10 Mbps
Bend radius	360 degrees/ft.
Installation/maintenance	Easy to install and reroute; flexible
Cost	Cheapest form of coaxial cable; prefabricated cables average \$1/ft.
Connector type	British Naval Connector (BNC)
Security	Susceptible to eavesdropping
Interference rating	Good: lower than thicknet, higher than twisted pair



Research shows numerous names for the BNC acronym for thinwire Ethernet and thicknet connectors, including British Naval Connector (preferred Microsoft use), bayonet nut connector, bayonet navy connector, and bayonet Neill-Concelman.

Thickwire Ethernet (Thicknet)

Thickwire Ethernet is a rigid coaxial cable about 0.4 inches (approximately 1 cm) in diameter. It often has a bright-yellow Teflon coating and is commonly described as “frozen yellow garden hose,” which accurately conveys its rigidity. Thicknet is sometimes described as Standard Ethernet because it was the first type of cable used for this networking technology. However, its expense and lack of ductility, or flexibility, have made it the least commonly used type of Ethernet cable.

Thickwire’s increased diameter does offer some advantages: better resistance to interference and better conductivity. It also means a longer maximum cable segment length and an increase in the number of devices that can be attached to a single segment. Thickwire’s capability to carry signals over longer distances, coupled with its superior interference resistance, help explain why this cable is most commonly used for backbones—heavy-duty, long-run cables—that interconnect smaller thinnet or twisted-pair network segments.

Thinwire Ethernet cables connect directly to network interfaces, but attaching to thickwire Ethernet takes a different approach. For thickwire, a device called a “vampire tap” is usually used to attach a device to the cable, which in turn attaches to a transceiver. The tap must be installed carefully because a hole has to be drilled into the wire, which can result in a short. The transceiver then attaches to a drop or transceiver cable that plugs into an attachment unit interface (AUI) on the computer’s NIC or on other devices to attach to the network.

Transceiver cables can be up to 50 meters long (approximately 164 feet), so there can be some latitude when running thickwire cable; its path doesn’t have to snake from system to system. Thinwire, however, must go from system to system because the network cable attaches directly to the network interface on the computer or other device. As long as the distance between the cable and computer remains under 50 meters, thickwire Ethernet requires less network cable than thinwire. On the other hand, the necessary transceivers and transceiver cables make thickwire more expensive than thinwire. The increased expense of using thickwire, its larger diameter, and its lack of flexibility explain why it’s rarely used now for new network installations. Table B-3 summarizes the characteristics of thickwire.

Table B-3 Thickwire Ethernet characteristics

Characteristic	Value
Maximum cable length	500 m (1640 ft.)
Bandwidth	10 Mbps
Bend radius	30 degrees/ft.
Installation/maintenance	Hard to install and reroute; rigid
Cost	More expensive than thinwire; cheaper than fiber
Connector type	BNC
Security	Susceptible to eavesdropping
Interference rating	Good: lowest of all electrical cable types

All types of Ethernet coaxial cable have an additional requirement to create a working network. A connector (a female BNC for thinwire and thickwire) must cap each end of a cable, and a terminator must screw into each end connector. Terminators “soak up” signals that arrive at the end of the cable; otherwise, they would bounce and reflect up the cable, interfering with network traffic. Without correct termination, a coax-based Ethernet network can’t work.

The two features that make coaxial cable an attractive medium are its capability to carry signals a long distance and its resistance to interference. However, its low bandwidth capability, coupled with its expense, make coaxial cable obsolete in LAN applications. Twisted-pair, fiber-optic, and wireless media rule in most networks now.

10Base5 Ethernet

10Base5 uses transceivers attached to thicknet by a vampire tap. When vampire taps are installed, a special jig fixture is used to drill through the covering and mesh, and the tap makes contact only with the center conductor. A vampire tap has small teeth that keep the

tap/transceiver from moving after it's installed. A drop cable connects the transceiver to the NIC's AUI or DIX port (standard Ethernet connectors). Each computer connected to the thicknet cable must have a transceiver and drop cable.

The distance limitations for 10Base5 Ethernet are more stringent than for other Ethernet types. Transceivers must be at least 2.5 meters (about 8 feet) apart. Each cable segment can be a maximum length of 500 meters (1640 feet). Up to five cable segments can be attached by using repeaters, creating a network with a total length of 2500 meters; the drop cable connecting the computer to the transceiver must be less than 50 meters (164 feet). However, the length of drop cables isn't figured into the total network length.

All coaxial Ethernet networks (10Base5 and 10Base2) are subject to the 5-4-3 rule, which states that a coaxial Ethernet network can consist of a maximum of five segments, with four repeaters, with devices attached to three of the segments. This configuration prevents signal loss caused by attenuation.

Note that the 5-4-3 rule is an "end to end" rule, not a "total population" rule. The difference has to do with how many segments and repeaters are used between any two machines, instead of stipulating the total number of elements in a network. This rule applies only to single pairs of segments, when a node on one segment seeks to transmit data to a node on another segment. Therefore, the 5-4-3 rule doesn't mean a 10Base5 or 10Base2 network can have only five segments, four repeaters, and so forth in total; this rule applies only when tracing a route from a node on one segment to a node on another segment. Plenty of networks with hundreds of segments and numerous repeaters don't violate the 5-4-3 rule because they're designed not to.

As mentioned, 10Base5 networks represent the original Ethernet architecture. However, 10Base5's limitations and the difficulties of working with thicknet cable have rendered it obsolete. Table B-4 lists the specifics for 10Base5 Ethernet.

Table B-4 10Base5 Ethernet summary

Category	Specification
IEEE specification	802.3
Advantages	Long maximum cable length
Disadvantages	Difficult to install; cost
Topology	Linear bus
Cable type	50 ohm thicknet
Channel access method	CSMA/CD
Transceiver location	Connected to cable at vampire tap
Maximum cable segment length	500 m (1640 ft.)
Maximum total network length	2500 m (8200 ft.)
Maximum drop cable length	50 m (164 ft.)
Minimum distance between transceivers	2.5 m (8 ft.)
Maximum number of segments	Five connected by four repeaters

Table B-4 10Base5 Ethernet summary (continued)

Category	Specification
Maximum number of populated segments	Three
Maximum devices per segment	100
Maximum devices per network	1024
Transmission speed	10 Mbps

10Base2 Ethernet

10Base2 was the second version of Ethernet. Following the IEEE naming convention, you would think 10Base2 could support a single 200-meter cable segment. The original IEEE specification for 10Base2 did permit a 200-meter cable segment, but that distance was shortened to 185 meters to improve performance and to account for patch cables. Like 10Base5, 10Base2 uses coaxial cable, but instead of thicknet, it uses thinnet, which is flexible and easier to manipulate. Also, unlike 10Base5, the transceiver is part of the NIC, so the cable attaches directly to the device. 10Base2 uses a BNC connector to connect the NIC to the cable and uses the bus topology with terminators at each end of the cable segment. The minimum cable length for 10Base2 is .5 meters (about 20 inches).

Although thinnet cable looks remarkably like the coaxial cable used for television, these two types aren't interchangeable. The IEEE specification states that thinnet must use RG-58A/U or RG-58C/U. Thinnet uses 50 ohm coaxial cable; the cable used for cable TV is 75 ohm cable. In addition, other RG-58 cable types (RG-58U, for example) can't support Ethernet 10Base2.

Like its 10Base5 predecessor, 10Base2 follows the 5-4-3 rule. The 10Base2 limitations on cable length from end to end on a network allow for five 185-meter segments connected by four repeaters, with three segments populated. This creates a maximum total network length of 925 meters from any one end to any other end of the network. In each of those five cable segments, a BNC barrel connector can be used to connect two shorter thinnet cables. Its use should be limited, however, because each barrel connector degrades the signal as it travels across the network. 10Base2 supports up to 30 devices per cable segment.

Because of its ease of installation and lower price, thinnet rapidly replaced thicknet as the preferred network medium. As new Ethernet standards developed, thinnet, too, was eventually replaced. Table B-5 summarizes the 10Base2 Ethernet standard.

Table B-5 10Base2 Ethernet summary

Category	Specification
IEEE specification	802.3
Advantages	Inexpensive; easy to install and configure
Disadvantages	Difficult to troubleshoot
Topology	Linear bus
Cable type	50 ohm thinnet (RG-58A/U or RG-58C/U)
Channel access method	CSMA/CD

Table B-5 10Base2 Ethernet summary (continued)

Category	Specification
Cable type	50 ohm thinnet (RG-58 A/V or RG-58v C/V)
Maximum cable segment length	185 m (607 ft.)
Maximum total network length, end to end	925 m (3035 ft.)
Minimum distance between devices	.5 m (20 in.)
Maximum number of segments	Five connected by four repeaters
Maximum number of populated segments	Three
Maximum devices per segment	30
Maximum devices per network	1024
Transmission speed	10 Mbps

100VG-AnyLAN

The Ethernet Standard 100VG-AnyLAN—also called 100BaseVG, 100VG, VG, or AnyLAN—was developed by Hewlett-Packard and AT&T. It combines elements of Ethernet and token ring architectures and uses a demand priority channel access method, in which intelligent hubs control network communication. When a computer has data to transmit, it sends a demand packet to the hub, which then tells the computer when the channel is free to send its data. These hubs can cascade, much like 10BaseT, creating a star topology network. A root hub or parent hub connects to multiple hubs, each of which can connect to other hubs.

100VG-AnyLAN is designed to run over any data-grade UTP cable and can be used with Category 3 or higher. In older facilities, it's an attractive option because existing cabling can be reused. However, one caveat is that 100VG requires all four pairs or wires in a typical UTP cable (two to transmit and two to receive), whereas 10BaseT uses only two pairs. In some existing 10BaseT installations, two pairs of wires in a cable can be used for data, and the other two pairs can be used for voice. In this case, the cabling must be upgraded before 100VG-AnyLAN can work (It might make even more sense to use a different, less expensive 100 Mbps Ethernet technology).

The biggest limitation of 100VG-AnyLAN is cost. The requirements for special NICs and hubs for demand priority channel access and for all four pairs of wires in UTP make costs much higher than with other 100 Mbps Ethernet implementations. Because of these disadvantages, 100VG-AnyLAN is a networking technology that came and went without making much impact on the networking world. Table B-6 summarizes this standard.

Table B-6 100VG-AnyLAN summary

Category	Specification
IEEE specification	802.12
Advantages	Fast; easy to configure and troubleshoot; supports token ring and Ethernet packets
Disadvantages	High cost; limited distance over UTP
Topology	Star
Cable type	Cat 3 or higher UTP and STP, fiber-optic
Channel access method	Demand priority
Transceiver location	On NIC
Maximum cable segment length	100 m (328 ft.) Cat 3 UTP; 150 m (492 ft.) Cat 5 UTP; 2000 m (6561 ft.) fiber-optic
Maximum number of segments	1023
Maximum devices per segment	One
Maximum devices per network	1024
Transmission speed	100 Mbps

Ethernet Frame Types

One major distinction between Ethernet and other network architectures is that Ethernet can structure data several different ways before placing it on the network medium. As discussed in Chapter 3, a computer places data on the network in frames, which define the data's structure. Ethernet supports four frame types, and these frame types don't work with each other. For communication to take place between Ethernet devices, their frame type settings must match. The Ethernet frame types are as follows and are discussed in more detail in the following sections:

- Ethernet 802.3 is generally used by IPX/SPX on Novell NetWare 2.x and 3.x networks.
- Ethernet 802.2 is the default frame type used by IPX/SPX on Novell NetWare 3.12 and 4.x networks. It's also the native frame type supported by default when Microsoft NWLink is installed.
- Ethernet SNAP is used in EtherTalk and mainframe environments.
- Ethernet II is used by TCP/IP and is discussed in Chapter 3.

All Ethernet frame types support a packet size between 64 and 1518 bytes and can be used by all network architectures mentioned previously. In most cases, a network requires only one frame type, but occasionally devices, such as file or database servers, must support

multiple frame types (for instance, when some clients use one frame type, but other clients use another).

When running a protocol that can use more than one frame type, such as IPX/SPX, there must be a method to select the frame type. In Windows, this selection is made in the Local Area Connection Properties dialog box. Windows defaults to auto-detection of the frame type, but this setting can cause undesirable results because non-server versions of Windows support only the first frame type detected. If resources on the network use different frame types, some resources aren't available to workstations that auto-detect a different frame type. Therefore, ensuring that only a single frame type is used or all workstations have been set to a common frame type is essential for allowing access to resources.



Always remember that mismatched frame types prevent network communication.

TIP

Ethernet 802.3 Frame Type

Sometimes called “Ethernet raw,” the Ethernet 802.3 frame type was developed before the IEEE 802.3 specification was completed. Therefore, the 802.3 frame type doesn't completely comply with the 802.3 specification, despite its name. Generally, Ethernet 802.3 frames occur only on Novell NetWare 2.x or 3.x networks.

The Ethernet 802.3 frame begins with a preamble and a start frame delimiter (SFD) statement, which indicates the beginning of the frame. The frame's destination and source addresses follow. Because Ethernet supports variable-length frames (64 to 1518 bytes), the next field specifies the length of the frame's data section. Then a 4-byte CRC follows to verify that the data reached its destination undamaged.

Ethernet 802.2

Ethernet 802.2 frames comply completely with the Ethernet 802.3 standard. The IEEE 802.2 group didn't address Ethernet, only the Logical Link Control (LLC) sublayer of the OSI model's Data Link layer. However, because Novell had already decided to use the term Ethernet 802.3 to describe Ethernet raw, it's generally accepted that Ethernet 802.2 means a fully 802.3- and 802.2-compliant Ethernet frame. Ethernet 802.2 frames contain fields similar to those in 802.3, with three additional LLC fields.

Ethernet SNAP

Ethernet SubNetwork Address Protocol (SNAP) is generally used on the AppleTalk Phase 2 networks discussed later in “The AppleTalk Environment.” It contains enhancements to the 802.2 frame, including a protocol type field, which indicates the network protocol used in the frame's data section.

The Token Ring Architecture

Developed by IBM in the mid-1980s, the token ring network architecture provides fast, reliable transport. Based on the IEEE 802.5 standard, token ring networks are cabled in a physical star topology but function as a logical ring. The token-passing channel access method,

rather than the network's physical layout, gives token ring its name. The original version of token ring operated at 4 Mbps, but later versions increased this speed to 16 Mbps.



If a 4 Mbps NIC is used in any workstation in an otherwise 16 Mbps token ring network, the entire network operates at 4 Mbps.

Token Ring Function

By using the token-passing access method, token ring networks ensure that all computers get equal time on the network. A small frame, called a token, passes around the ring. A computer receives the token from its nearest active upstream neighbor (NAUN). If the token isn't in use at the time—no nearby computer is sending data—and the computer has data to send, it attaches data to the token and sends it to its nearest active downstream neighbor (NADN). Each computer thereafter receives the token, determines that the token is in use, and verifies that it's not the data's destination station. If not, the computer re-creates the token and the data exactly as it received them and sends them to its NADN.

When data reaches its destination, the receiving computer sends the data to the upper-layer protocols (Network, Transport, Session, Presentation, and Application layers) for processing. Then the receiving computer toggles two bits in the data packet to indicate it received the data and sends the token and data along the network to its NADN. Eventually, both token and data reach the original sender; the sender sees that the data was received successfully, frees the token, and then passes it along.

Beaconing

One unique aspect of the token ring network architecture is its capability to isolate faults automatically by using a process called “beaconing.” The first computer powered on in a token ring network is assigned the responsibility of ensuring that data can travel along the ring. This computer, the active monitor, manages the beaconing process. All other computers on the network are standby monitors.

Every 7 seconds, the active monitor sends a special packet to its NADN announcing the address of the active monitor and the fact that it's the upstream neighbor. The station examines the packet and passes it along to its NADN, changing the upstream address. The third station then has a packet listing the active monitor's address and the address of its upstream neighbor. The third station repeats the process, sending to its NADN a packet containing the active monitor's address and its own address. When the active monitor receives the packet, it knows that the packet has navigated the ring successfully and the ring is intact. In addition, all stations know the address of their upstream neighbor.

If a station doesn't hear from its upstream neighbor in 7 seconds, it sends a packet down the ring containing its address, the address of its NAUN (from which it received no packet), and a beacon type. As the other computers in the network receive this packet, they check their configurations. If the NAUN doesn't answer, the ring can reconfigure itself to avoid the problem area. Beaconing allows some level of automatic fault tolerance in a token ring network, something many other network architectures lack.

Although this process seems laborious, it's fairly efficient. Unlike Ethernet, there are no collisions, so data seldom has to be sent again, and much larger data packets can be

sent—between 4000 and 17,800 bytes. Because all computers on the network have equal access to the token, traffic is consistent, and token ring handles increases in network size and bandwidth use efficiently. Another advantage of token ring is that it's considered deterministic, meaning you can determine the maximum time a device must wait between accesses to the token and, therefore, the network medium.

Hardware Components

In a token ring network, a hub can be referred to as a multistation access unit (MSAU) or smart multistation access unit (SMAU). IBM's implementation of token ring is the most popular adaptation of the IEEE 802.5 standard. Although there are some minor differences in IBM and IEEE specifications, such as the maximum number of computers on an STP ring, they're very similar. When discussing hardware components of the token ring architecture, IBM equipment is most often used.

A typical IBM token ring hub, such as the 8228 MSAU, has 10 connection ports, 8 of which can be used for connecting computers. The other two ports are used to connect the hubs in a ring. The ring out (RO) port on one hub connects to the ring in (RI) port on the next hub, and so on to form a ring among the hubs. New hubs must also be added to the ring in this manner. IBM's implementation of token ring allows connecting 33 hubs in this fashion. The original token ring hubs allowed a total of 260 stations per network. However, newer hubs that allow 16 computers per hub double this number.

Cabling in a Token Ring Environment In 1984, IBM defined a comprehensive cabling system that specified cable types, connectors, and all other components required for networking. This cabling system categorizes cables based on the American Wire Gauge (AWG) standards that specify wire diameters. When token ring was introduced, it followed these standards for cabling and equipment. Table B-7 shows the cable types included in the IBM system and used by token ring.

Table B-7 IBM/token ring cabling

Cable type	Description
Type 1	STP with two pairs of 22-AWG solid copper wire surrounded by a braided shield and casing. This cable is used to connect computers to MSAUs and can be run through conduit or inside walls.
Type 2	STP with two pairs of 22-AWG solid copper wire for data and four pairs of 26-AWG wire for voice. This cable is used to connect both data and voice without running two cables.
Type 3	UTP voicegrade cable with 22-AWG or 24-AWG, with each pair twisted twice every 3.6 m (12 ft.). It's a cheaper alternative to Type 1 but is limited to 4 Mbps.
Type 5	Fiber-optic cable, 62.5- or 100-micron diameter; used for linking MSAUs over long distances.
Type 6	STP cable with two twisted pairs of 26-AWG stranded wire surrounded by braided shield and casing. Similar to Type 1, except the stranded wire allows more flexibility but less distance (two thirds that of Type 1). This cable is generally used as a patch cable or for extensions in wiring closets.
Type 8	STP cable for use under carpets. It's similar to Type 6 but is flat.
Type 9	Plenum-rated Type 6 cable.



AWG numbers are inversely related to the cable's diameter, so larger AWG numbers indicate smaller diameters. For example, standard phone wire has a thickness of 22 AWG, and thicknet cable is 12 AWG.

Table B-8 summarizes the token ring network architecture.

Table B-8 Token ring summary

Category	Specification
IEEE specification	802.5
Advantages	Fast and reliable
Disadvantages	More expensive than Ethernet; difficult to troubleshoot
Topology	Ring; cabled as star
Channel access method	Token passing
Maximum cable segment length	45 m (150 ft.) for UTP; 101 m (330 ft.) for STP
Maximum number of segments	33 hubs
Maximum devices per segment	Depends on hub
Maximum devices per network	72 with UTP, 260 with STP
Transmission speed	4 Mbps or 16 Mbps

The AppleTalk Environment

Apple Computer, Inc., designed the AppleTalk architecture for use in Macintosh networks. AppleTalk can run over a variety of other physical architectures, including Apple's own LocalTalk, Ethernet, and token ring. Because of Ethernet's superior speed and large installation base, AppleTalk is most commonly run over Ethernet, sometimes referred to as EtherTalk.

First introduced in 1983, AppleTalk is a simple, easy-to-implement network architecture. Because all Macintoshes have a built-in network interface, setting up AppleTalk is as easy as attaching all the computers with cable. Therefore, AppleTalk networks were popular in early Macintosh environments.



At its introduction, "AppleTalk" referred to the networking protocols and the hardware used to connect computers. In 1989, Apple changed AppleTalk's definition to refer to the network's overall architecture and added the term "LocalTalk" to refer to the cabling system.

Unlike Ethernet and token ring, which use the NIC's address, AppleTalk applies a dynamic scheme to determine a device's address. When a computer is powered on, it chooses a numeric address—generally, the last address it used. It then broadcasts this address to the network to determine whether the address is available. If the address isn't taken, the computer starts transmitting from this address. If, however, another device on the network is

using the address, the computer chooses another address randomly and broadcasts it to the network. This process continues until the computer finds an unused address.

The original version of AppleTalk, now called AppleTalk Phase 1, supported only 32 computers per network, and these computers could use only LocalTalk cabling. Including hubs and repeaters increased the number of computers to 254. When Apple introduced AppleTalk Phase 2 in 1989, it also introduced EtherTalk and TokenTalk, which allow AppleTalk protocols to operate over Ethernet and token ring networks, respectively. These architectures increased the number of computers an AppleTalk network can include to more than 16 million. In practice, standards governing AppleTalk networks, token ring, or Ethernet limit the number of computers to well below 16 million.

LocalTalk

Apple Computer, Inc., designed the LocalTalk network architecture, which uses STP in a bus topology, to enable users to share peripherals and data in a SOHO environment. LocalTalk uses the CSMA/CA channel access method, which prevents more collisions but is cumbersome. Using CSMA/CA is like having to mail a postcard announcing that you're sending a letter before you can send the letter.

A LocalTalk network's maximum transmission speed is only 230.4 Kbps. When compared with other network architectures' speeds (10, 100, 1000 Mbps, and even 10,000 Mbps for Ethernet and 4 or 16 Mbps for token ring), it's easy to see why this architecture was used mainly in small, Macintosh-only environments.



Starting with Mac OS X, LocalTalk hardware is no longer supported. However, if a LocalTalk network is required, purchasing a third-party adapter to convert from Ethernet to LocalTalk might be possible.

EtherTalk and TokenTalk

In an effort to overcome LocalTalk's speed limitation, Apple created EtherTalk and TokenTalk. EtherTalk is the AppleTalk protocol running over a 10 Mbps IEEE 802.3 Ethernet network. TokenTalk is much the same—AppleTalk running over a 4 or 16 Mbps IEEE 802.5 token ring network.

Both implementations require using a different NIC, one that includes all drivers and protocols needed to run EtherTalk or TokenTalk. With extra software, each protocol can be used to connect Macintosh computers to a PC Ethernet or token ring environment. Since 1996, Apple Computer has offered systems with built-in Ethernet interfaces or with options to add Ethernet or token ring to its systems at a low cost. Note that a Macintosh running Mac OS X with an Ethernet interface can participate in a Windows-based network, accessing Windows file and printer shares and giving Windows clients access to Macintosh file and printer shares.

The Fiber Distributed Data Interface Architecture

Fiber Distributed Data Interface (FDDI) uses the token-passing channel access method and dual counter-rotating rings (usually physical rings of fiber-optic cable) for redundancy. It transmits at 100 Mbps and can include up to 500 nodes over a distance of 100 km

(60 miles). FDDI full-duplex technology, an extension to standard FDDI, can support up to 200 Mbps. Like token ring, FDDI uses token passing; however, FDDI networks are often wired as a physical ring, not as a star. An FDDI network has no hubs; devices generally connect directly to each other. However, devices called concentrators can serve as a central connection point for buildings or sites in a campus setting.

An FDDI network handles token passing differently than token ring does. An FDDI token passes around the ring, but unlike token ring, when the computer possessing the token has more than one frame to send, it can send the next frame before the initial frame fully circles the ring. This transmission is possible because the sender has the token, so no other senders can become active. The computer can avoid data collisions by calculating the network latency and waiting a suitable interval before sending the next packet. This process transmits data more quickly around the network. Also, after a computer finishes sending its data, it can immediately pass the token along without waiting for confirmation of the data's receipt; the data doesn't need to make a complete circuit of the ring before the token can be passed on. Unlike token ring, FDDI supports assigning a priority level to a station or type of data. For example, a server can get higher priority than workstations, and video or time-sensitive data can get even higher priority.

As mentioned, FDDI uses two physical rings operating in different directions to avoid cable problems. In a token ring network, beaconing and network reconfiguration resolve cable breaks. In an FDDI network, all data transmission occurs along the primary ring, and the secondary ring circumvents a cable break. When a computer determines it can't communicate with its downstream neighbor, it sends data along the secondary ring. When data reaches the other end of the ring where the cable break is located, it's transferred to the primary ring, where it continues its journey.

An FDDI uses two types of NICs: dual attachment station (DAS) and single attachment station (SAS). DASs attached to both rings are intended for use in servers, concentrators, and other devices that require full reliability. SASs connected to only one ring are intended for workstations attached to concentrators. These stations still benefit from the reliability of the dual rings in FDDI because the concentrators they're attached to are usually attached to both rings, too. Table B-9 outlines the FDDI architecture.

Table B-9 FDDI summary

Category	Specification
IEEE specification	No IEEE; ANSI X3T9.1
Advantages	Very fast and reliable; can span long distances; highly secure
Disadvantages	Expensive; difficult to install
Topology	Ring
Cable type	Fiber optic
Channel access method	Token passing
Maximum total network length	100 km (60 miles)
Maximum number of devices per network	500
Transmission speed	100 Mbps

Routable Versus Nonroutable Protocols

As mentioned in Chapter 6, the Network layer of the OSI model is responsible for moving data across multiple networks. Devices called routers are responsible for this routing process. However, not all protocol suites operate at the Network layer. Protocol suites that do function at the Network layer are called routable or routed protocols, and protocol suites that don't are called nonroutable. Because routing operates at the Network layer, the routable/nonroutable attribute applies mostly to protocols operating at this layer.

A protocol suite's capability to be routed (or not) has a major impact on its effectiveness in any network requiring a router's services, such as internetworks, MANs, and WANs. TCP/IP and IPX/SPX are routable protocols well suited for these types of networks. An older and nearly obsolete protocol, NetBEUI (discussed later in "NetBIOS and NetBEUI"), is a nonroutable protocol that works well in small networks, but its performance drops considerably as a network grows. When choosing a protocol suite for your network, consider the network's current size and possibilities for future expansion.

The IPX/SPX Protocol Suite

Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) is the original protocol suite developed for use with the Novell NetWare network OS. Novell continues to support this protocol suite, even in its latest version of NetWare, but does so mainly for backward-compatibility with older NetWare networks. Currently, TCP/IP is the protocol suite of choice for networking, even in NetWare.

The Microsoft implementation of IPX/SPX was originally called NWLink, but in Windows 98, this terminology for IPX/SPX protocols was changed to Novell IPX ODI Protocol (Windows 98, Release 2) or IPX/SPX-Compatible Protocol (Windows 98, Release 1). Windows Server 2000/2003 and Windows XP include NWLink, mostly to support connections to older NetWare servers. However, NWLink can also provide transport for NetBIOS. Because it's a routable protocol suite, network expansion is easier with NWLink than with NetBEUI.

One major consideration when using IPX/SPX or NWLink is which Ethernet frame type to use. (A frame type defines the format of the Data Link-layer protocol data unit.) Remember that all computers on a network must use the same frame type to communicate successfully. If computers on a network use IPX/SPX and communication doesn't occur, verify that all computers are using the same frame type.

IPX/SPX Implementations

Open Data-link Interface (ODI) is similar to the Microsoft Network Device Interface Specification (NDIS). It allows a single network driver to support multiple protocols, thereby enabling a computer to use multiple protocols for network communication through a single NIC.

IPX is a Transport- and Network-layer protocol that handles all addressing and routing on a network. Workstations use the NIC's hardware (MAC) address for identification. IPX is a connectionless protocol that provides fast but unreliable service.

IPX Routing Information Protocol (IPX RIP) is used by servers and routers to exchange information about network addresses and topology. It's a distance-vector protocol that uses the number of hops between points to determine a packet's best path from sender to receiver. In addition to hops, IPX RIP uses ticks, a value based on the expected delay between routers. Chapter 7 discusses RIP and other routing protocols in more detail.

SPX works with IPX to provide connection-oriented services. As with all connection-oriented protocols, transmission is slower but more reliable.

NetWare Core Protocol (NCP) functions at the Transport layer and all upper layers (Session, Presentation, and Application) to provide a wide range of client/server functions. NCP handles client redirection through IPX/SPX or NWLink, including printing and file sharing.

Service Advertising Protocol (SAP) is used by file and print servers to advertise their services to computers on the network. SAP packets are broadcast periodically (usually every 60 seconds) to ensure that all computers know what services are available and the addresses of servers. Newer NetWare versions don't use SAP; instead, they use Novell eDirectory and related protocols because SAP's once-per-minute update interval can cause problems on large networks with many service advertisers.

Service Lookup Protocol (SLP) is a new IP-based NetWare protocol that applies when clients want to look up the services available on an IP-only network. SLP packets locate the nearest identifiable directory tree and ensure that all directory-enabled computers can inquire about available network services.

NetBIOS and NetBEUI

In the early 1980s, IBM hired a third-party company, Sytek, to build a simple, basic set of network programming interfaces. The result became Network Basic Input/Output System (NetBIOS). This interface persisted much longer than anyone expected and remains in wide use today. It was first deployed in a basic networking product that IBM called PC-Net and Microsoft later remarketed as MS-Net.

By the mid-1980s, Microsoft, 3Com, and IBM developed a protocol suite for use with OS/2 and LAN Manager. Using NetBIOS to provide Application-layer capabilities, this consortium developed a lower-layer protocol called NetBIOS Extended User Interface (NetBEUI) that spans Layers 2, 3, and 4 of the OSI model. NetBIOS and NetBEUI were designed to work in small to medium networks of 2 to 250 computers. NetBIOS used with TCP/IP or IPX/SPX is still around, but Microsoft has discontinued support for NetBEUI, starting with Windows XP, because TCP/IP is the protocol of choice and NetBEUI isn't routable.

Although NetBIOS and NetBEUI work closely together and are often confused with each other, they're neither inseparable nor the same. The Microsoft protocol suite defines four components above the Data Link layer. For this reason, Microsoft protocols can run on any NIC or physical medium.

The redirector interprets requests from computers and determines whether they're local or remote. It passes a local request to the local OS and a request for remote network service to the protocol below: in this case, Server Message Block (SMB). SMB is the message format that DOS and Windows use to share files, directories, and devices, and SMB file sharing is also supported by most Linux and UNIX systems. SMB passes information between networked computers, and the redirector is responsible for repackaging SMB requests for transmission to other devices for processing.

As mentioned in Chapter 6, the Session layer is responsible for managing communication between two computers. NetBIOS works at this layer to establish and maintain these connections. NetBEUI works at the Transport layer to manage communication between two computers. It's often shown operating at the Network layer, but it's actually a nonroutable protocol and skips this layer. A NetBEUI packet has no fields for source or destination network information.

NetBIOS also operates at the Session layer to provide peer-to-peer network application support. A unique 15-character name identifies each computer in a NetBIOS network, and a NetBIOS broadcast advertises this computer name. Periodically, a computer broadcasts its NetBIOS name so that other computers can communicate with it. All computers on the network keep a cache of names and hardware addresses of computers they receive broadcasts from. If a computer wants to communicate with a computer whose name isn't in its cache, it sends a broadcast requesting that computer's hardware address.

NetBIOS is a connection-oriented protocol responsible for establishing, maintaining, and terminating network connections but can use connectionless communication, if necessary. Although closely related to NetBEUI, NetBIOS can use other lower-layer protocols, including TCP/IP and IPX/SPX, for transport and lower-layer services. It's a nonroutable protocol, but it can be routed when using a routable protocol for transport.

NetBEUI is a small, fast, nonroutable Transport- and Data Link-layer protocol designed for use with NetBIOS on small networks. NetBEUI 3.0 is the Microsoft improvement on IBM's version of NetBEUI (and, therefore, works only on Microsoft networks). Its low overhead makes NetBEUI ideal for DOS-based computers that require network connectivity, and its speed and size also make it a good choice for slow serial links. Because NetBEUI isn't routable, its use is typically limited to small networks. However, as mentioned, Microsoft no longer supports NetBEUI.

AppleTalk

Although the AppleTalk standard defines physical transport in Apple Macintosh networks, it also establishes a suite of protocols these computers use to communicate. Apple created AppleTalk Phase II to allow connectivity outside the Macintosh world. Instead of defining networks, AppleTalk divides computers into zones that network administrators can use to logically group computers and other resources that communicate frequently, in a manner similar to subnetting.



Network Troubleshooting Guide

This appendix lists basic questions you can ask when you approach network problems. Guidelines for troubleshooting specific areas of networking technology are also included.

General Questions for Troubleshooting

When troubleshooting, the first question you should ask is “Has this piece of equipment or procedure ever worked correctly?” If it did in fact work once, your next question should be “Since then, what has changed?” The following is a list of other useful questions:

- Was only one user affected, or were many users affected?
- Were users affected randomly or all at once?
- Is only one computer down, or is the whole network down?
- Does this problem happen all the time or only during specific times?
- Does this problem affect only one application, more than one, or all applications?
- Does this problem resemble any past problems?
- Have you added any users to the network?
- Have you added any new equipment to the network?
- Did you install a new application just before the problem occurred?
- Have you moved any equipment recently?
- Are any vendor products involved in this problem? If so, who are the vendors?
- Does this problem occur in components, such as disk drives, hubs, application software, cards, or network software?
- Has anyone else attempted to remedy this problem?
- Can the computer having the problem function as a stand-alone computer if it’s not functioning on the network?
- If the computer can’t function on the network, have you checked whether its NIC is working?
- Is the amount of traffic on the network normal?

The following sections list questions that are helpful when troubleshooting specific components of your network.

Cabling Problems

If you suspect a problem with cabling, check the following items:

- Missing or loose connections
- Frayed or broken sections
- Correct length
- Cable and connectors match (Cat 5E cable and Cat 5E connectors, for example)
- NIC specifications
- Crimped or bent cables
- Location of the cable routing near a transformer, large electric motor, or air conditioner
- Correct termination at jacks and patch panels

Problems with NICs

Here are some things to check with NICs:

- If you have more than one NIC in a computer, do any settings conflict?
- Are the type (half- or full-duplex) and signaling speed (for example, 10 Mbps, 100 Mbps, 1000 Mbps) set correctly?
- In a wireless environment, are the SSID and encryption key set correctly?

Driver Problems

Check the following to isolate driver problems:

- How old is the equipment?
- Have any changes been made to the equipment since it was working correctly?
- Has anyone moved any hardware?
- Has software been installed recently?
- Are old drivers being used with new equipment?
- Have you checked the manufacturer's Web site for the newest drivers?

Problems with Network Operations

You can follow this checklist for network operation problems:

1. Inspect the hardware in servers and verify the following:
 - It's on the OS vendor's compatibility list.
 - It has the correct, most current drivers installed.
 - It contains enough memory for the network operations you perform.
 - It has adequate hard drive space for the amount of information stored on it.
 - It has plenty of processing power to support your network.
2. Check all network bindings to make sure they're correct, and verify that the most used bindings are listed first.

3. Double-check client computers to verify that they have the correct client software (redirectors) installed.
4. Check that the installed protocol matches the protocol already in use on the network.
5. In a wireless network, make sure matching standards are in use; for example, ensure that all devices use 802.11b, 802.11g, 802.11n, or 802.11a as appropriate.
6. In a wireless network, check that all clients have a strong enough signal from the access point and are using the correct SSID and encryption type and key.

Problems with Network Printing and Fax Services

Check the following if there's a problem with network printing or faxing:

- Is the shared fax's or printer's power on?
- Is the selected shared printer or fax machine the correct one for the client computer's driver?
- Are permissions correct for the shared printer or fax that users and printer/fax managers are using?
- Are the cables used by the shared printer or fax in good condition and connected correctly?

Problems with Client/Server Computing

Check the following for problems with clients or servers in a client/server environment:

- Is the client front end configured correctly and working?
- Is the server software configured correctly and working?
- Is the network application doing what it's supposed to?
- Does the server running the network application have enough RAM, space on its shared disk, and processing power?
- Have users received training in using the network application, and are they using the correct methods to get the most out of it?

Problems with Network Accounts

Check the following if a user can't log on with a certain account:

- Is the user entering the correct username?
- Is the name of the domain (if logging on to a Windows domain) correct?
- Is the user typing the correct password? Remember that passwords are case sensitive.
- Has the user account been disabled or locked out?

Problems with Data Security

Use the following checklist if you suspect problems with data security:

1. If a user can access a resource that should be unavailable or can't access a resource that should be available, check the following:
 - Does the user have the correct permissions for the resource?
 - Does the user belong to a group that has the correct access to the resource?

- Do any trustee assignments to the resource conflict? (Check share-level permissions versus user-level permissions).
2. Check whether the user belongs to any group assigned a “no access” permission.
 3. If the user can access previously secured data or there’s a problem with data theft, alteration, or contamination, check the following:
 - If the server is in a locked room, who has access to it?
 - Are any computers being left as logged on and unattended?
 - Are any passwords written on paper and left in obvious places, such as on the monitor, in a desk drawer, or under the keyboard?
 - Are any users using obvious passwords, such as names of children, pets, or spouses?
 - Are any users using the same password with a revision number (that is, Dawn1, Dawn2, Dawn3, and so on)?
 - Do any users have a regular logon name equivalent to an administrator or superuser?
 - Are any users storing confidential data on their local hard drives?
 - Do any users have their OSs configured to log them on automatically, bypassing the username and password process?

Problems with WAN Communication

To start, you troubleshoot a WAN in the same way you do a LAN. However, some considerations are specific to WANs. These types of problems usually require the assistance of vendors or service providers. Here are some questions related to WAN troubleshooting:

1. Did any vendor replace, add, or remove anything from the WAN?
2. Is the power to the following components turned on, and are the components turned on?
 - Switch
 - Router
 - Modem
 - CSU/DSU
3. For the same components, check the following:
 - Are all cables connected correctly and in good condition?
 - Is the component compatible with the communication medium and the communication device at the other end of the link?
 - Is the software configured correctly, and does it match the configuration of the connected communication equipment?

Networking Resources, Online and Offline

Numerous resources are available to help you find information you need to implement a network successfully. This appendix identifies many valuable networking resources. In addition to the resources listed here, you can find helpful networking information with an Internet search or by visiting the Web sites listed in the “Online/Electronic Materials” section.

Printed Materials

Ciampa, Mark. *Security+ Guide to Networking Security Fundamentals, Third Edition*. Course Technology, Cengage Learning, Boston: 2004 (ISBN 0619215666).

Eckert, Jason W. and M. John Schitka. *Linux+ Guide to Linux Certification, Second Edition*. Course Technology, Cengage Learning, Boston: 2006 (ISBN 0619216212).

Tomsho, Greg. *Guide to Network Support and Troubleshooting*. Course Technology, Cengage Learning, Boston: 2002 (ISBN 061903551X).

Tomsho, Greg. *MCTS Guide to Microsoft Windows Server 2008 Active Directory Configuration*. Course Technology, Cengage Learning, Boston: 2010 (ISBN 1423902351).

Tulloch, Ingrid and Mitch Tulloch. *Microsoft Encyclopedia of Networking*. Microsoft Press, Redmond, WA: 2002 (ISBN 0735613788).

Online/Electronic Materials

Acronym Finder (www.acronymfinder.com)

Gigabit Ethernet Resources (www.ethermanage.com/ethernet/gigabit.html)

How Stuff Works (www.howstuffworks.com): A great Web site to learn how just about anything works, from networks to yo-yos

IEEE Local and Metropolitan Area Network Standards (<http://ieeexplore.ieee.org/Xplore/guesthome.jsp>): The place to find information on all IEEE networking standards, purchase the standards, or download the 802 standard in PDF format

O'Reilly Safari Bookshelf (<http://safari.oreilly.com>): A subscription service to hundreds of current information technology books you can read online

Overview of Cable Modem Technology and Services (www.lightreading.com/lr-cable/)

Protocols.com (www.protocols.com): A helpful reference for reviewing packet structure and header fields of all networking protocols

TechFest Networking tutorials (www.techfest.com)

TechNet online version: (<http://technet.microsoft.com/en-us/default.aspx>)

TechWeb (www.techweb.com)

Wi-Fi News (www.wi-fiplanet.com): A Web site dedicated to information on the 802.11b wireless networking standard

A Step-by-Step Guide to Using Server Virtualization Software

Virtualization enables a school or student to get the most out of computer resources. Schools can use virtualization to turn a single server-grade computer into a virtual server that can host two, three, or more operating systems. For example, one computer can house three virtual servers running Windows Server 2008. This capability saves the school money on servers and enables more students to be able to work on their own OSs.

Another capability of virtualization is making it possible for a school or student to turn a single PC into a virtual system on which to run another OS—without having to alter the current OS running on the PC. A single computer lab PC or a student’s home PC can be turned into a host for Windows Server 2008. This is ideal, for example, when your textbook comes with an evaluation copy of Windows Server 2008. You can install virtualization software and then install Windows Server 2008 for doing hands-on projects and activities. You can use your original OS, such as Windows XP or Windows 7, and then use Windows Server 2008 in a virtual window or session, for example. When you’re finished learning Windows Server 2008, you simply remove the virtualization software and you’re back where you started with your original OS.

This appendix is a step-by-step guide for turning a computer into a virtual system hosting one or more virtual machines. The main focus is on three popular virtualization systems that are available free:

- *Microsoft Virtual PC*—Intended for a workstation-grade PC to host another OS, such as a Windows Server 2008 virtual machine
- *Microsoft Virtual Server*—Intended for a server-grade computer to host multiple virtual machines, including Windows Server 2008 and other OSs
- *VMware Server*—Intended for server-grade computers to host multiple virtual machines

For each of these virtualization systems, you learn how to do the following:

- Download a free version and install it.
- Create a virtual machine.
- Install a guest OS, such as Windows Server 2008, in the virtual machine, and then access the virtual machine’s OS.
- Install ISO images.
- Configure virtual networking.
- Configure hardware components.

At the end of this appendix is a brief look at VMware Workstation 6 and Microsoft Hyper-V.

Microsoft Virtual PC

Microsoft Virtual PC can be installed in Windows XP, Vista, 7, and Server 2003 OSs. At this writing, it's not adapted to be installed in Windows Server 2008. Although Microsoft Virtual PC is intended to host workstation OSs as virtual machines, you can also use it to create a Windows Server 2008 Standard Edition virtual machine.

Microsoft Virtual PC is available as a free download. From a student's perspective, it's ideal for running the Windows Server 2008 Standard Edition evaluation DVD (available at www.microsoft.com) on a Windows XP or Vista computer. It works equally well on Windows XP or Vista computers in a student lab.

Requirements for Microsoft Virtual PC

At this writing, Microsoft Virtual PC 2007 with Service Pack 1 (SP1) is the most recent version. It can be loaded on the following host OSs:

- Windows XP Professional with SP2 or SP3
- Windows Server 2003 Edition SP2 (x86 or x64)
- Windows Vista Business Edition (x86 or x64 versions with or without SP1)
- Windows Vista Enterprise Edition (x86 or x64 versions with or without SP1)
- Windows Vista Ultimate Edition (x86 or x64 versions with or without SP1)

The hardware requirements for Microsoft Virtual PC 2007 SP1 are as follows:

- *CPU*—Intel Celeron, Pentium II, Pentium III, Pentium 4, Core Duo, or Core 2 Duo CPU or an AMD Athlon or Duron CPU (400 MHz or faster; x86 or x64).
- *RAM*—Enough RAM for at least the minimum requirements of the total number of OSs you're running. For example, if you're running Windows XP Professional (128 MB minimum) and want to load Windows Server 2008 (512 MB minimum) as a virtual machine, you need a minimum of 640 MB to 1 GB RAM. If Windows Vista is the host and you want to run a Windows Server 2008 Standard Edition virtual machine, you need a minimum of 1 GB RAM.
- *Disk space*—Enough disk storage for the OSs you plan to run. For example, Windows XP requires at least 1.5 GB, Windows Vista requires at least 15 GB, and Windows Server 2008 requires at least 10 GB (but 15 GB to 20 GB is better for using different roles and services).

Guest OSs Supported

After Virtual PC 2007 SP1 is loaded, you can run any of the following OSs as virtual machines (guests) in Virtual PC 2007 SP1:

- Windows 98 and 98 SE
- Windows Me
- Windows 2000 Professional
- Windows XP Home or Professional with SP1, SP2, SP3 (or no service pack)
- Windows Vista Business Edition (x86 or x64 versions with or without SP1)
- Windows Vista Enterprise Edition (x86 or x64 versions with or without SP1)

- Windows Vista Ultimate Edition (x86 or x64 versions with or without SP1)
- Windows Server 2008 Standard Edition
- OS/2 Warp

Downloading Microsoft Virtual PC

Microsoft Virtual PC can be downloaded from Microsoft's Web site for no cost. The steps to download Microsoft Virtual PC 2007 SP1 are as follows:

1. Log on to your computer, and create a folder for storing the setup.exe file for Microsoft Virtual PC.
2. Start a Web browser, and go to www.microsoft.com/downloads or www.microsoft.com/downloads/Search.aspx?displaylang=en (for English).



Web links and specific instructions change periodically. You might need to search www.microsoft.com for the most current link if these links don't work.

3. Look for Microsoft Virtual PC in the Popular Downloads or Recommended Downloads sections. (Also check the New Downloads section in case a new version is available.) If you find it in one of these sections, click the **Microsoft Virtual PC** link. If you don't see a link, click **Windows** under the Product Families heading. Click the **Show downloads** for list arrow, click **Microsoft Virtual PC**, and then click **Go**.
4. Click the **VPC 2007 SP1** link.



To use Microsoft Virtual PC 2007 with Windows Server 2008 or Windows Vista as the guest OS, you must use the download containing SP1.

5. Click the **Download** button for the setup.exe file that matches your computer, which is 32 BIT\ setup.exe for an x86 computer or 64 BIT\ setup.exe for an x64 computer.
6. Save the file in the folder you created in Step 1, close the Download complete dialog box, and exit your Web browser.

Installing Microsoft Virtual PC

To install Microsoft Virtual PC 2007 SP1, follow these steps:

1. Browse to the folder where you saved the setup.exe file for Microsoft Virtual PC, and double-click **setup.exe**.
2. In the welcome window for the Microsoft Virtual PC 2007 SP1 Wizard (see Figure E-1), click **Next**.
3. Click the **I accept the terms in the license agreement** option button, and then click **Next**.
4. Enter your username and name of your organization or school (if needed). The product key should already be entered. Make sure **Anyone who uses this computer (All Users)** is selected (if this option is available), and then click **Next**.
5. Click **Install**. The installation process takes a few minutes. When it's completed, click **Finish**.



Figure E-1 The Microsoft Virtual PC 2007 SP1 Wizard

Courtesy of Course Technology/Cengage Learning

Creating a Virtual Machine and Installing a Guest OS

After Microsoft Virtual PC 2007 SP1 is installed, the next step is to create a virtual machine in which to install a guest OS.



Microsoft Virtual PC 2007 SP1 might not be compatible with hardware virtualization on some CPUs. If you notice a crash dump when configuring the virtual machine or loading the guest OS, first make sure you have enabled hardware virtualization as described in the following steps. If this setting doesn't work, try disabling hardware virtualization in the BIOS and restarting these steps from the beginning.

The following steps describe the general procedure for setting up a virtual machine with Windows Server 2008 Standard Edition as the guest OS:

1. In the host OS, click **Start**, point to **All Programs**, and click **Microsoft Virtual PC**.
2. In the welcome window of the New Virtual Machine Wizard (see Figure E-2), click **Next**.



Figure E-2 The New Virtual Machine Wizard

Courtesy of Course Technology/Cengage Learning

3. Make sure **Create a virtual machine** is selected, and then click **Next**.
4. Enter a name for the virtual machine, such as **Windows Server 2008**, and then click **Next**.
5. Click **Windows Server 2008** as the OS to install, and then click **Next**.
6. Make sure at least 512 MB to 1 GB RAM is allocated for the virtual machine. If necessary, click **Adjusting the RAM** and use the slider bar to change the amount of memory. Click **Next**.
7. Verify that **A new virtual hard disk** is selected, and then click **Next**.
8. Make sure the virtual hard disk is sized to meet your needs, or leave the default size. (You need 15 GB for Windows Server 2008 and might use at least 20 to 40 GB, for example.) Click **Next**.
9. Click **Finish**. You should see the Virtual PC Console open. If it doesn't, click **Start**, point to **All Programs**, and click **Microsoft Virtual PC**.
10. You can configure some options at this point by clicking **File, Options** from the menu. Click each option to see what it does and configure any options as necessary. When you're finished, click **OK**. The options are as follows:
 - *Restore at Start*—Pauses a running virtual machine when you exit the console and restores the virtual machine when you open the console again.
 - *Performance*—Specifies how CPU time is allocated to virtual machines and what happens when Virtual PC is a process running in the background.
 - *Hardware virtualization*—Enable hardware virtualization, if your CPU has this capability.
 - *Full-Screen Mode*—Allows adjusting the screen resolution so that it's the same for the host and guest OSs. (Note the previous caution if this setting is enabled.)
 - *Sound*—Configures virtual machine sound, which is muted by default. If you enable it, sounds from the host and guest OS can be difficult to differentiate.
 - *Messages*—Turn off error and informational messages from Virtual PC.
 - *Keyboard*—Specifies the host key for the guest OS. The default is the right Alt key. When you press this key, you can switch the mouse between the guest and host windows and use guest key combinations, such as pressing Alt+Delete to send the Ctrl+Alt+Delete key combination to the guest OS for logging on.
 - *Mouse*—Specifies how the pointer is captured in the virtual machine window.
 - *Security*—Determines how to control access to Virtual PC functions.
 - *Language*—Specifies the language to use for Virtual PC.
11. Insert the Windows Server 2008 Standard Edition installation DVD.



At this point, you could install any supported guest OS. If you were installing a different OS, you would insert the CD/DVD now, and the remaining steps would be unique to the OS you're installing.

- Click the **Start** button in the Virtual PC Console. This opens a second larger window, which is the Microsoft Virtual PC 2007 console. Wait a few minutes for the DVD to start loading. Click in the console to enable the mouse to operate in this window. (If necessary, you can switch the mouse movement back so that it can go all over the screen by pressing the right Alt key, which is the “host” key.)



Occasionally, the mouse might seem stuck, move slowly, or stop functioning in the active portion of the console. If this happens, close all windows and open the Virtual PC Console from the Start menu again. Also, some installation processes take longer in a virtual machine. Don't close the window or stop the installation prematurely, even if you seem to be stuck on a black screen for several minutes.

- Click the language you want in the Language to install drop-down list. In the Time and currency format list box, make your selection, such as English (United States). In the Keyboard or input method list box, make your selection, such as US. Click **Next**, and then click **Install now**.
- Click **Windows Server 2008 Standard (Full Installation)**, and then click **Next**.
- Read the license terms, click the **I accept the license terms** check box, and click **Next**.
- Click **Custom (advanced)**. You'll see the amount of unallocated disk space selected, which is the disk space you specified when you configured the virtual machine. Make sure it's selected, and then click **Next**.
- The Windows Server 2008 installation starts. You'll see progress information about copying files, expanding files, installing features, installing updates, and completing the installation. This part of the installation can take 30 minutes or longer. When it's finished, your computer restarts automatically.
- The message “Please wait while Windows sets up your computer” is displayed, and then you see a window indicating that the installation is completing.
- Your system restarts again. When you see a message that the user's password must be changed before logging on the first time, click **OK**. (You might have to click inside the console first to have the mouse function in it.)
- Enter a new password for the Administrator account, and then enter the same password again to confirm it. Click the **blue circle** with the white arrow inside.



If you enter a password that isn't a strong password, you'll see the message “Unable to update the password.” This means the password doesn't meet the length, complexity, or history requirements of the domain. Click OK and enter a different password that's longer than seven characters and uses letters, numbers, and special symbols, such as &.

- When you see the message “Your password has been changed,” click **OK**. The Windows desktop is displayed and the Initial Configuration Tasks applet opens. You can configure Windows Server 2008 as you would in a nonvirtual environment.

22. When you close the Microsoft Virtual PC 2007 console, you can turn off the virtual machine or save its current state. Unless you want to save its state, a good practice is to shut down the server before closing the window. (Saving the state means to keep the server in its current state without shutting it down.) When you shut down the server in this way, the Microsoft Virtual PC 2007 console closes but leaves the Virtual PC Console open. Also, to restart the virtual machine, open the Virtual PC Console, click **Start**, and wait for the system to boot in the Microsoft Virtual PC 2007 console.



When you log on to Windows Server 2008 from the console, the normal Ctrl+Alt+Delete key sequence doesn't work. Instead, click the Action menu and press Ctrl+Alt+Delete. Another alternative is to press right Alt+Delete.

Installing an OS from an ISO Image

An ISO file is an optical disc (CD/DVD) image file with the .iso file extension. It can be accessed in several ways, such as from a CD/DVD, from a hard drive, or as a shared network file. Typically, when you download an OS, such as an evaluation copy of a Windows OS, you download an ISO file. One advantage of using an ISO file for installing a guest OS on a virtual machine is that the installation process can go faster. Virtual PC enables you to install from an ISO file by using the following general steps:

1. Follow Steps 1 to 10 in the previous section, “Creating a Virtual Machine and Installing a Guest OS.”
2. Click the **Start** button in the Virtual PC Console.
3. After the Microsoft Virtual PC 2007 console opens, press the **right Alt** key if necessary to access the menu at the top of the window.
4. Click **CD, Capture ISO Image** from the menu.
5. Navigate to and click the ISO file, and click the **Open** button.
6. You return to the Microsoft Virtual PC 2007 console, and then you should restart the virtual machine.

Configuring Networking and Hardware Options

You can configure a range of networking and hardware options in Microsoft Virtual PC. For example, if the host computer has two or more NICs, you can specify which NIC to use for a virtual machine. You might also need to create additional virtual hard disks for a virtual machine. Use these steps to configure networking and hardware options:

1. Open the Virtual PC Console, if necessary. Make sure the virtual machine is turned off before you start.
2. Click the **Settings** button, or click **Action, Settings** from the menu to open the dialog box shown in Figure E-3.
3. Click **Networking** in the left pane. If your computer has multiple adapters, you can select the adapter (or multiple adapters) to associate with a virtual machine.

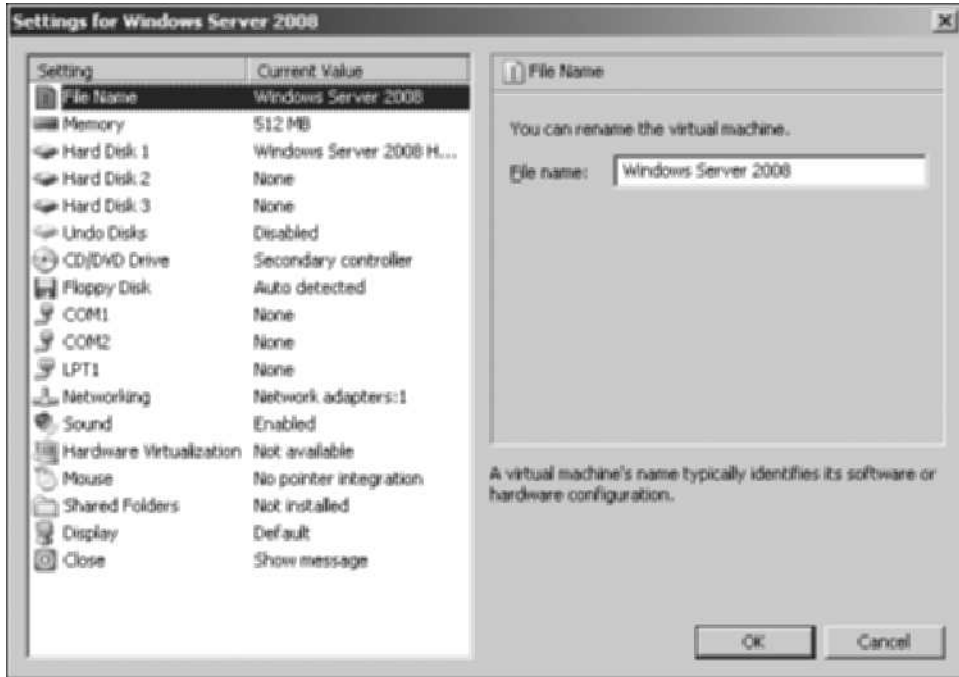


Figure E-3 Settings for a virtual machine

Courtesy of Course Technology/Cengage Learning

4. In the right pane, click the list arrow for the adapter that's selected by default. The following options are available; make the selections that are suitable for your network:
 - *Not connected*—Used if you don't want the virtual machine to access a network (including the Internet) and be accessed from a network.
 - *Local*—If two or more virtual machines are set up, they can access each other; however, virtual machines can't access the network.
 - *NetworkInterfaceName*—The actual name of a NIC model, such as an Intel or a Broadcom NIC, that the virtual machine is connected to for regular network and Internet access. With this selection, network configuration tasks that apply to other network computers also apply to the virtual machine. If a DHCP server is on the network or if the network uses a router with NAT, the virtual machine's network connection can be configured to use these services. The same applies if a DNS server is set up.
 - *Shared Networking (NAT)*—Used to create a private Virtual PC network that has a virtual DHCP server and a virtual NAT-enabled router or firewall. Typically, the first virtual computer created acts as the DHCP server and provides NAT services. In this arrangement, Microsoft Virtual PC performs as a virtual DHCP server, leasing IP addresses for virtual machines in the range of 192.168.131.1 to 192.168.131.253. Further, the virtual machines appear as computers in a private NAT-protected network. A connection to the Internet is shared among virtual machines and is protected in a way similar to a NAT-enabled router or firewall.
 - *Loopback Adapter*—You see this option if the OS is configured to have a Microsoft loopback adapter (configured as a network adapter, such as through the Add

Hardware option in Control Panel). This option is used in two contexts. One is when no physical network connection is present, but you want to simulate network connectivity between the host and all virtual machines. A second is when you're creating a network with many routers and firewalls as well as many virtual machines.

5. Click **Memory** in the left pane. Notice that you can increase the memory allocation for the virtual machine by using the slider bar in the right pane.
6. In the left pane, click **Hard Disk 1**. The right pane shows the path to the virtual hard disk file, and you can configure the Hard Disk 2 and Hard Disk 3 options for additional virtual hard disks. To do this, click **Hard Disk 2** in the left pane, for example, and click the **Virtual Disk Wizard** button in the right pane. (A virtual machine can have up to three hard disks.)
7. Click **CD/DVD Drive** in the left pane, and notice that in the right pane, you can attach a CD or DVD drive.
8. Click **Hardware Virtualization** in the left pane. In the right pane, you can enable hardware virtualization, if your computer supports it.
9. You can configure additional hardware, such as communication (COM) ports, a floppy disk, printer (LPT) ports, sound, the mouse, the display, and other devices. When you're finished with the configurations, click **OK**.

Host Key Options

Because a virtual machine represents an OS running inside an OS, you need a way to use the keyboard so that the keys you press communicate with the guest OS. For example, pressing Ctrl+Alt+Delete opens the Windows Security dialog box or a menu of options, depending on which Windows version is the host OS. It doesn't take you to a logon screen in the guest OS.

Microsoft Virtual PC enables you to communicate with the guest OS by using the host key, which is the right Alt key by default. Table E-1 lists important host key combinations you can use while you're accessing a virtual machine.

Table E-1 Host key options for Microsoft Virtual PC

Key combination	Result
HostKey	Enables you to move the mouse outside the window area used by the guest OS. (Move the mouse back into the guest OS display and click when you want to work in the guest OS.)
HostKey+Delete	The virtual machine OS responds to this key combination as Ctrl+Alt+Delete.
HostKey+P	Toggles the virtual machine between pause and resume.
HostKey+R	Causes the virtual machine to reset.
HostKey+A	Selects all items in the active window in the guest OS.
HostKey+C	Copies selected text and items in the guest OS's active window.
HostKey+V	Pastes text and items in the guest OS's active window.
HostKey+Enter	Switches between full screen and window modes.
HostKey+down arrow	Minimizes the virtual machine window.
HostKey+I	Enables you to install virtual machine additions.

Microsoft Virtual Server

Microsoft Virtual Server 2005 is intended to host server OSs as virtual machines. At this writing, Microsoft Virtual Server 2005 R2 SP1 is the most recent version. It supports hardware (integrated in the CPU) virtualization, such as AMD CPUs equipped with AMD-V and Intel CPUs with Intel VT. Other new features include the following:

- Can be installed in x64 OSs
- Supports Internet small computer system interface (iSCSI), a technology used in storage area networks (SANs)
- Can cluster virtual servers on a single computer
- Enhances Active Directory support by publishing Virtual Server binding data through service connection points

Other features of Microsoft Virtual Server include the following:

- Allows dynamic expansion of virtual disks
- Supports most popular x86 OSs
- Can mount a virtual disk on a different OS
- Enables use of Volume Shadow Copy Service (VSS) for backups (used in newer versions of Windows OSs, such as Windows Server 2008 and Vista)
- Offers virtual server management through the Virtual Server Web console
- Can use scripting to control virtual machine setups
- Memory access can be resized

Guest OSs Supported

Microsoft Virtual Server can house virtual machines for Windows and Linux server and workstation OSs. The following can be guest OSs:

- Windows Server 2008 Standard, Enterprise, Datacenter, and Web Server (x86 and x64)
- Windows Server 2003 Standard, Enterprise, Datacenter, and Web Server SP1 or SP2 (x86 or x64)
- Windows Server 2003 Standard, Enterprise, Datacenter, and Web Server R2 (x86 or x64)
- Windows Small Business Server 2003 (Standard and Premium Editions)
- Windows 2000 Server
- Windows XP Professional SP2
- Windows Vista Business, Ultimate, and Enterprise
- Red Hat Enterprise Linux versions 2.1 to 4.0
- SUSE Linux Enterprise Server 9.0
- SUSE Linux versions 9.2 to 10.0



Other OSs can run experimentally in Microsoft Virtual Server.

Host OSs Supported

Microsoft Virtual Server can be installed on the following Windows host OSs:

- Windows Server 2008 Standard and Enterprise (x86 or x64)
- Windows Server 2003 Standard, Enterprise, and Web Server with SP1 or SP2 (x86 or x64)
- Windows Server 2003 Standard, Enterprise, and Web Server R2 (x86 or x64)
- Windows Small Business Server 2003 (Standard and Premium Editions, also R2 versions)
- Windows 2000 Server with SP3 or SP4
- Windows XP Professional (x86 and x64)
- Windows Vista Business, Ultimate, and Enterprise Editions

Requirements for Microsoft Virtual Server

The hardware requirements for Microsoft Virtual Server 2005 R2 with SP1 are as follows:

- *CPU*—Intel Celeron, Pentium III, Pentium 4, Xeon, or AMD Opteron, Athlon, Athlon 64, Althon X2, Duron, or Sempron (550 MHz or faster; x86 or x64).
- *RAM*—Enough RAM to match at least the minimum requirements of the total number of OSs you'll be running. For example, if you're running Windows XP Professional (256 MB minimum required for Virtual Server) and want to load Windows Server 2008 (512 MB minimum) as a virtual machine, you need a minimum of 768 MB to 1 GB RAM. If Windows Server 2003 R2 Standard Edition is the host and you want to run a Windows Server 2008 Enterprise Edition virtual machine, you need a minimum of 768 MB to 1 GB RAM.
- *Disk space*—Enough disk storage for the OSs you plan to run. For example, Windows Server 2003 R2 Standard Edition requires at least 3 GB, and Windows Server 2008 requires at least 10 GB (but with 15 to 20 GB, you can load more roles and services).

Downloading Microsoft Virtual Server

To download Microsoft Virtual Server from Microsoft's Web site free, follow these steps:

1. Log on to your computer, and create a folder for storing the setup.exe file for Microsoft Virtual Server.
2. Start a Web browser, and go to www.microsoft.com/downloads or www.microsoft.com/downloads/Search.aspx?displaylang=en (for English).



Web links and specific instructions change periodically. You might need to search www.microsoft.com for the most current link if these links don't work.

3. Look for Microsoft Virtual Server in the Popular Downloads or Recommended Downloads sections. (Also check the New Downloads section in case there's a new version.) If you find it in one of these sections, click the **Microsoft Virtual Server** link. If you don't see a link, set the search box near the top to Windows, if necessary. Enter **Virtual Server** in the text box next to the Go button, and click **Go**.
4. Click the **Virtual Server 2005 R2 SP1** link, and then click the **Continue** button to register for the free download.
5. The information you enter next depends on whether you have already signed up for Windows Live ID or have an MSN Hotmail, MSN Messenger, or Passport account. If you already have an account, enter your e-mail address and password for Windows Live ID, click **Sign in** to verify your information (and answer any required questions), and click **Continue**. If you don't have an account or a Windows Live ID, follow the steps to sign up for a Windows Live ID.
6. Click the **Download** button for the setup.exe file that matches your computer, which is 32 BIT\ setup.exe for an x86 computer or 64 BIT\ setup.exe for an x64 computer.
7. Save the file in the folder you created in Step 1. When the download is completed, click **Close** in the Download complete dialog box, and exit your Web browser.

Installing Microsoft Virtual Server

The general steps for installing Microsoft Virtual Server on the host OS are as follows:

1. Browse to the folder where you saved the setup.exe file for Microsoft Virtual Server, and double-click **setup.exe**.
2. Click **Install Microsoft Virtual Server 2005 R2 SP1** (see Figure E-4).
3. Click **I accept the terms in the license agreement**, and then click **Next**.
4. Enter your username and the name of your organization or school (if necessary). The product key information is entered by default. Click **Next**.
5. In the Setup Type window, make sure **Complete** is selected, as shown in Figure E-5, and then click **Next**.
6. The Virtual Server Administration Website is added to Internet Information Services (IIS), and the default port is 1024. If it's available, make sure **Configure the Administration Website to always run as the authenticated user (Recommended for most users)** is selected, and then click **Next**.



After you click Next, you might see a message stating that the installed version of IIS doesn't allow multiple Web sites. The Virtual Server Administration Website is added as a virtual directory under the default site.

7. If Windows Firewall is enabled on your computer, you can have the setup process create firewall exceptions for Virtual Server. Make sure **Enable Virtual Server exceptions in Windows Firewall** is selected, and then click **Next**.
8. Click **Install**. If the IIS components needed for the Virtual Server Administration Website aren't already installed, click Yes to install them. Click **Install** again, if necessary. You'll see a dialog box showing that the components are being installed.



Figure E-4 Installing Microsoft Virtual Server 2005 R2 SP1

Courtesy of Course Technology/Cengage Learning



Figure E-5 Selecting the setup type

Courtesy of Course Technology/Cengage Learning



If you see a message that the installation program needs the IIS World Wide Web service installed and there's no option to install it, typically it means the Virtual Server installation program can't install IIS. Click OK in the message box, click Cancel to stop the installation, and follow the steps for your host OS to install IIS. (You might need the host OS installation CD/DVD.) Start the Virtual Server installation again from Step 1.

9. You'll see a window showing that Microsoft Virtual Server 2005 R2 SP1 is being installed. Click **Finish**, and then close any open windows.

Creating a Virtual Machine and Installing a Guest OS

After Microsoft Virtual Server is installed, you can use the Virtual Server Administration Website to configure Microsoft Virtual Server, configure a virtual machine, and install a guest OS. Here are the steps for creating a virtual machine and installing a guest OS (using Windows Server 2008 as the guest OS):

1. Click **Start**, point to **All Programs**, and click **Microsoft Virtual Server**.
2. Click **Virtual Server Administration Website**. In the Connect to dialog box, enter a username and password (for an account with administrator privileges), and click **OK**.
3. If you're using a recent version of Windows Firewall, you might see the Internet Explorer dialog box open so that you can add this Web site to the list of trusted sites. (You're likely to see this dialog box the first time you access the Virtual Server Administration Website.) Click the **Add** button. In the Trusted sites dialog box, click the **Add** button for the site you're adding, and then click **Close**. Also, if you see the Microsoft Phishing Filter dialog box, select whether you want to enable the Phishing Filter (which is recommended), and then click **OK**.
4. The Virtual Server Administration Website is displayed in Internet Explorer, as shown in Figure E-6. Notice that the left pane contains options to navigate, create and add virtual machines, manage virtual disks, manage virtual networks, and manage the virtual server. In the left pane under Virtual Machines, click **Create**.
5. Enter the name for the virtual machine and set the virtual machine memory. For Windows Server 2008, you should set it for at least 512 MB to 1024 MB. Also, click **Create a new virtual hard disk** and set at least 15 GB (more is better) for Windows Server 2008. Finally, specify the virtual network adapter, such as an external network interface, and then click **Create**.



The virtual network adapter options are Not connected, External Network, and Internal Network. Not connected (the default) doesn't provide any type of connection, so you can access the virtual machine only from the server. External Network means users can connect to the virtual machine through the computer's NIC. Internal Network means there can be a connection between virtual machines on the same computer.

6. If you see the option to enable AutoComplete (to remember your entries in Web forms), click **Yes** or **No**.
7. In the right pane, review the configuration information for your test server. You can use this pane to make changes to the configuration. (See "Configuring Networking and Hardware Options" later in this appendix for more information about configuring these options.)

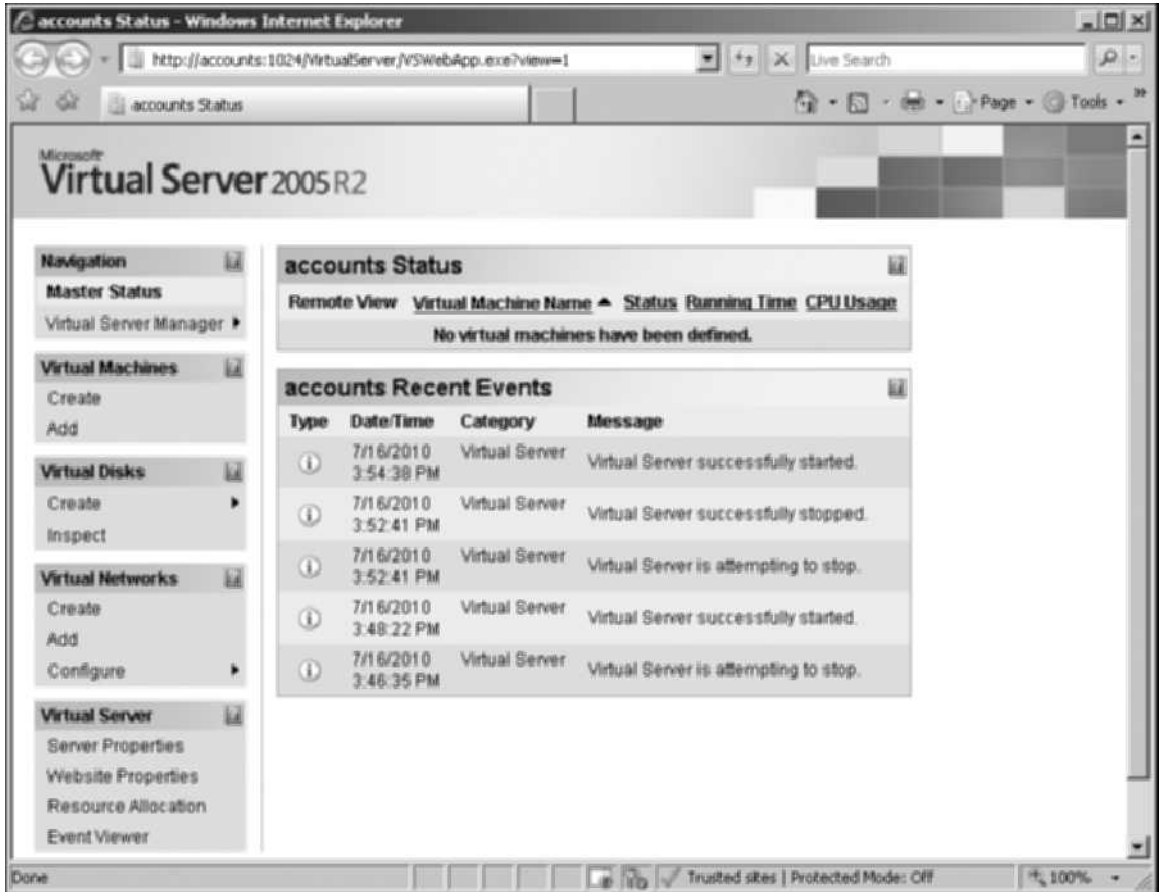


Figure E-6 The Virtual Server Administration Website

Courtesy of Course Technology/Cengage Learning

8. So that you can access a window in which to use the virtual server, click **Server Properties** under **Virtual Server** in the left pane.
9. In the right pane, click **Virtual Machine Remote Control (VMRC) Server**. Make sure the **VMRC server** check box is selected for the **Enable** option and the TCP/IP address of the host server is entered. (If you have trouble connecting after entering the host server's TCP/IP address, try leaving this setting as "All unassigned.") Also, verify that the **Authentication** option is set to **Automatic**. Click to clear the **Enable** check box for **Disconnect idle connections** (so that you aren't disconnected during OS installation). Click the **Enable** check box for **Multiple VMRC connections** and for **SSL 3.0/TLS 1.0 encryption**. If necessary, set the **SSL 3.0/TLS 1.0 certificate** to **Keep** or **Request** (if **Keep** is disabled). Make sure the hostname is the same as the name of the computer you're using. Click **OK** at the lower right. (If you have any problems using VMRC Server, you can come back to this window to adjust settings.)
10. In the left pane, point to **Configure** under **Virtual Machines** and click the name of the virtual machine you created.

11. Next, you need to turn on the virtual machine. In the right pane, click the thumbnail image for the virtual machine to turn on the virtual machine.



You might see a message that you need to configure Internet Explorer security to proceed. Make the necessary security configurations. Also, if you see an IE message to install an add-on, click the message and click **Install ActiveX Control**, and then follow the directions to continue.

12. Insert the Windows Server 2008 installation DVD. If necessary, click the thumbnail again for the virtual machine. If you see a security message, click **Yes** to proceed.
13. Enter your username and password (using an account with administrator privileges), and then click **OK**. If you see another security message, such as for NTLM Authentication, click **Yes** to proceed.
14. If necessary, scroll down to view information for working in the Remote Control window. Notice the options **Pause**, **Save State**, **Turn Off**, and **Reset** for the virtual machine. Scroll back to the top of the Remote Control window.
15. You should see a beginning installation window for Windows Server 2008. Move the mouse pointer into this window. (The mouse pointer becomes a small black dot.) Click in the window until you see the normal arrow for the mouse pointer. Notice that you can work only in the console for the virtual machine. Press the **right Alt** key (the default host key) to be able to use the mouse in the Remote Control window. Remember that you can always use the right Alt key to leave the console as needed. (Also, to work inside the console again, click the mouse inside the console.) At the upper right of the Remote Control window, click the **Remote Control** down arrow. Review the options on the menu, such as **Special Keys** and **Connect To Server**.



When you point to **Special Keys**, note that you can press the host key (right Alt key) with the Delete key to send the Ctrl+Alt+Delete key sequence to the virtual machine. (This is important to know later for logging on after you have installed Windows Server 2008.)

16. Move the mouse pointer back into the console and click so that you can work in this area again. You can now proceed with installing Windows Server 2008.
17. In the **Install Windows** window, click the language you want in the **Language to install** drop-down list. In the **Time and currency format** list box, make your selection, such as **English (United States)**. In the **Keyboard or input method** list box, make your selection, such as **US**, and then click **Next**. Click **Install now**.



If your connection stops before the installation is finished, use the left arrow at the top to go back to the main **Status** window. Click the virtual machine thumbnail to open a new connection via the Remote Control window. Respond to any security messages, log back on, and respond to any other security messages. The installation should still be running.

18. Click **Windows Server 2008 Enterprise (Full Installation)** or select another edition (such as **Standard Edition**, if it's available), and then click **Next**. Read the license terms, click the **I accept the license terms** check box, and then click **Next**.

19. Click **Custom (advanced)**. You'll see the amount of unallocated disk space selected, which is the disk space you specified when you configured the virtual machine. Make sure it's selected, and then click **Next**.
20. The Windows Server 2008 installation starts. You'll see progress information about copying files, expanding files, installing features, installing updates, and completing the installation. This part of the installation can take 30 minutes or longer. When it's finished, your computer restarts automatically.
21. The message "Please wait while Windows sets up your computer" is displayed, and then you see a window indicating that the installation is completing.
22. Your system restarts again. When you see a message that the user's password must be changed before logging on the first time, click **OK**. (You might have to click inside the console first to have the mouse function in it.)
23. Enter a new password for the Administrator account, and then enter the same password again to confirm it. Click the **blue circle** with the white arrow inside.



If you enter a password that isn't a strong password, you'll see the message "Unable to update the password." This means the password doesn't meet the length, complexity, or history requirements of the domain. Click OK and enter a different password that's longer than seven characters and uses letters, numbers, and special symbols, such as &.

24. When you see the message "Your password has been changed," click **OK**. The Windows desktop is displayed and the Initial Configuration Tasks applet opens. You can configure Windows Server 2008 as you would in a nonvirtual environment.
25. You can close the Remote Control window (the Virtual Machine Remote Control Server) or the Status window (the Virtual Server Administration Website) at any time. The virtual machine continues running in the background. Also, when in the Remote Control window, you can go back to the Administrator window by clicking the left arrow at the top of the Remote Control window.



You can shut down a server by first logging on through the Remote Control window. Also, you can use this window and the Status window to turn off a virtual machine (but make sure you shut down the server first).



To access the documentation for Microsoft Virtual Server, click Start, point to All Programs, click Microsoft Virtual Server, and click Virtual Server Administrator's Guide.

Installing an OS from an ISO Image

If you have an ISO image file for the guest OS, you have the option to install it instead of performing a traditional installation with an installation DVD. Here are the general steps for installing an ISO image file on a virtual machine in Microsoft Virtual Server:

1. Follow Steps 1 to 10 in the previous section, "Creating a Virtual Machine and Installing a Guest OS."

2. The bottom portion of the right pane should now show the configuration options for the virtual machine. Click the **CD/DVD** option.
3. Under **Virtual CD/DVD Drive 1**, click the **Known image files** option button. Next, click the **Known image files** list arrow and click the image file. If the ISO image file isn't listed, enter the path to the ISO image file in the "Fully qualified path to file" text box.
4. Click **OK** to return to the Master Status listing.

Configuring Networking and Hardware Options

You can use the Virtual Server Administration Website to configure virtual networks. As you learned earlier, a connected network has two default virtual network options: external network and internal network. You can customize settings for both types of networks, such as settings for a virtual DHCP server. You can also create a virtual network with properties you define.



A virtual network is one used by virtual machines in a network and is independent of other virtual networks. In Microsoft Virtual Server, the number of virtual machines connected to a virtual network is unlimited.

The Virtual Server Administration Website also has options for configuring hardware settings, such as adding memory for use by a virtual server. In the next sections, you learn how to configure virtual networking and configure hardware for a virtual machine.

Configuring Virtual Networking In the following steps, you see how to configure virtual networking:

1. Open the Virtual Server Administration Website, if necessary. (Click **Start**, point to **All Programs**, click **Microsoft Virtual Server**, and click **Virtual Server Administration Website**.)
2. In the left pane under **Navigation**, click **Master Status**, if necessary. Select each virtual server that's running (if any) and shut it down. To do this, point to the server name (that has a right arrow) under **Virtual Machine Name** in the right pane, click **Turn Off**, and click **OK**. (You can configure virtual networking while virtual machines are running, but turning them off first is recommended.)
3. In the left pane under **Virtual Networks**, point to **Configure** and click **View All**.
4. In the right pane, point to **External Network (NICname)** and click **Edit Configuration**. Review this information.
5. In the right pane, click the **Network Settings** link. Review the properties information, including information about the NIC, and then click **OK**.
6. In the right pane, click the **DHCP Server** link. You can use this pane to configure a virtual DHCP server that leases IP addresses through Microsoft Virtual Server (see Figure E-7). To enable the virtual DHCP server, click the **Enabled** check box. When you enable the virtual DHCP server, you can configure the following:
 - *Network address*—Enter the network address for the virtual network.
 - *Network mask*—Enter the network mask.

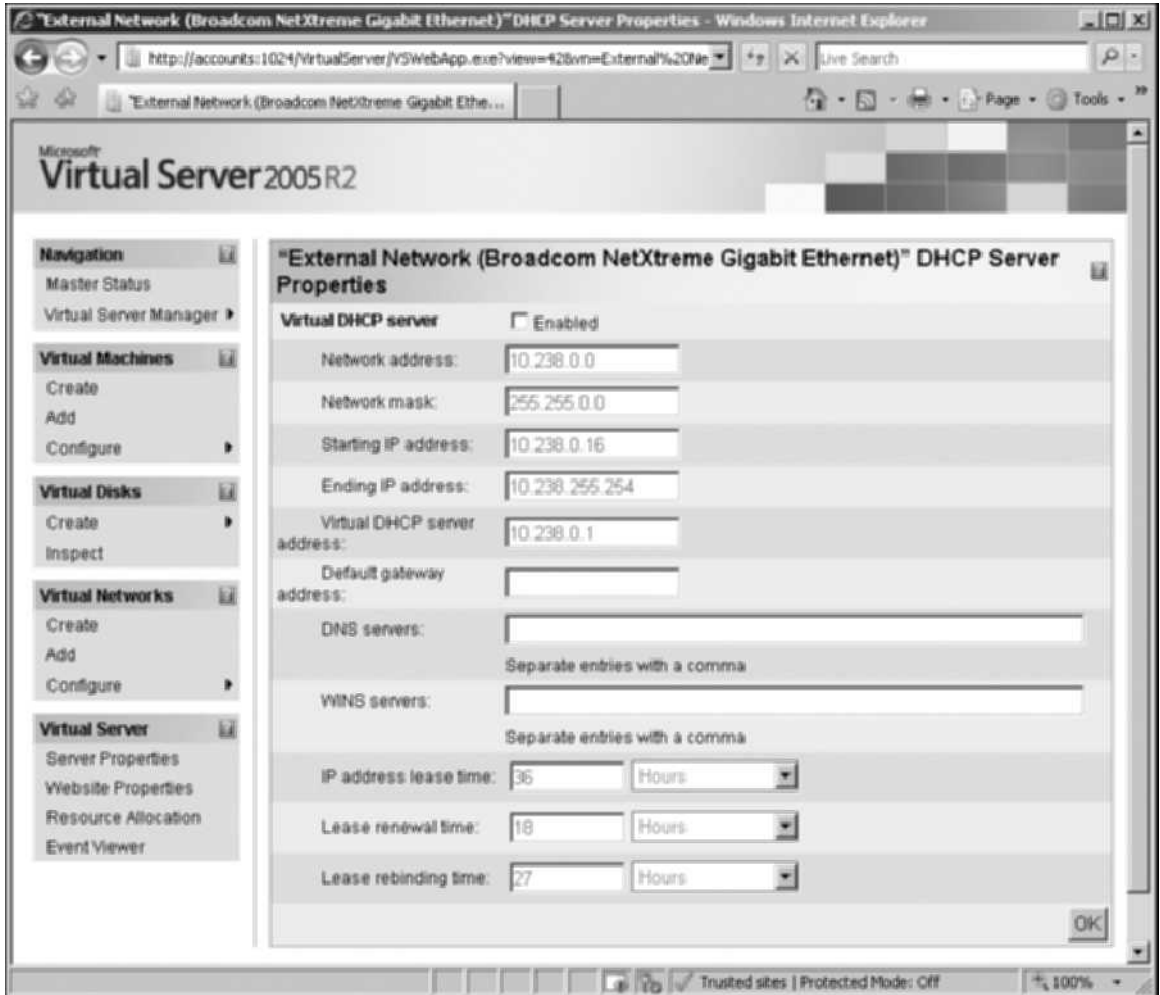


Figure E-7 Virtual DHCP server configuration options

Courtesy of Course Technology/Cengage Learning

- *Starting IP address*—Enter the beginning address for the range (scope) of IP addresses that can be leased.
- *Ending IP address*—Enter the ending address for the range of IP addresses that can be leased.
- *Virtual DHCP server address*—Enter the IP address of the virtual DHCP server.
- *Default gateway address*—Enter the IP address of a router that transports packets beyond the virtual network.
- *DNS servers*—Enter the IP address of one or more DNS servers already on the network.
- *WINS servers*—Enter the IP addresses of any Windows Internet Naming Service (WINS) servers (for converting NetBIOS computer names to IP addresses).
- *IP address lease time*—Enter the amount of time an IP address can be leased, which can be set in days, hours, minutes, or seconds. (Typically, you set it for one or more days.)

- *Lease renewal time*—Enter the amount of time in which the client can contact the virtual DHCP server to renew a lease (in days, hours, minutes, or seconds, but with a minimum of 30 seconds).
 - *Lease rebinding time*—Enter the amount of time it takes to allow the client to contact another server to renew its lease, when the main leasing server can't be reached (in days, hours, minutes, or seconds, but with a minimum of 45 seconds).
7. In the left pane under Virtual Networks, point to **Configure** and click **Internal Network**. Review the information for virtual network properties.
 8. Click **Network Settings** in the right pane and review the information, and then click the **back arrow** at the top.
 9. Click **DHCP Server** in the right pane, and notice that you can enable a virtual DHCP server and configure it. Leave this window open for the next steps.

Configuring Hardware for a Virtual Machine In addition to configuring a virtual network, you can configure hardware and other options for a virtual machine. In the following steps, you examine the options that can be configured:



The virtual machine you select in the steps that follow should be turned off before you start.

1. Make sure the Virtual Server Administration Website is open.
2. In the left pane under Virtual Machines, point to **Configure** and click the name of the virtual server you have configured.
3. Scroll to the configuration section in the right pane. Review the options that can be configured, which include the following:
 - General properties
 - Virtual Machine Additions
 - Memory
 - Hard disks
 - CD/DVD
 - SCSI adapters
 - Network adapters
 - Scripts
 - Floppy drive
 - COM ports
 - LPT ports
4. In the right pane, click **General properties**. If your computer supports hardware-assisted virtualization, notice that you can enable it here. You can also specify a user account under which to run the virtual machine, and you can specify what action to take when the virtual server stops. If you make changes, click **OK** at the lower left.

5. Click the **back arrow** at the top of the window to return to the previous configuration display in the right pane.
6. In the right pane, click **Memory**. Now you can change the amount of memory allocated to the virtual machine. If you make changes, click **OK**.
7. Click the **back arrow** at the top of the window.
8. In the right pane, click the link for **Hard disks**. In the right pane, you see the configuration of the virtual disk used by the virtual machine. Notice the option “Enable undo disks.” When you select this option, configuration and other changes on the virtual machine are saved so that you can undo those changes, if necessary. Also, notice that you can add a new virtual disk by clicking the Add disk button. If you make changes, remember to click **OK** so that they take effect.
9. Click the **back arrow**.
10. Click **CD/DVD** in the right pane. In the right pane, you can click the Remove check box to remove a CD/DVD drive, and you can click the Add CD/DVD Drive button to add a new drive. If you make changes, click **OK**.
11. Click the **back arrow**.
12. Click each of the remaining configuration options in the right pane to view what they cover. In particular, notice that you can add NICs by using the Network adapters option.
13. Close the Virtual Server Administration Website when you’re finished (or restart your virtual server so that it’s in use).

Host Key Options

Microsoft Virtual Server designates the right Alt key as the default host key and offers host key options that are similar to those in Microsoft Virtual PC. Table E-2 lists important host key combinations you can use while accessing a virtual machine.

Table E-2 Host key options for Microsoft Virtual Server

Key combination	Result
HostKey	Enables you to move the mouse outside the window used by the guest OS. (Move the mouse back into the guest OS display and click when you want to work in the guest OS.)
HostKey+Delete	The virtual machine OS responds to this key combination as Ctrl+Alt+Delete.
HostKey+C	Displays the Connect to server dialog box for connecting to a specific virtual machine. (If you have selected text first, it can be used to copy text.)
HostKey+A	Toggles to the Administrator display window.
HostKey+I	Shows the VMRC Connection Properties dialog box with information about the connected virtual machine.
HostKey+B	Provides information about the VMRC client software.
HostKey+V	Pastes text and items saved in the Clipboard into the active window in the guest OS.
HostKey+H	Enables you to configure a different key as the host key.

VMware Server

VMware Server enables you to set up virtual machines to run Windows or Linux OSs. VMware Server version 2 is a major update compared with previous 1.x versions. The new features of VMware Server 2 include the following:

- Managing virtual machines from the Web Access management interface or the VMware Remote Console
- Configuring different levels of permissions
- Configuring which OSs are started when VMware is started
- Editors for hardware devices
- New support for Windows Vista, Windows Server 2008, Red Hat Enterprise 5.0, and Ubuntu Linux through version 8.x
- Handling increased memory (to 8 GB) and more NICs (up to 10) in the host machine
- Supporting 64-bit guest OSs on 64-bit (x64) host computers
- Hot-add capability for new SCSI and tape devices (without shutting down a virtual machine)
- Supporting VSS for backups on Microsoft guest OSs
- Using Firefox 3 or Internet Explorer for the Web Access management interface
- Supporting hardware virtualization, such as through AMD CPUs with AMD-V capability and Intel CPUs with Intel VT
- Supporting multiple monitors (to see different virtual machines on different displays)

Guest OSs Supported

VMware Server supports the following guest OSs:

- Windows Server 2008 Standard, Enterprise, Datacenter, and Web Server (x86 or x64)
- Windows Server 2003 Standard, Enterprise, Datacenter, and Web Server with SP1 or SP2 (x86 or x64)
- Windows Server 2003 Standard, Enterprise, Datacenter, and Web Server R2 (x86 or x64)
- Windows Small Business Server 2003 (Standard and Premium Editions)
- Windows 2000 Server and Professional
- Windows XP Professional
- Windows Vista Business and Ultimate (x86 and x64)
- Red Hat Enterprise Linux Server and Desktop versions up through version 5 (x86 and x64)
- Ubuntu Linux 6.x to 8.x
- SUSE Linux Enterprise Server up to 10.x (x86 and x64)
- SUSE Linux versions up to 10.x (x86 and x64)
- Novell NetWare
- Solaris

Host OSs Supported

VMware Server 2.x runs on more different host OSs than Microsoft Virtual PC or Server because it can run on several Linux distributions. It also runs on x86 and x64 computers. VMware host OSs includes the following:

- Windows Server 2008 Standard, Enterprise, Datacenter, and Web Server (x86 or x64)
- Windows Server 2003 Standard, Enterprise, Datacenter, and Web Server with SP1 or SP2 (x86 or x64)
- Windows Server 2003 Standard, Enterprise, Datacenter, and Web Server R2 (x86 or x64)
- Windows Small Business Server 2003 (Standard and Premium Editions)
- Windows 2000 Server and Professional with SP3 or SP4
- Windows XP Professional and Home through the current service pack
- Windows Vista Business and Ultimate (x86 and x64)
- Red Hat Enterprise Linux Server and Desktop versions up through version 5 (x86 and x64)
- Ubuntu Linux 6.x to 8.x
- SUSE Linux Enterprise Server up to 10.x (x86 and x64)
- SUSE Linux versions up to 10.x (x86 and x64)
- Mandrake Linux up to 10.x

VMware Server also can run on other Windows and Linux distributions, such as other Windows Vista editions or Fedora Linux, but they should be considered “experimental” because they might not be fully tested.



For Windows host OSs, you must download the VMware Server version for Windows, which is in .exe format. For Linux host OSs, you must download the VMware Server version for Linux, which is in .tar format.



Windows Server Core isn't a supported host at this writing.

Requirements for VMware Server

VMware Server has the following hardware requirements:

- *CPU*—Any standard x86 or x64 computer, including the following processors: dual-core or quad-core Intel Xeon, Intel Core 2, AMD Opteron, or Athlon (733 MHz or faster)
- *RAM*—A minimum of 512 MB but must include enough RAM for at least the minimum requirements of the total number of OSs you'll be running (host and guest)
- *Disk space*—Enough disk storage for the OSs you plan to run (host and guest)
- *Web Access console*—Internet Explorer 6.0 or later (for Windows hosts) or Mozilla Firefox 2.0 or later (for Linux hosts)



VMware Server 2.x virtual machines can connect to hard, optical, and floppy drives. VMware 2.x also supports USB 2.x connections.

Downloading VMware Server

VMware Server can be downloaded from VMware's Web site at no cost by following these steps.

1. Log on to your computer, and create a folder for storing the download file for VMware Server.
2. Start your Web browser, and go to www.vmware.com/products/server.



Web links and instructions change periodically. You might need to search for the most current link at www.vmware.com if this link doesn't work.

3. Click **Download Now**. Find the latest version of VMware Server (if multiple versions are listed) and click **Download** or **Download Now**.
4. If asked to provide registration information, fill out the registration form. Read the licensing information and click **Yes** or **Accept**. Record the serial number for the Windows version (used later when you install VMware Server).
5. Click the link to download the Binary (.exe) file for VMware Server for Windows Operating Systems.
6. Save the file in the folder you created in Step 1. When the download is completed, click **Close** in the Download complete dialog box, and exit your Web browser.

Installing VMware Server

The general steps for installing VMware Server on the host OS are as follows:

1. If possible, connect to the Internet so that updates can be installed automatically during the installation process.
2. Browse to the folder where you saved the downloaded file for VMware Server, and double-click **VMware-server-2.x.x-xxxxxx** (replacing 2.x.x-xxxxxx with the VMware Server version you downloaded).
3. You see a message about preparing for the installation, and then the Windows Installer dialog box opens. In the welcome window of the Installation Wizard for VMware Server (see Figure E-8), click **Next**.
4. Read the license agreement, click **Yes, I accept the terms in the license agreement**, and then click **Next**.
5. Verify that the VMware server files will be written to the correct destination folder (or click the **Change** button to select another destination), and then click **Next**.
6. Verify the fully qualified domain name for the host computer, and make sure the server HTTP (8222) and server HTTPS (8333) ports are selected by default. Make any changes as needed, such as the host and domain names (but leave the defaults for the ports), and then click **Next**.
7. Make sure the shortcuts you want are selected, as shown in Figure E-9, and then click **Next**.



Figure E-8 The Installation Wizard for VMware Server

Courtesy of Course Technology/Cengage Learning




Figure E-9 Shortcut options

Courtesy of Course Technology/Cengage Learning

8. Click **Install**. You'll see a message that the installation might take several minutes. If you see any messages about installing device software, click **Install**.
9. For the registration information, enter your name and the name of your company or school (if necessary). Enter the serial number you wrote down in the previous steps, and then click **Enter**

- Click **Finish**. Make sure you exit all other programs, and then click **Yes** to restart the system.

Creating a Virtual Machine and Installing a Guest OS

Now that VMware Server is installed, the next step is to create a virtual machine and install the guest OS. Here are the general steps, using Windows Server 2008 as the guest OS:



The VMware Remote Console you use later in these steps requires that the VMware virtual server (host computer) be resolved via DNS. Before you start, make sure your server can be resolved via DNS on your network (or DNS is installed on the host). For example, there should be a host address (A) resource record in the DNS server for the host computer.

- Double-click the **VMware Server Home Page** icon on the desktop or the taskbar. (You can also click **Start**, point to **All Programs**, click **VMware Server**, and click **VMware Server Home Page**.)



You might need to configure security requirements for Internet Explorer, such as providing a digital certificate, specifying whether to set up a phishing filter, and adding this site as a trusted site.

- Log on with your host computer account name (or the administrator account) and enter the password. (Use the same account you used to install VMware Server.) You see the VMware Infrastructure Web Access console, shown in Figure E-10.



Figure E-10 The VMware Infrastructure Web Access console

Courtesy of Course Technology/Cengage Learning



A certificate error is reported in Figure E-10 because this new site doesn't yet have a trusted certificate. If you have this problem, you might be able to import a certificate by clicking the Certificate Error box at the top, clicking the View certificates link, and clicking Install Certificate. Another option is to talk to your network administrator about importing a certificate.

3. Make sure your host computer is selected in the left pane, and click the **Virtual Machines** tab.
4. In the right pane under the Commands heading, click **Create Virtual Machine**. Enter the virtual machine name, and then click **Next**.
5. Make sure **Windows operating system** is selected for the guest OS, click the OS (in the Version list box), and then click **Next**.
6. Set the memory size to **512 MB** or higher. (1024 MB is the default when installing Windows Server 2008.) Also, if your system has a dual-core or quad-core CPU or is an SMP system, you can select the number of processors to use. Notice, however, that you shouldn't reconfigure the setting for number of processors after the virtual machine is set up. Click **Next**.
7. Select the virtual disk to use, such as by clicking **Create a New Virtual Disk** (a disk on the current computer). (The other option is Use an Existing Virtual Disk, which is a disk on a shared drive or hard disk on a different computer.) Enter the capacity for the virtual disk, such as **20 GB** (see Figure E-11). Adjust any of the following settings as needed, and then click **Next**:
 - *Location for the virtual disk file*—Specify a file location other than the default.
 - *File Options*—Allocate disk space now or split the disk into two files.
 - *Disk Mode*—Create independent disks not affected by snapshots.
 - *Virtual Device Node*—Select the SCSI or IDE adapter and device.
 - *Policies*—Optimize for safety (the default) or for performance.

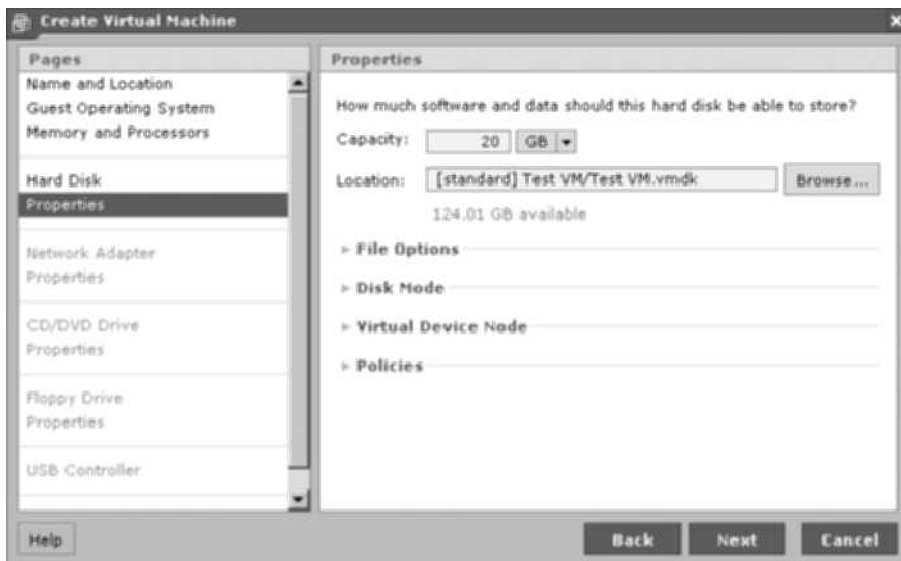


Figure E-11 Configuring virtual disk properties

Courtesy of Course Technology/Cengage Learning

8. In the next window, you can add a network adapter for access over a network. Click **Add a Network Adapter**. If you don't want to use the default settings for Network Connection (Bridged) and Connect at Power On (Yes), change these settings. The network settings you can configure are as follows; when you're finished, click **Next**:
 - *Bridged*—This setting gives the virtual machine its own network identity (so that it's seen as a different computer from the host), which enables other computers on the network to communicate with it. It also means the virtual machine can access the Internet through the local network.
 - *HostOnly*—With this setting, only the host computer and other virtual machines on the same host can access the virtual machine, which means the virtual machine isn't accessible through the local network.
 - *NAT*—The virtual machine and host use the same IP and MAC addresses, which means the virtual machine doesn't have its own identity on the local network. You might choose this setting if IP addresses are in short supply for the network or because an organization's network policy is to allow only one IP address per computer.
9. You can specifying a CD/DVD drive or an ISO image for the OS installation. For this activity, click **Use a Physical Drive**. Make sure the correct CD/DVD drive is selected, such as drive E, and verify that Connect at Power On is set to **Yes**. Click **Next**.
10. If your computer has a floppy drive, you can configure it to provide an image for the OS. Select the configuration options. (To install Windows Server 2008, click **Don't Add a Floppy Drive**.) Click **Next**, if necessary. (Depending on your selection, you might need to configure additional properties.)
11. In the next window, you can specify adding a USB controller, such as for accessing a flash drive. Make your selection and click **Next**, if necessary.
12. Review your selections, and then click **Finish**. In the bottom pane, you should see **Success** displayed in the Status column to show that you created the virtual machine.



TIP

In some cases, if you have selected different configuration options and then clicked the **Back** button to return to the preceding steps, VMware Server might give you an error message or you might not end up with an installed virtual machine. If this happens, start from scratch and avoid undoing selections you have made.

13. Insert the Windows Server 2008 installation DVD. In the left pane, click the new virtual machine name under the host server name. (You might have to expand entries under the host server name first.)
14. Click the **Summary** tab, if necessary. In the right pane, scroll to the **Hardware** section. Click the **CD/DVD Drive 1** down arrow and click **Edit**.
15. Review the settings for the host media (CD/DVD drive), make any needed changes, and click **OK**.
16. Click the **Console** tab, and then click **Install plug-in** to install the Remote Console plug-in.



If you see a message box about noticing the Information Bar, click Close. Also, if the plug-in isn't installed successfully in Internet Explorer, you might see a message at the top that you must click to continue. Click the message and click to install the elements IR requires, such as Install the ActiveX Control. Next, click Install plug-in again, and, if necessary, click Install.

17. In the right pane, click **Powered off** (which is like a switch to turn the virtual machine on or off). Click anywhere in the reduced console area in the right pane.
18. In the Install Windows window, click the language you want in the Language to install drop-down list. In the Time and currency format list box, make your selection, such as English (United States). In the Keyboard or input method list box, make your selection, such as US, and then click **Next**. Click **Install now**.
19. Click **Windows Server 2008 Enterprise (Full Installation)** or select another edition (such as Standard Edition, if it's available), and then click **Next**. Read the license terms, click the **I accept the license terms** check box, and then click **Next**.
20. Click **Custom (advanced)**. You'll see the amount of unallocated disk space selected, which is the disk space you specified when you configured the virtual machine. Make sure it's selected, and then click **Next**.
21. The Windows Server 2008 installation starts. You'll see progress information about copying files, expanding files, installing features, installing updates, and completing the installation. This part of the installation can take 30 minutes or longer. When it's finished, your computer restarts automatically.
22. The message "Please wait while Windows sets up your computer" is displayed, and then you see a window indicating that the installation is completing.
23. Your system restarts again. When you see a message that the user's password must be changed before logging on the first time, click **OK**. (You might have to click inside the console first to have the mouse function in it.)
24. Enter a new password for the Administrator account, and then enter the same password again to confirm it. Click the **blue circle** with the white arrow inside.



If you enter a password that isn't a strong password, you'll see the message "Unable to update the password." This means the password doesn't meet the length, complexity, or history requirements of the domain. Click OK and enter a different password that's longer than seven characters and uses letters, numbers, and special symbols, such as &.

25. When you see the message "Your password has been changed," click **OK**. The Windows desktop is displayed and the Initial Configuration Tasks applet opens. You can configure Windows Server 2008 or log off and use the Remote Control window later to access Windows Server 2008.
26. You can close the VMware Remote Console window at any time (but note that the virtual machine keeps running). Close the VMware Infrastructure Web Access console when you are finished using it. (The virtual machine continues running, unless you shut it down in the VMware Remote Console window and power it off in the VMware Infrastructure Web Access window.)



To access online help documentation while you're in the VMware Infrastructure Web Access console, click the Help option at the upper right.

Installing an OS from an ISO Image

VMware Server supports installing an OS via an ISO image file. The general steps are as follows:

1. Follow the steps to create a virtual machine. In the Inventory pane of the VMware Infrastructure Web Access console, click the virtual server you have created.
2. Click the **Summary** tab, and scroll down to view the Hardware section.
3. Click the **CD/DVD Drive 1** down arrow and click **Edit**.
4. In the Connection section, click the **ISO Image** option button. Enter the optical disk image path or use the **Browse** option to find and select it.
5. If necessary, select the device node in the Virtual Device Node section. Click **OK**.
6. Click the **Console** tab.
7. Power on the virtual machine, if necessary. Click inside the console and follow the OS installation instructions.

Configuring Networking Options

As you have learned, the three network connection options are Bridged, HostOnly, and NAT, which have the following default names:

- Bridged is called VMnet0.
- HostOnly is called VMnet1.
- NAT is called VMnet8.

You can configure virtual networking, including VMnet0, VMnet1, and VMnet8, in the Virtual Network Editor. For example, you can configure VMware internal DHCP server capability for HostOnly and NAT networks. Bridged networks use an external DHCP server, such as a Windows Server 2008 server configured for this service. To explore the Virtual Network Editor, follow these steps:

1. Click **Start**, point to **All Programs**, click **VMware**, click **VMware Server**, and click **Manage Virtual Networks**. The Virtual Network Editor has the following tabs (see Figure E-12):
 - *Summary*—Shows a summary of the virtual networks, including VMnet0, VMnet1, and VMnet8
 - *Automatic Bridging*—Controls bridging between the VMnet0 network and the network adapter
 - *Host Virtual Network Mapping*—Enables you to link virtual networks to physical network adapters and virtual network adapters as well as configure subnet and DHCP properties

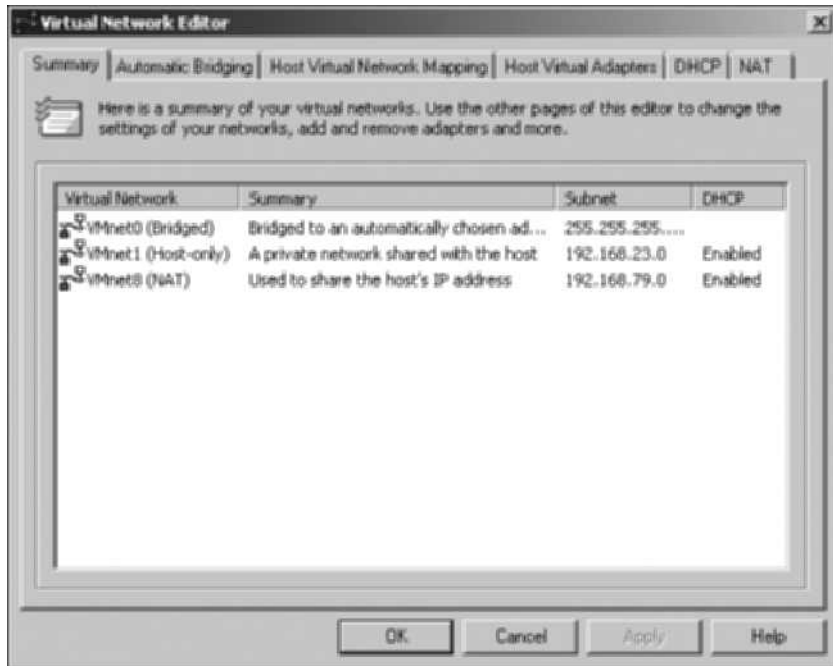


Figure E-12 The Virtual Network Editor

- *Host Virtual Adapters*—Shows virtual adapter connections, virtual networks, and the status of connections
 - *DHCP*—Enables you to configure DHCP for VMnet1 and VMnet8
 - *NAT*—Enables you to control the NAT service and configure NAT settings
2. Click each tab to view what's available. When you're finished, click the **DHCP** tab again. Notice that you can start, stop, and restart the DHCP service in the DHCP tab.
 3. Click **VMnet1** and click **Properties**. In the DHCP Settings dialog box, notice that you can configure the range of IP addresses to use. You can also configure the lease duration parameters for clients. Click **Cancel**.
 4. Click the **NAT** tab. You can use this tab to associate the NAT service with a virtual network and to start, stop, and restart the NAT service.
 5. When you're finished, close the Virtual Network Editor.

Configuring Hardware Options

After you set up a virtual machine, you might want to go back and configure hardware options. For example, you might change the configuration of the network and decide to go from a Bridged network to a HostOnly network. Follow these general steps to configure hardware:

1. Open the VMware Infrastructure Web Access console. In the Inventory pane, expand the virtual machines under the host server, if necessary. Click the virtual machine you want to configure.

2. To configure hardware, first make sure the virtual machine is turned off. Click the **Console** tab, and shut down the virtual machine. You can also click **Virtual Machine** on the toolbar and click **Power Off**.
3. Click the **Summary** tab, and scroll down to view the Hardware section.
4. Click the **Processors** down arrow and click **Edit**. You'll see a note that advises against changing the number of virtual processors, if you have more than one processor. Click **Cancel**.
5. Click the **Memory** down arrow and click **Edit**. Notice the recommended size information for memory allocation. You can use the Size (in multiples of 4) text box to change the memory allocation. Click **Cancel**.
6. Click the **Hard Disk 1** down arrow and click **Edit**. You can increase the virtual disk capacity, configure the virtual device node, configure the disk mode, and configure policies. Click **Cancel**.
7. Click **Network Adapter 1** and click **Edit**. You can change the type of network connection and view information about the connection status, MAC address, and virtual device. Click **Cancel**.
8. Click the **CD/DVD Drive 1** down arrow and click **Edit**. Review the properties you can set and the connection status information. Click **Cancel**.
9. Review information about any other hardware devices, and when you're finished, restart the virtual machine.

Installing VMware Tools

VMware Tools is an add-on that gives you more ways to manage a virtual machine and improve its performance. Its components include the following:

- A control panel to change virtual machine settings and connect devices conveniently
- VMware user processes for Linux and Solaris guest OSs
- Device drivers for enhanced video, audio, mouse, network, and SCSI disk performance
- Tools service that includes a variety of tools for messaging, mouse performance, screen resolution, and others

When you install VMware Tools, the virtual machine must be started and you should be logged on to the guest OS account from which you manage VMware Server. The reason is that VMware Tools, including drivers, is installed on the guest OS, and you can access it from Control Panel in Windows Server 2008 (and other Windows OSs). To install VMware Tools, follow these steps:

1. Open the VMware Infrastructure Web Access console. In the Inventory pane, click a virtual machine.
2. If the guest OS isn't running, start it. Log on to the Administrator account or an account with administrator privileges.
3. In the VMware Infrastructure Web Access console, click **Install VMware Tools** in the Status column in the right pane, and then click **Install**.
4. Open the virtual machine console by clicking the **Console** tab and clicking inside the console.

- When the AutoPlay message box opens in the guest OS desktop (which might take several minutes), click **Run setup.exe**. You see the message “Preparing to install,” and this process might take several minutes.
- Click **Next** in the Welcome to the installation wizard for VMware Tools window (see Figure E-13).



Figure E-13 The installation wizard for VMware Tools

Courtesy of Course Technology/Cengage Learning

- Select the setup type option from the following options, and then click **Next**:
 - Typical*—If you plan to use only VMware Server
 - Complete*—If you plan to use VMware Server and other VMware products
 - Custom*—If you want to choose the specific features to install
- Click **Install**. If you see the message “Windows can’t verify the publisher of the driver software,” click the option **Install this driver software anyway**. (You might see this message several times.)

9. If you see a Windows Security dialog box asking whether you want to install this device software, click the **Always trust software from “VMware, Inc.”** check box. Click **Install**, and then click **Finish**.
10. Save any work you have open on the virtual machine and click **Yes** to restart.
11. Log back on to the guest OS in the console window. In the guest OS (Windows Server 2008), click **Start, Control Panel**.
12. Click **Classic View** and click the new applet **VMware Tools**. The VMware Tools Properties dialog box opens (see Figure E-14). Click each tab to see what settings are available.



Figure E-14 The VMware Tools Properties dialog box

Courtesy of Course Technology/Cengage Learning

13. Click the **Help** button to learn more about VMware Tools capabilities. When you're finished, close the VMware Tools Help window.
14. Click **Cancel** to close the VMware Tools Properties dialog box. Notice that a new icon is displayed on the guest OS's taskbar, which can be used to open the VMware Tools Properties dialog box. Close Control Panel in the guest OS.

Other Virtual Systems

This appendix has focused on free virtualization systems. Other systems are available at a cost. On the desktop side, VMware Workstation has grown in use along with desktop virtualization. Another system is Microsoft Hyper-V, which is new to Windows Server 2008. The

following sections give you a brief overview of these systems but are not intended to provide instructions about how to use them.



VMware Workstation is free for academic institutions approved in the VMware Academic Program. Entry in this program is free for two-year and four-year degree-granting higher education institutions and accredited technical schools. For more information, visit <http://vmware.com/partners/academic>.

VMware Workstation

VMware Workstation is popular among software developers and testers because it provides a safe environment in which to write and test development software before it's released to live production. It's also used by people who need to run multiple OSs on one workstation-class computer, including legacy OSs. This can be useful for running old software without having to convert it for a new OS. It's also useful for learning a new OS.

VMware Workstation 6.04 (and later) supports Windows, Linux, and other OSs as host and guest OSs. Newer OSs supported as both hosts and guests include the following:

- Windows Server 2008 Standard, Enterprise, and Datacenter editions (x86 and x64)
- Windows Vista Home Basic, Home Premium, Enterprise, Business, and Ultimate (x86 and x64)
- Red Hat Enterprise Linux up to 4.6 (x86 and x64)
- Ubuntu Linux up to 7.10 (x86 and x64)
- SUSE Linux Enterprise Server 10 (x86 and x64)
- openSUSE Linux up to 10.3

VMware Workstation has several of the same new features as VMware Server, which include the following:

- Handles increased memory (to 8 GB)
- Supports 64-bit guest OSs on 64-bit host computers
- Supports hardware virtualization, such as through AMD CPUs that have AMD-V capability and Intel CPUs with Intel VT
- Supports USB 2.0 (including on Linux OSs)
- Supports multiple monitors (to see different virtual machines on different displays)

As with VMware Server, you can configure hardware for the virtual machine, including multiple processors, memory, hard disks, USB access, floppy access, and other hardware elements. You can also configure Bridged, HostOnly, and NAT virtual networks. A virtual DHCP server can be configured when you use HostOnly and NAT virtual networking. Setting up a virtual machine is also done with a step-by-step wizard.

Also, as in VMware Server, you can install VMware Tools, which includes specialized drivers, such as drivers for enhanced video and audio functions for the guest OS. VMware Workstation has a console for accessing the guest OS that resembles the VMware Server console.

VMware Workstation is specifically designed for workstation host machines and offers a wider range of host and guest OS compatibility than Microsoft Virtual PC (at this writing). You can download a 30-day free evaluation version at www.vmware.com/products/ws.

Microsoft Hyper-V

Microsoft Hyper-V was released just a few months after Windows Server 2008. Unlike the other virtualization systems discussed in this appendix, Microsoft Hyper-V is intended to run only on Windows Server 2008. It's loaded through Server Manager like any other role in Windows Server 2008. In this regard, Windows Server 2008 offers perhaps the smoothest installation process of any virtual system discussed in this appendix. Also, unlike the other systems discussed, Hyper-V runs only on x64 computers, which means the host OSs include only the following:

- Windows Server 2008 Standard Edition x64
- Windows Server 2008 Enterprise Edition x64
- Windows Server 2008 Datacenter Edition x64

You can purchase any of Windows Server 2008 Standard, Enterprise, or Datacenter Editions with Hyper-V (for an extra \$28 at this writing), or you can purchase Hyper-V separately (also for \$28). The low cost and seamless installation and integration with Windows Server 2008 are designed to make this virtualization system particularly appealing to Windows Server 2008 users. The guest OSs that can be installed in Hyper-V include the following:

- Windows Server 2008 Standard, Enterprise, Datacenter, and Web Server (x86 or x64)
- Windows Server 2003 Standard, Enterprise, and Datacenter (x86 or x64)
- Windows Server 2003 Web Edition
- Windows 2000 Server and Advanced Server with SP4
- Windows Vista Business, Enterprise, and Ultimate (x86 and x64)
- Windows XP Professional with SP2 or SP3 (x86)
- Windows XP Professional with SP2 (x64)
- SUSE Linux Enterprise Server 10 with SP1 or SP2 (x86 or x64)

After Hyper-V is installed as a server role, you can open Hyper-V Manager as a Microsoft Management Console (MMC) snap-in or from the Administrative Tools menu—steps familiar to Windows Server 2008 administrators. Hyper-V Manager is easy to use because it's designed in the same format as most Windows Server 2008 administrative tools. For example, to create a virtual machine, click the New option in the right pane and follow the steps in the New Virtual Machine Wizard.

To configure hardware and management settings for a virtual machine, click Settings under the name of the virtual machine in the right pane of Hyper-V Manager. You can use the Settings dialog box (see Figure E-15) to add hardware, configure hardware, and configure management capabilities.

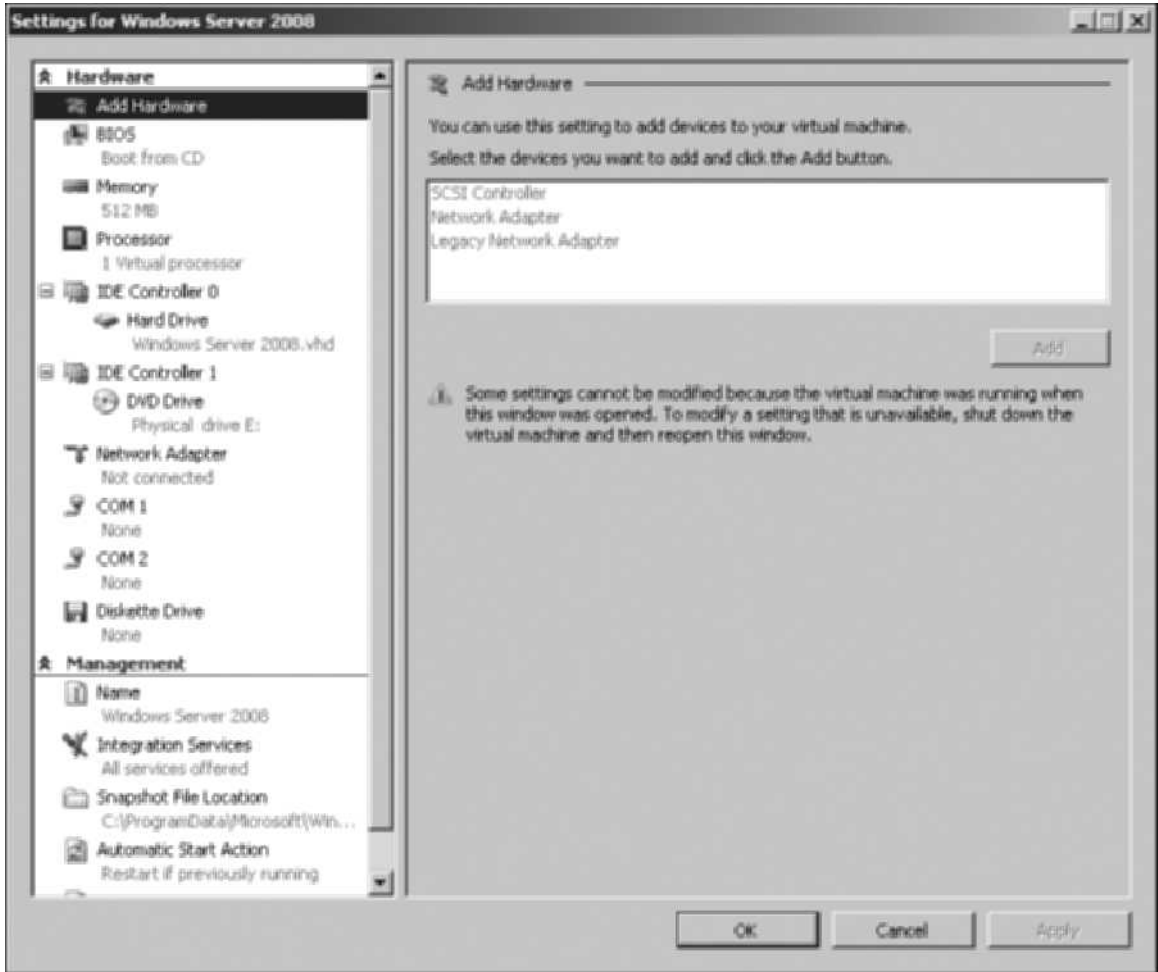


Figure E-15 Configuring settings for a virtual machine

Courtesy of Course Technology/Cengage Learning

You can access the Virtual Network Manager dialog box from Hyper-V Manager to configure a virtual network. There are three types of virtual networks:

- *Private*—Communication only between virtual machines on the same virtual server
- *Internal*—Communication between virtual machines and the host virtual server
- *External*—Communication between virtual machines and the physical network (using a network adapter)

For an external virtual network, you can specify a virtual LAN identification number. This unique number is used for communication through the network adapter that distinguishes the virtual network from other networks.

The guest OS appears in a console that has an Action menu, where you can send a Ctrl+Alt+Delete keystroke for logging on and start, turn off, shut down, or pause a virtual

machine (as well as other options). You also can expand the console to fill the desktop display. To start the console, click its thumbnail. When the console opens, it displays a message about how to start the guest OS.

At this writing, Hyper-V doesn't have as extensive a range of guest and host OSs as other virtualization systems. However, it's a good fit with Windows Server 2008 environments, and more guest OSs will likely be added in the future. Hyper-V installation and administration are consistent with how other server roles are installed and administered.

Glossary

1000BaseT Ethernet A technology defined by the IEEE 802.3ab standard, supports 1000 Mbps Ethernet (usually called Gigabit Ethernet) over Category 5 or higher UTP cable, using baseband signaling.

100BaseFX 100 Mbps Ethernet using baseband signaling over two strands of fiber-optic cabling.

100BaseTX A technology defined by IEEE 802.3u, it's the most commonly used Ethernet variety today. It runs over Category 5 or higher UTP cable and uses two of the four wire pairs: one to transmit data and the other to receive data. It runs at 100 Mbps, using baseband signaling.

10BaseT A technology defined by IEEE 802.3i, it's Ethernet running at 10 Mbps, using baseband signaling over Category 3 or higher twisted-pair cabling. Although still seen in older networks, newer networks use 100BaseT or faster technology.

10GBaseT A technology defined by IEEE 802.3an, it's 10 Gigabit Ethernet running over four pairs of Category 6A UTP cabling, using baseband signaling. Unlike the other BaseT Ethernet standards, 10GBaseT operates only in full-duplex mode.

802.11i A security extension to 802.11 and a successor to Wi-Fi Protected Access; currently the strongest security protocol for wireless networks. *See also* Wi-Fi Protected Access (WPA).

access control In the context of the Network layer and routing, the process by which a router consults a list of rules before forwarding an incoming packet. The rules determine whether a packet meeting certain criteria (such as source and destination address) should be permitted to reach the intended destination.

access control list (ACL) A set of rules configured on a router's interface for specifying which addresses and protocols can pass through the interface and to which destinations.

access point (AP) A wireless device that serves as the central connection point of a wireless LAN and mediates communication between wireless computers.

active partition A partition that can hold boot files the BIOS loads before it can start the OS.

Address Resolution Protocol (ARP) An Internetwork-layer protocol used to resolve a host's IP address to its MAC address. ARP uses a broadcast frame containing the target host's IP address, and the host that's assigned the address responds with its MAC address.

address space The number of addresses available in an IP network number that can be assigned to hosts.

ad hoc mode Sometimes called peer-to-peer mode, it's a wireless mode of operation typically used only in small or

temporary installations. There's no central device, and data travels from one device to another to reach the destination device.

aging time The amount of time a switch maintains a switching table entry that hasn't been updated.

analog signal A signal, represented by a sine wave, that varies over time continually and smoothly.

Application layer Layer 7 in the OSI model provides interfaces that enable applications to request and receive network services. *See also* Open Systems Interconnection (OSI) reference model.

application servers Computers that supply the server side of client/server applications, and often the data that goes along with them, to network clients.

ARP cache A temporary storage location in an IP host's RAM that keeps recently learned IP address/MAC address pairs so that the ARP protocol isn't necessary for each packet sent to a host.

Asymmetric DSL (ADSL) A DSL variation in which the download and upload speeds differ substantially, so the data rates aren't symmetrical. Typical connection speeds for downloading data range from 256 Kbps to 8 Mbps; upload speeds are typically much slower, in the range of 16 Kbps to 640 Kbps. *See also* Digital Subscriber Line (DSL).

Asynchronous Transfer Mode (ATM) A high-speed, cell-based packet-switching technology designed for both LAN and WAN use; uses connection-oriented switches to allow senders and receivers to communicate over a network.

attenuation Weakening of a signal as it travels the length of the medium.

authentication The process of identifying who has access to the network. The most common form of authentication is a logon with a username and password.

authorization The process of granting or denying an authenticated user's access to network resources.

automatic link aggregation A feature that enables you to install multiple NICs in one computer and aggregate the bandwidth so that, for example, you can install two 1 Gbps NICs and have a total bandwidth of 2 Gbps to and from that computer.

Automatic Private IP Addressing (APIPA) A private range of IP addresses assigned to an APIPA-enabled computer automatically when an IP address is requested via DHCP but no DHCP server responds to the request. *See also* Dynamic Host Configuration Protocol (DHCP).

auto-MDIX A switch port option used to detect the type of device and cable the switch port is connected to; if necessary,

the port swaps its transmit and receive pins, which enables you to use a straight-through or crossover cable regardless of the type of device you're connecting to the port.

auto-negotiate mode Communication between a switch and a device connected to a switch port, in which the switch attempts to set the port's operating mode to the highest performance setting the device supports.

backbone cabling Network cabling that interconnects telecommunications closets and equipment rooms. This cabling runs between floors or wings of a building and between buildings to carry network traffic destined for devices outside the work area. It's often fiber-optic cable but can also be UTP. Also called "vertical cabling."

backdoor A program installed on a computer that permits access to the computer, thus bypassing the normal authentication process.

bandwidth sharing A network design in which interconnecting devices allow only one connected device to transmit data at a time, thus requiring devices to share available bandwidth.

bare-metal virtualization The hypervisor implements OS virtualization by running directly on the host computer's hardware and controls and monitors guest OSs. *See also* virtualization.

baseband A type of signaling used in networks, in which each bit of data is represented by a pulse of electricity (on copper media) or light (on fiber-optic media). These signals are sent at a single fixed frequency, using the medium's entire bandwidth. LAN technologies use baseband signaling.

baseline A record of performance data gathered when a system is performing well under normal operating conditions. The baseline can then be compared with data collected during peak resource demands to give you insight into your system's capabilities and limitations.

basic disk A disk configuration in which the space on the disk can be divided into one to four partitions.

Basic Rate Interface (BRI) An ISDN format that consists of two 64-Kbps B-channels and a 16-Kbps D channel; generally used for remote connections. *See also* Integrated Services Digital Network (ISDN).

batch file A text file containing a list of commands you ordinarily type at the command prompt.

blocking mode A mode on a switch port that prevents the switch from forwarding frames out the blocked port, thereby preventing a switching loop. *See also* switching loop.

boot partition The partition or logical drive holding Windows OS files.

broadband A type of signaling that uses analog techniques to encode binary 1s and 0s across a continuous range of values. Broadband signals move across the medium in the form of

continuous electromagnetic or optical waves rather than discrete pulses. Signals flow at a particular frequency, and each frequency represents a channel of data, allowing multiple streams of data on a single wire. TV and cable Internet use broadband signaling.

broadcast domain The scope of devices to which broadcast frames are forwarded. Router interfaces delimit broadcast domains because they don't forward broadcasts, whereas switches and hubs do.

broadcast frame A network message intended to be processed by all devices on a LAN; has the destination address FF:FF:FF:FF:FF:FF.

broadcast storm A condition that occurs when a broadcast frame is forwarded endlessly in a switching loop. *See also* switching loop.

bus A collection of wires that carry data from one place to another on a computer's motherboard.

bus mastering A feature that allows a network adapter to take control of the computer's bus to initiate and manage data transfers to and from the computer's memory, independent of the CPU.

cable plant The collection of all cables and connectors tying a network together.

cable segment A length of cable between two network devices, such as a NIC and a switch. Any intermediate passive (unpowered) devices, such as wall jacks, are considered part of the total segment length.

Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) An access control method used by Wi-Fi networks, in which an acknowledgement is required for every packet sent, thereby avoiding most possibilities of a collision (collision avoidance).

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) A media access method in which a device must first listen (carrier sense) to the medium to be sure no other device is transmitting. If two devices transmit at the same time (multiple access), a collision occurs and is detected (collision detection). In this case, all devices involved in the collision wait for a random period of time before transmitting again.

channel service unit/data service unit (CSU/DSU) A device that creates a digital connection between a LAN device, such as a router, and the WAN link from the service provider.

circuit-switched WAN A type of WAN connection in which a temporary dedicated connection is established between sender and receiver on demand.

Classless Interdomain Routing (CIDR) A method of IP addressing in which the network and host IDs are determined by a prefix number that specifies how many bits of the IP address are network bits; the remaining bits are host bits.

clear to send (CTS) A signal an AP generates in response to a request-to-send signal. A CTS signal indicates that the computer that sent an RTS can transmit data. *See also* access point (AP) *and* request to send (RTS).

client Term used to describe an OS designed mainly to access network resources, a computer's primary role in a network (running user applications and accessing network resources), and software that requests network resources from servers.

client-to-gateway VPN mode This VPN mode establishes a VPN connection between a single client computer and a VPN device.

cloud computing A networking model in which data, applications, and processing power are managed by servers on the Internet, and users of these resources pay for what they use rather than for the equipment and software needed to provide resources.

cloud storage A data storage method in which some or all of an organization's data is stored on servers located offsite and maintained by a storage hosting company.

collision The result of two or more devices on the same medium transmitting simultaneously when CSMA/CD is the media access method in use. *See also* Carrier Sense Multiple Access with Collision Detection (CSMA/CD).

collision domain The extent to which signals in an Ethernet bus topology network are propagated. All devices connected to a logical bus topology network are in the same collision domain. Switch and router ports delimit collision domains.

Committed Information Rate (CIR) A guaranteed minimum transmission rate offered by the service provider.

communication servers Computers that provide a mechanism for users to access a network's resources remotely.

connectionless protocol A type of network communication in which data is transferred without making a connection between communicating devices first, and the receiving station gives no acknowledgement that the data was received.

context switching Occurs when the OS suspends one process and activates another process.

cooperative multitasking In this form of multitasking, the OS can't stop a process; when a process gets control of the CPU, it maintains control until it satisfies its computing needs and informs the OS that another process can be activated.

copy backup A backup that copies selected files to the selected medium without marking files as backed up.

core An instance of a processor inside a single CPU chip. *See also* multicore CPU.

cracker Someone who attempts to compromise a network or computer for personal gain or to cause harm.

credentials A username and password or another form of identity used to access a computer.

crossover cable A type of patch cable that uses the 586B standard on one end and the 586A standard on the other end. This arrangement crosses the transmit and receive wires so that transmit on one end connects to receive on the other end. Often used to connect two devices of the same type to one another—for example, connecting a switch to a switch.

crossstalk Interference one wire generates on another wire when both wires are in a bundle.

customer premises equipment (CPE) The equipment at the customer site that's usually the responsibility of the customer.

customer relationship management (CRM) A category of software designed to help businesses manage customers and sales prospects.

cut-through switching With this switching method, the switch reads only enough of the incoming frame to determine its source and destination addresses. After the forwarding location is determined, the frame is switched internally from the incoming port to the outgoing port, and the switch is free to handle additional frames.

Cyclic Redundancy Check (CRC) The error-checking code in an Ethernet frame's trailer; it's the result of a mathematical algorithm computed on the frame data. When the destination device receives the frame, the calculation is repeated. If the results of this calculation don't match the CRC in the frame, it indicates the data was altered in some way.

daily backup A backup that copies all files changed the day the backup is made; doesn't mark files as backed up.

data circuit-terminating equipment (DCE) The device that sends data to (and receives data from) the last mile; usually a CSU/DSU or modem. *See also* channel service unit/data service unit (CSU/DSU) *and* last mile.

data collector set A feature of Performance Monitor that specifies the performance counters you want to collect, how often to collect them, and the time period.

datagrade A grade of cable suitable for data networking.

Data Link layer Layer 2 in the OSI model is responsible for managing access to the network medium and delivery of data frames from sender to receiver or from sender to an intermediate device, such as a router. *See also* Open Systems Interconnection (OSI) reference model.

Data Over Cable Service Interface Specification (DOCSIS) The official standard governing cable modem operation.

data terminal equipment (DTE) The device that passes data from the customer LAN to the DCE; usually a router. *See also* data circuit-terminating equipment (DCE).

dedicated bandwidth A property of switches in which each port's bandwidth is dedicated to the devices connected to the

port; on a hub, each port's bandwidth is shared between all devices connected to the hub.

deencapsulation The process of stripping the header from a PDU as it makes its way up the communication layers before being passed to the next higher layer. *See also* protocol data unit (PDU).

default gateway The address configured in a computer's IP address settings specifying the address of a router to which the computer can send all packets destined for other networks.

default groups Special groups with rights already assigned; created during installation in a Windows environment.

default route A routing table entry that tells a router where to send a packet with a destination network address that can't be found in the routing table.

demarcation point The point at which the CPE ends and the provider's responsibility begins. *See also* customer premises equipment (CPE).

denial-of-service (DoS) attack An attempt to tie up network bandwidth or services so that network resources are rendered useless to legitimate users.

destination network The network address of a network to which the router can forward packets.

device driver Software that provides the interface between the OS and computer hardware.

Dial-Up Networking (DUN) The Windows software component (beginning with Windows 95) for setting up a connection to an RRAS server or connecting computers to ISPs for dial-up Internet access.

differential backup A backup that copies all files changed since the last full backup; doesn't mark files as backed up.

differential signal A method for transmitting data in which two wires of opposite polarity are used. One wire transmits using positive voltage and the other uses negative voltage. Differential signals enhance signal reliability by providing a canceling affect on EMI and crosstalk.

digital certificates Digital documents used in encryption and authentication protocols that identify a person or computer and can be verified by a certification authority.

digital signal Represented as a square wave, a signal that uses binary 1s and 0s to represent two possible states.

Digital Subscriber Line (DSL) A broadband technology that uses existing phone lines to carry voice and data simultaneously.

directory service The software that manages centralized access and security in a server-based network.

disk mirroring A fault-tolerant disk configuration in which data is written to two hard drives rather than one so that if one disk fails, the data isn't lost.

disk quotas A feature available in some file systems that allows an administrator to set a limit to how much disk space a user's files can occupy.

disk striping with parity A fault-tolerant disk configuration in which parts of several physical disks are linked in an array, and data and parity information are written to all disks in this array. If one disk fails, data can be reconstructed from the parity information written on the others.

distance-vector protocol A routing protocol that routers use to share information about an internetwork's status by copying their routing table to other routers with which they share a network.

DNS zone A database of primarily hostname and IP address pairs that are related by membership in an Internet or a Windows domain.

domain A collection of users and computers in a server-based network whose accounts are managed by Windows servers called domain controllers. *See also* domain controller.

domain controller A computer running Windows Server with Active Directory installed; maintains a database of user and computer accounts as well as network access policies in a Windows domain. *See also* directory service.

Domain Name System (DNS) An Application-layer protocol that resolves computer and domain names to their IP addresses. DNS uses the UDP Transport-layer protocol. *See also* User Datagram Protocol (UDP).

dotted decimal notation The format used to express an IPv4 address; four decimal numbers separated by periods.

dynamic disk A disk configuration in Windows that can be divided into one or more volumes. You can create up to 1000 volumes per dynamic disk (although no more than 32 is recommended). A dynamic disk offers features that a basic disk doesn't, namely RAID and disk spanning.

Dynamic Host Configuration Protocol (DHCP) An Application-layer protocol used to configure a host's IP address settings dynamically. It uses the UDP Transport-layer protocol because DHCP messages consist of a single packet and are used on the local LAN.

electromagnetic interference (EMI) A disturbance to the operation of an electronic circuit or its data, caused by devices that emit an electromagnetic field.

encapsulation The process of adding header and trailer information to chunks of data.

encoding Representing 0s and 1s as a physical signal, such as electrical voltage or a light pulse.

encoding term The method used to represent bits on a medium.

encryption A technology used to make data unusable and unreadable to anyone except authorized users of the data.

end user license agreement (EULA) A license that governs how an application can be used. It specifies how many users are allowed to use an application, how many times it can be installed, and whether the software can be copied, among other things.

entrance facility The location of cabling and equipment that connects a corporate network to a third-party telecommunications provider. It can also serve as an equipment room and the main cross-connect for all backbone cabling.

equipment room A room that houses servers, routers, switches, and other major network equipment and serves as a connection point for backbone cabling running between telecommunications closets.

exclusion A configuration option for the IP address scope; excludes specified IP addresses from the DHCP IP address scope. *See also* IP address scope.

extended LANs A LAN that's expanded beyond its normal distance limitations with wireless communication.

extended partition A partition type that can be divided into one or more logical drives, each of which can be formatted and assigned a drive letter.

extended star topology An extension of the physical star topology, in which a central switch or hub is the central connecting point for other switches or hubs that have computers and other network devices attached, forming a star of stars. *See also* physical star topology.

failover cluster A server cluster configuration used for fault tolerance so that if one server fails, the other takes over its functions immediately, with no or little downtime.

fault tolerance A feature available on some high-end NICs. By installing a second NIC in a PC, failure of the primary NIC shifts network traffic to the second NIC instead of cutting off the PC from the network.

Fiber Distributed Data Interface (FDDI) A technology that uses the token-passing media access method and dual rings for redundancy. The rings in an FDDI network are usually a physical ring of fiber-optic cable. FDDI transmits at 100 Mbps and can include up to 500 nodes over a distance of 100 kilometers.

fiber-optic cable A cable type that carries data over thin strands of glass by using optical (light) pulses to represent bits.

file and print servers Computers that provide secure centralized file storage, sharing, and access to networked printers.

file system The method by which an OS stores, organizes, and manages access to files on a storage device, such as a hard drive.

firewall A hardware device or software program that inspects packets going into or out of a network or computer and then discards or forwards packets based on a set of rules.

flood The process whereby a switch forwards a frame out all connected ports.

flow control A mechanism network protocols use to prevent a destination device from becoming overwhelmed by data from a transmitting computer, resulting in dropped packets.

fractionalized The term used to describe a T-carrier line in which portions are dedicated for different purposes. *See also* T-carrier lines.

fragment-free switching With this switching method, the switch reads enough of the frame to guarantee that it's at least the minimum size for the network type, reducing the possibility that the switch will forward a frame fragment.

frame A packet with source and destination MAC addresses added and an error-checking code added to the back end. Frames are generated by and processed by the network interface. *See also* packet.

frame fragment An invalid frame that's damaged because of a collision or a malfunctioning device.

frame relay A PVC packet-switching technology that offers WAN communication over a fast, reliable, digital link. Throughput is usually improved because error checking is done on endpoint devices instead of on the digital link. *See also* permanent virtual circuits (PVCs).

frame types The frame formats that describe the content and length of a frame header.

frequently asked questions (FAQ) A document with two parts to each entry. The first part is a question the manufacturer has anticipated or received from customers; the second part is an answer to the question.

full backup A backup that copies all selected files to the selected medium and marks files as backed up. Also called a normal backup.

full-duplex mode A communication mode in which a device can simultaneously transmit and receive data on the same cable connection. Switches can operate in full-duplex mode, but hubs can't.

fully qualified domain name (FQDN) A name that includes the hostname, subdomain names (if applicable), second-level domain name, and top-level domain name, separated by periods.

gateway-to-gateway VPN mode This VPN mode establishes a connection between two routers that support VPNs.

Gigabit Ethernet *See* 1000BaseT Ethernet.

hacker Sometimes a derogatory term to describe an unskilled or undisciplined programmer; it can also mean someone who's highly skilled with computer systems and programs and can use some of the same tools attackers use to poke around networks or systems, but not for evil purposes.

half-duplex mode A communication mode in which a device can send or receive data but can't do both simultaneously. Hubs operate only in half-duplex mode; switches can operate in both half-duplex and full-duplex modes.

header Information added to the front end of a chunk of data so that the data can be correctly interpreted and processed by network protocols.

hertz (Hz) A unit expressing how many times per second a signal or electromagnetic wave occurs.

hoax virus A type of virus that's not really a virus but simply an e-mail announcement of a made-up virus. Its harm lies in people believing the announcement and forwarding the e-mail on to others.

homegroup A peer-to-peer networking feature introduced in Windows 7 that simplifies sharing files and printers between computers.

hop Each router a packet must go through to get to the destination network.

hop count The total number of routers a packet must travel through to get to its destination network.

horizontal wiring The network cabling running from the work area's wall jack to the telecommunications closet, usually terminated at a patch panel. The total maximum distance for horizontal wiring is 100 meters.

host computer The physical computer on which virtual machine software is installed and virtual machines run.

hosted applications A category of cloud computing in which a customer pays for the use of applications that run on a service provider's network; also called software as a service (SaaS).

hosted infrastructure A category of cloud computing in which a company can use a provider's storage or virtual servers as its needs demand; also called infrastructure as a service (IaaS).

hosted platforms A category of cloud computing in which a customer develops applications with the service provider's development tools and infrastructure; also called platform as a service (PaaS). After applications are developed, they can be delivered to the customer's users from the provider's servers.

hosted virtualization The hypervisor implements OS virtualization by being installed in a general-purpose host OS, such as Windows 7 or Linux, and the host OS accesses host hardware on behalf of the guest OS. *See also* virtualization.

hot-swappable device A computer device that can be removed, replaced, or added to a server while it's running.

hub A network device that performs the same function as a repeater but has several ports to connect a number of devices; sometimes called a multiport repeater. *See also* repeater.

hypervisor The component of virtualization software that creates and monitors the virtual hardware environment, which allows multiple VMs to share physical hardware resources.

incremental backup A backup that copies all files changed since the last full or incremental backup and marks files as backed up.

infrared (IR) A very long wavelength light source in the invisible spectrum that can be used to transmit data wirelessly.

infrastructure as a service (IaaS) *See* hosted infrastructure.

infrastructure mode An operational mode for Wi-Fi networks, in which wireless stations connect through a wireless access point before they can begin communicating with other devices.

Integrated Services Digital Network (ISDN) A digital WAN technology developed to replace the analog phone system. It defines communication channels of 64 Kbps and is most often used by OSHO users.

International Organization for Standardization (ISO) The international standards-setting body based in Geneva, Switzerland, that sets worldwide technology standards.

Internet Control Message Protocol (ICMP) An Internet-layer protocol used to send error and control messages between systems or devices. It's an encapsulated IP protocol, meaning it's wrapped in an IP header.

Internet Message Access Protocol (IMAP) An Application-layer protocol used by an e-mail client to download messages from an e-mail server; operates on TCP port 143. IMAP also provides fault-tolerance features. It downloads only message headers from the server initially, and then downloads the message body and attachments after the message is selected.

Internet Protocol Security (IPSec) An extension to IP that provides security by using authentication and encryption. It authenticates the identity of computers transmitting data with a password or some other form of credentials, and it encrypts data so that if packets are captured, the data will be unintelligible.

Internet Protocol Version 4 (IPv4) A connectionless Internet-layer protocol that provides source and

destination addressing and routing for the TCP/IP protocol suite. Uses 32-bit dotted decimal addresses.

Internet Protocol version 6 (IPv6) A connectionless Internetwork-layer protocol that provides source and destination addressing and routing for the TCP/IP protocol suite. Uses 128-bit hexadecimal addresses and has built-in security and QoS features.

internetwork A networked collection of LANs tied together by devices such as routers. *See also* local area network (LAN).

intrusion detection system (IDS) Usually a component of a firewall, a hardware device or software that detects an attempted security breach and notifies the network administrator. An IDS can also take countermeasures to stop an attack in progress.

IP address A 32-bit dotted-decimal address used by IP to determine the network a host resides on and to identify hosts on the network at the Internetwork layer.

IP address scope A component of a DHCP server, it's a range of IP addresses the server leases to clients requesting an IP address.

IP prefix A value used to express how many bits of an IP address are network ID bits. Usually expressed as */Prefix Number*—for example, 192.168.1.24/27, with 27 as the IP prefix.

IP Security (IPSec) An extension to the IP protocol suite that creates an encrypted and secure conversation between two hosts.

IrDA devices Devices that use infrared signals to communicate. IrDA stands for Infrared Device Association.

Kerberos authentication An authentication protocol used in a Windows domain environment or on a Linux system; uses OS-generated keys, which makes this protocol more secure than having an administrator enter keys.

last mile The connection between a WAN's demarcation point and the central office (CO); also called the local loop. *See also* demarcation point.

link-state protocol A routing protocol that a router uses to share information with other routers by sending the status of all its interface links to all other routers in the internetwork. The status includes link speed, whether the link is up or down, and the link's network number.

load-balancing cluster A server cluster configuration that provides high-performance computing and data access by spreading the workload among multiple computers.

local area network (LAN) A small network, limited to a single collection of machines and linked by interconnecting devices in a small geographic area.

local loop *See* last mile.

local profile A user profile stored on the same system where a user logs on; created from a hidden profile called Default the first time a user logs on to the system. *See also* user profile.

localhost The name used to refer to the loopback address in an IP network. *See also* loopback address.

Logical Link Control (LLC) sublayer The upper sublayer of the IEEE Project 802 model for the OSI model's Data Link layer. It handles error-free delivery and controls the flow of frames between sender and receiver across a network.

logical topology The path data travels between computers on a network. The most common logical topologies are switched, bus, and ring.

loopback address An address that always refers to the local computer; in IPv4, 127.0.0.1 is the loopback address.

MAC address filtering A security method often used in wireless networks, in which only devices with MAC addresses specified by the administrator can gain access to the wireless network.

mail servers Computers that handle sending and receiving e-mail messages for network users.

malware Any software designed to cause harm or disruption to a computer system or otherwise perform activities on a computer without the consent of the computer's owner.

managed switch A high-end switch with many advanced features that can be configured.

management information base (MIB) A collection of network data stored by Simple Network Management Protocol software agents. *See also* software agents.

maximum transmission unit (MTU) The maximum frame size allowed to be transmitted across a network medium.

MDI crossed (MDI-X) devices Network devices that connect by using RJ-45 plugs over twisted-pair cabling; they transmit over pins 3 and 6 and receive over pins 1 and 2 of an RJ-45 connector.

media access control *See* media access method.

Media Access Control (MAC) sublayer The lower sublayer of the IEEE Project 802 model for the OSI model's Data Link layer. It handles accessing network media and mapping between logical and physical network addresses for NICs.

media access method A set of rules governing how and when the network medium can be accessed for transmission. The rules ensure that data is transmitted and received in an orderly fashion, and all stations have an opportunity to communicate. Also called media access control.

medium dependent interface (MDI) devices Network devices that connect by using RJ-45 plugs over twisted-pair cabling; they transmit on pins 1 and 2 and receive on pins 3 and 6 of an RJ-45 connector.

mesh topology A topology in which each device in the network is connected to every other device, providing multiple pathways in the event of a device or cable failure.

metric A numeric value that tells the router how “far away” the destination network is. It can be composed of values such as the bandwidth of links between the source and destination, the hop count, and the link’s reliability.

metropolitan area network (MAN) An internetwork confined to a geographic region, such as a city or county; uses third-party communication providers to provide connectivity between locations. *See also* internetwork.

modem A device that converts a sending computer’s digital signals to analog signals for transmission over phone lines and then converts analog signals to digital signals for the receiving computer.

multicore CPU A CPU containing two or more processing cores. *See also* core.

multiplexing A technology that supports simultaneous communication links over the same set of cables, so data transmissions from several sources can be combined and delivered over a single cable.

multiprocessing A feature of some OSs that allow two or more threads to be run concurrently by separate CPUs or CPU cores. *See also* thread.

Multiprotocol Label Switching (MPLS) A highly scalable, flexible WAN technology that works with any Network-layer protocol and is independent of the Data Link layer technology; used exclusively in IP networks. It creates a connection-oriented virtual circuit, using labels assigned to each packet that make it unnecessary to view packet contents.

multitasking An operating system’s capability to run more than one application or process at the same time.

multithreaded application An application that has two or more threads that can be scheduled separately for execution by the CPU. *See also* thread.

name server A computer that stores names and addresses of computers on a network, allowing other computers to use computer names rather than addresses to communicate with one another.

narrowband radio Low-powered, two-way radio communication systems, such as those used in taxis, police radios, and other private radio systems. Also called “single-frequency radio.”

neighbor In an internetwork, routers sharing a common network.

network Two or more computers connected by a transmission medium that enables them to communicate.

Network Address Translation (NAT) A service that translates a private IP address to a public IP address in packets destined for the Internet, and then translates the public IP address in the reply to the private address.

network appliance A device equipped with specialized software that performs a limited task, such as file sharing. Network appliances are often packaged without video interfaces, so you don’t configure them with an attached keyboard and monitor.

network backbone The cabling used to communicate between LANs or between hubs or switches. The backbone cabling often runs at a faster speed than the cabling used to connect computers because the backbone must carry data from many computers to other parts of the network.

network bandwidth The amount of data that can be transferred on a network during a specific interval; usually measured in bits per second.

network client software The application or OS service that can request information stored on another computer.

Network File System (NFS) The native Linux file-sharing protocol.

Network Information Service (NIS) A Linux directory service that supports centralized logon.

network interface card (NIC) A device that creates and mediates the connection between a computer and the network medium.

Network layer Layer 3 of the OSI model handles logical addressing and routing of PDUs across internetworks. *See also* Open Systems Interconnection (OSI) reference model *and* protocol data unit (PDU).

network model A model defining how and where resources are shared and how access to these resources is regulated.

network monitors Programs that monitor network traffic and gather information about packet types, errors, and packet traffic to and from each computer.

network protocols The software defining the rules and formats a computer must use when sending information across the network.

network server software The software that allows a computer to share its resources by fielding requests generated by network clients.

network-attached storage (NAS) A dedicated server device designed solely for providing shared storage for network users.

next hop An interface name or the address of the next router in the path to the destination network.

NTFS permissions A feature in Windows NTFS that gives administrators fine-grained control over file and folder access for both network users and interactive users.

octet A grouping of 8 bits, often used to identify the four 8-bit decimal numbers that compose an IP address (as in “first octet,” “second octet,” and so forth).

onboard co-processors A feature included on most NICs that enables the card to process incoming and outgoing network data without requiring service from the CPU.

Open Systems Interconnection (OSI) reference model ISO Standard 7498 defines a frame of reference for understanding networks by dividing the process of network communication into seven layers. Each layer is defined in terms of the services and data it handles on behalf of the layer above it and the services and data it needs from the layer below it.

packet A chunk of data with source and destination IP addresses (as well as other IP information) added to it. Packets are generated by and processed by network protocols.

packet filtering A process whereby a router blocks a packet from being forwarded based on rules specified by an access control list. *See also* access control list (ACL).

packet forwarding The process of a router receiving a packet on one port and forwarding it out another port based on the packet’s destination network address and information in the routing table.

packet-switched WAN A type of WAN network in which data is transmitted in frames or packets, and each packet is transmitted through the provider’s network independently. Instead of having a dedicated circuit over which data travels, a provider’s customers share the bandwidth.

patch cable A short cable for connecting a computer to an RJ-45 jack or connecting a patch-panel port to a switch or hub. *See also* straight-through cable.

PCI Express (PCIe) A bus standard that uses a high-speed serial communication protocol of one or more lines or lanes. Each lane of PCIe 1.0 can operate at 250 MBps in each direction. *See also* Peripheral Component Interconnect (PCI).

PCMCIA cards Credit card-sized expansion cards used mainly to add functionality to laptop computers. The main standards are Cardbus and ExpressCard. Cardbus operates at 33 MHz and supports a 32-bit bus; ExpressCard uses PCIe technology to provide data transfer speeds up to 500 MBps.

peer communication In the layered approach, each layer on one computer behaves as though it were communicating with its counterpart on the other computer. This means each layer on the receiving computer sees network data in the same format its counterpart on the sending computer did.

peer-to-peer network A network model in which all computers can function as clients or servers as needed, and there’s no centralized control over network resources.

penetration tester A term used to describe a security consultant who detects holes in a system’s security for the purpose of correcting these vulnerabilities.

Peripheral Component Interconnect (PCI) A bus standard used to connect I/O devices to the memory and CPU of a PC motherboard. PCI is implemented in both 32-bit and 64-bit versions at speeds of 33 and 66 MHz, respectively.

Peripheral Component Interconnect-Extended (PCI-X) A bus standard that’s backward-compatible with PCI and supports speeds of 66 to 533 MHz with 32-bit or 64-bit bus widths. *See also* Peripheral Component Interconnect (PCI).

permanent virtual circuits (PVCs) Pathways between two communication points that are established as permanent logical connections; therefore, the pathway exists even when it’s not in use. *See also* virtual circuit.

physical bus topology A network topology in which a continuous length of cable connects one computer to another in daisy-chain fashion. There’s no central interconnecting device.

Physical layer Layer 1, the bottom layer of the OSI model, transmits and receives signals and specifies the physical details of cables, NICs, connectors, and hardware behavior. *See also* Open Systems Interconnection (OSI) reference model.

physical ring topology A cabling arrangement in which each device is connected to another device in daisy-chain fashion, and the last device connects back to the first device forming a ring. Used by token ring and FDDI, the physical ring is rarely used now.

physical star topology A network topology that uses a central device, such as a hub or switch, to interconnect computers in a LAN. Each computer has a single length of cable going from its NIC to the central device. It’s the most common physical topology in LANs.

physical topology The arrangement of cabling and how cables connect one device to another in a network. The most common physical topology is a star, but bus, ring, point-to-point, and mesh topologies are also used.

ping scanner An automated method for pinging a range of IP addresses.

platform as a service (PaaS) *See* hosted platforms.

Point-to-Point Protocol (PPP) A remote access protocol that supports many protocols and is used to carry data over a variety of network connections.

point-to-point topology A topology in which cabling creates a direct link between two devices; used most often in WANs or in wireless networks to create a wireless bridge.

Port Address Translation (PAT) An extension of NAT, a service that allows several hundred workstations to access the Internet with a single public Internet address by using

Transport-layer port numbers to differentiate each host conversation. *See also* Network Address Translation (NAT).

port forwarding The process by which a router forwards a request for a TCP or UDP port to a specified computer.

port number A field in the Transport-layer protocol header that specifies the source and destination Application-layer protocols that are used to request data and are the target of the request, respectively.

port scanner Software that determines which TCP and UDP ports are available on a computer or device.

Post Office Protocol version 3 (POP3) An Application-layer protocol used by a client e-mail application to download messages from an e-mail server; operates on TCP port 110.

power conditioning A method of cleaning the power input, removing noise caused by other devices on the same circuit.

preemptive multitasking A form of multitasking in which the OS controls which process gets access to the CPU and for how long.

Presentation layer At Layer 6 of the OSI model, data can be encrypted and/or compressed to facilitate delivery. Platform-specific application formats are translated into generic data formats for transmission or from generic data formats into platform-specific application formats for delivery to the Application layer. *See also* Open Systems Interconnection (OSI) reference model.

preshared key A series of letters, numbers, and special characters, much like a password, that both communicating devices use to authenticate each other's identity.

primary partition A partition type that can be formatted with a file system and assigned a drive letter or mounted in an empty folder on an existing drive letter; also called a volume. *See also* volume.

Primary Rate Interface (PRI) An ISDN format that consists of 23 64-Kbps B-channels and one 64-Kbps D-channel. *See also* Integrated Services Digital Network (ISDN).

process A program that's loaded into memory and run by the CPU. It can be an application a user interacts with or a program with no user interface that communicates with and provides services to other processes.

promiscuous mode An operational mode of a NIC in which all frames are read and processed rather than only broadcast and unicast frames addressed to the NIC. Protocol analyzer software sets a NIC to promiscuous mode so that all network frames can be read and analyzed.

protocol Rules and procedures for communication and behavior. Computers must use a common protocol and agree on the rules of communication.

protocol analyzers Programs or devices that can capture packets traversing a network and display packet contents in a form useful to the user.

protocol data unit (PDU) A unit of information passed as a self-contained data structure from one layer to another on its way up or down the network protocol stack.

protocol stack *See* protocol suite.

protocol suite A set of protocols working cooperatively to provide network communication. Protocols are "stacked" in layers in which each layer performs a unique function required for successful communication. Also called a protocol stack.

quality of service (QoS) A term that describes a network's capability to prioritize data packets based on the type of information they contain (for example, voice, video, or file data) or urgency of the information.

radio frequency interference (RFI) Similar to EMI, but RFI is usually interference caused by strong broadcast sources. *See also* electromagnetic interference (EMI).

RAM buffering A NIC feature for including additional memory to provide temporary storage for incoming and outgoing data.

redirector An OS client component that intercepts resource requests and determines whether the resource is local or remote.

redundant array of independent disks (RAID) A storage configuration of two or more disks, usually in a fault-tolerant arrangement so that if one disk fails, data is preserved and the server can continue to operate.

redundant power supply A second power supply unit in the computer case. Each unit is capable on its own of maintaining adequate power to the computer, so if one power supply fails, the other unit takes on the full load.

reflection *See* signal bounce.

Remote Monitoring (RMON) An advanced network-monitoring protocol that extends Simple Network Management Protocol's capabilities; contains software agents called probes that collect data and communicate with a management station by using SNMP.

repeater A network device that takes incoming signals and regenerates, or repeats them to other parts of the network.

request to send (RTS) A signal used in wireless networks indicating that a computer has data ready to send on the network. *See also* access point *and* clear to send (CTS).

reservation A configuration option for an IP address scope that ties an IP address to a MAC address. When a client

requests an IP address from the DHCP server, if the client's MAC address matches an address specified by a reservation, the reserved IP address is leased to the client instead of getting it from the scope. *See also* IP address scope.

rights In Windows, they define the types of actions a user can perform, such as creating file shares or installing software.

RJ-45 jack A device used in the work area in wall plates and surface-mounted boxes to plug a patch cable that connects a computer to the horizontal wiring.

RJ-45 plug A connector used to terminate twisted-pair cable for making patch cables. It has eight wire traces to accommodate a standard twisted-pair cable with four wire pairs.

roaming profile A user profile in a Windows environment that's stored on a server and can be accessed from any computer the user logs on to. *See also* user profile.

rollback plan The part of an upgrade plan with instructions on how to undo the upgrade if problems occur during or after the upgrade.

rootkits Forms of Trojan programs that can monitor traffic to and from a computer, monitor keystrokes, and capture passwords. They're among the most insidious form of malware because they can mask that the system has been compromised by altering system files and drivers required for normal computer operation. *See also* malware.

router A device that enables multiple LANs to communicate with one another by forwarding packets from one LAN to another. Routers also forward packets from one router to another when LANs are separated by multiple routers; they have multiple interfaces, and each interface communicates with a LAN.

Routing and Remote Access Service (RRAS) A software component included in Windows Server 2008 that provides remote access through dial-up and VPN connections as well as routing and packet filtering.

Routing Information Protocol (RIP) A distance-vector protocol that uses hop count as the metric to determine the best path to a destination network.

Routing Information Protocol version 2 (RIPv2) A newer version of RIP that supports a more complex IP addressing scheme and uses multicast packets rather than broadcasts to transmit routing table updates. *See also* Routing Information Protocol (RIP).

routing protocol A set of rules routers use to exchange information so that all routers have accurate information about an internetwork to populate their routing tables.

satellite microwave Microwave communication systems that send and receive data from satellites that maintain fixed positions in the sky.

segment The unit of information used by the Transport layer. A segment is passed up to the Application layer as data and passed down to the Internetwork layer, where it becomes a packet.

server Term used to describe an OS designed mainly to share network resources, a computer with the primary role of giving client computers access to network resources, and the software that responds to requests for network resources from client computers.

server cluster Two or more servers configured to operate as a single unit. The most common types of server clusters are failover clusters and load-balancing clusters.

Server Message Block (SMB) The Windows file-sharing protocol.

server-based network A network model in which servers take on specialized roles to provide client computers with network services and to provide centralized control over network resources.

service A process that runs in the background and provides services to other processes; for example, DNS client and server components are services.

service set identifier (SSID) The name assigned to a wireless network so that wireless clients can distinguish between them when more than one is detected.

Session layer Layer 5 of the OSI model is responsible for setting up, maintaining, and ending communication sequences (called sessions) across a network. *See also* Open Systems Interconnection (OSI) reference model.

shadow passwords A secure method of storing user passwords on a Linux system.

shared adapter memory A feature on some NICs in which the NICs buffers map directly to RAM on the computer. A computer actually writes to buffers on the NIC instead of writing to its own memory.

shared system memory A feature on some NICs in which a NICs onboard processor selects a region of RAM on the computer and writes to it as though it were buffer space on the adapter.

signal bounce The result of electricity bouncing off the end of a cable and back in the other direction. It causes corruption of data as the bouncing signal collides with signals behind it. A terminator at each cable end is needed to prevent signal bounce. Also called reflection.

signal propagation Signals traveling across a medium and through any connectors and connecting devices until the signal weakens enough to be undetectable or is absorbed by a termination device.

Simple Mail Transfer Protocol (SMTP) The standard protocol for sending e-mail over the Internet.

smart switch A mid-range switch with some advanced features, typically multicast processing, Spanning Tree Protocol, VLANs, and port security. *See also* Spanning Tree Protocol (STP) *and* virtual local area networks (VLANs).

snapshot A partial copy of a virtual machine made at a particular moment, used to restore the virtual machine to its state when the snapshot was taken. *See also* virtual machine (VM).

social engineering A tactic attackers use to get users to perform an action, such as opening an infected e-mail attachment, sending a hoax virus, or providing a password, without being aware that they're aiding the attacker. *See also* hoax virus.

software agents Simple Network Management Protocol components that are loaded on network devices; they monitor network traffic and device status information and send it to a management station.

software as a service (SaaS) *See* hosted applications.

spam Unsolicited e-mail. The harm in spam is the loss of productivity when people receive dozens or hundreds of spam messages daily and the use of resources to receive and store spam on e-mail servers.

Spanning Tree Protocol (STP) A communication protocol switches use to ensure that they aren't connected in a way that creates a switching loop. *See also* switching loop.

special identity groups A type of group in Windows in which membership is controlled dynamically by Windows, can't be viewed or changed manually, and depends on how an account accesses the OS. For example, membership in the Authenticated Users group is assigned to a user account automatically when the user logs on to a computer or domain.

spoofed address A source address inserted into a packet that's not the sender's actual address.

spread-spectrum radio A radio communication system that uses multiple frequencies simultaneously, thereby improving reliability and reducing susceptibility to interference over narrowband radio.

spyware A type of malware that monitors or controls part of your computer at the expense of your privacy and the gain of some third party. *See also* malware.

stand-alone computer A computer that doesn't have the necessary hardware or software to communicate on a network.

stateful packet inspection (SPI) A filtering method used in a firewall, whereby packets aren't simply filtered based on packet properties but are checked for the context in which

they're being transmitted. If a packet isn't part of a legitimate, ongoing data conversation, it's denied.

static route A routing table entry that's entered manually by an administrator.

storage area network (SAN) A specialized network that uses high-speed networking technologies to give servers fast access to large amounts of disk storage.

store-and-forward switching This switching method requires the switch to read the entire frame into its buffers before forwarding it. It examines the frame check sequence (FCS) field to be sure the frame contains no errors before it's forwarded.

straight-through cable A standard patch cable that uses the same wiring standards on both ends so that each wire is in the same location on both ends of the cable (pin 1 goes to pin 1, pin 2 to pin 2, and so forth). *See also* patch cable.

structured cabling A specification for organizing cabling in data and voice networks, regardless of the media type or network architecture.

subnet A subdivision of an IP network address space.

subnet mask A 32-bit number in dotted decimal format, consisting of a string of eight or more binary 1s followed by a string of 0s, that determines which part of an IP address is the network ID and which part is the host ID. A binary 1 in the subnet mask signifies that the corresponding bit in the IP address belongs to the network ID, and a binary 0 signifies that the corresponding bit in the IP address belongs to the host ID.

subnetting The process of dividing an IP network address into two or more subnetwork addresses. *See also* subnet.

supernetting Reallocating bits from the network portion of an IP address to the host portion, effectively combining two or more smaller subnets into a larger supernet.

surge protection Power protection that evens out spikes or sags in the main current and prevents them from affecting a computer.

switch A network device that reads the destination MAC addresses of incoming frames to determine which ports should forward the frames.

switched virtual circuits (SVCs) A communication circuit that's established when needed and then terminated when the transmission is completed. *See also* virtual circuit.

switching loop A condition that occurs when switches are connected in such a way that frames can be forwarded endlessly from switch to switch in an infinite loop.

switching table A table containing MAC address and port pairs that a switch uses to determine which port to forward frames it receives.

Symmetric DSL (SDSL) A DSL variation in which the download and upload speeds are equivalent, or symmetrical. *See also* Digital Subscriber Line (DSL).

Synchronous Optical Network (SONET) A flexible, highly fault-tolerant technology that can carry signals of different capacities over a fiber-optic network at high speeds. It defines optical carrier (OC) levels for incrementally increasing data rates, and SONET networks can be arranged in a variety of physical topologies.

system partition The active primary partition storing the Windows boot loader.

T-carrier lines Communication lines that use one pair of wires for transmitting data and another pair for receiving data. They use the TDM signaling method, making it possible to extract any number of channels for a particular purpose. *See also* time division multiplexing (TDM).

telecommunications closet (TC) Usually an enclosed space or room that provides connectivity to computer equipment in the nearby work area; can also serve as the entrance facility in small installations. Typical equipment includes patch panels to terminate horizontal wiring runs, hubs, and switches.

termination The attachment of RJ-45 plugs on a cable to make a patch cable or punching down the cable wires into terminal blocks on a jack or patch panel.

terminator An electrical component called a resistor, placed at the ends of a physical bus network to absorb the signal instead of allowing it to bounce back up the wire.

terrestrial microwave Line-of-sight transmissions between microwave towers or between transmitters and receivers mounted on tall buildings, mountaintops, or other locations with long, clear lines of sight.

thread The smallest unit of software that can be scheduled to run.

three-way handshake A series of three packets used between a client and server to create a TCP connection. After the three-way handshake has been completed successfully, a connection is established between client and server applications, and data can be transferred.

time division multiplexing (TDM) A signaling method that allocates a time slot for each channel, making it possible to transmit multiple streams, or channels, of data on a single physical medium.

time-domain reflectometer (TDR) A network troubleshooting device that can determine whether there's a break or short in the cable and, if so, approximately how far down the cable it's located. Also shows the total cable length.

time slicing The process by which a CPU's computing cycles are divided between more than one process.

token ring A technology based on the IEEE 802.5 standard; its cabling is in a physical star topology, but it functions as a logical ring. It uses the token-passing media access method, and only the computer holding the token can send data.

total cost of ownership (TCO) The cost of a product or service when intangibles, such as support costs and productivity gains or losses, are factored in.

trailer Information added to the back end of a chunk of data so that the data can be correctly interpreted and processed by network protocols.

transceiver A device that transmits and receives. In wireless networking, an access point is a transceiver.

Transmission Control Protocol (TCP) A connection-oriented Transport-layer protocol designed for reliable transfer of information in complex internetworks.

Transmission Control Protocol/Internet Protocol (TCP/IP) The most common protocol suite in use, TCP/IP is the default protocol in contemporary OSs and the protocol of the Internet.

Transport layer Layer 4 of the OSI model is responsible for reliable delivery of data streams across a network. Layer 4 protocols break large streams of data into smaller chunks and use sequence numbers and acknowledgements to provide communication and flow control. *See also* Open Systems Interconnection (OSI) reference model *and* protocol data unit (PDU).

Trojan program A program that appears to be useful, such as a free utility, but in reality contains some type of malware.

trunk port A switch port configured to carry traffic from all VLANs to another switch or router. *See also* virtual local area networks (VLANs).

twisted-pair (TP) cable A cable containing one or more pairs of insulated strands of copper wire twisted around one another and housed in an outer sheath.

unicast frame A network message addressed to only one computer on the LAN.

uninterruptible power supply (UPS) A power protection device that includes a battery backup to take over if the main current fails; usually incorporates power conditioning and surge protection.

Universal Serial Bus (USB) An external PC bus interface for connecting I/O devices. Speeds range from 12 Mbps in USB 1.0 to 3.2 Gbps in USB 3.0.

uplink port A designated port on a hub or switch used to connect to another hub or switch without using a crossover cable.

User Datagram Protocol (UDP) A connectionless Transport-layer protocol designed for efficient communication of generally small amounts of data.

user profile A collection of a user's personal files and settings that define his or her working environment.

virtual circuit A logical connection created between two devices in a shared network, with bandwidth allocated for a specific transmission pathway through the network.

virtual local area networks (VLANs) A feature on some switches that allows configuring one or more switch ports into separate broadcast domains.

virtual machine (VM) A software environment that emulates a physical computer's hardware and BIOS.

virtual private networks (VPNs) Temporary or permanent connections across a public network that use encryption technology to transmit and receive data. *See also* encryption.

virtualization A process that creates a software environment to emulate a computer's hardware and BIOS, allowing multiple OSs to run on the same physical computer at the same time.

virus A malicious program that spreads by replicating itself into other programs or documents; usually aims to disrupt computer or network functions by deleting and corrupting files.

voicegrade A grade of cable that's not suitable for data networking but is suitable for voice communication.

volume Part or all of the space on one or more disks that contains or is ready to contain a file system. In Windows, volumes with file systems are usually assigned a drive letter. In Linux, volumes are mounted in the file system and accessed as though they were just another folder.

wardrivers Attackers who drive around with a laptop or PDA looking for wireless LANs to access.

Web server A computer running software that allows users to access HTML and other document types with a Web browser.

wide area networks (WANs) Internetworks that are geographically dispersed and use third-party communication providers to provide connectivity between locations. *See also* internetwork.

Wi-Fi Protected Access (WPA) A wireless security protocol that's the successor to Wired Equivalency Protocol. It has enhancements that make cracking the encryption code more difficult.

Wired Equivalent Privacy (WEP) A form of wireless security that encrypts data so that unauthorized people receiving wireless network signals can't interpret the data easily.

wireless bridge An operational mode of wireless networking usually used to connect two wired LANs that are separated from each other in such a way that using physical media is impractical. Can also be used to extend the reach of a wireless network.

Wireless Fidelity (Wi-Fi) The name given to the 802.11 series of IEEE standards that define four common varieties of wireless LANs: 802.11a, 802.11b, 802.11g, and 802.11n.

wireless personal area network (WPAN) A short-range networking technology designed to connect personal devices to exchange information.

work area The location of workstations and other user devices—in short, the place where people work with computers and other network devices.

Worldwide Interoperability for Microwave Access

(WiMAX) A wireless broadband technology defined in 802.16-2004 for fixed WiMAX and 802.16e for mobile WiMAX. WiMAX is considered a fourth-generation (4G) technology for bringing wireless Internet access to remote areas, large areas up to a mile radius, and mobile users.

worm A self-replicating program, similar to a virus, that uses network services such as e-mail to spread to other systems. *See also* virus.

X.25 A packet-switching technology that provides an interface between public packet-switching networks and their customers; it has the advantage of running effectively over older copper phone lines. X.25 networks are SVC networks, meaning they create the best available pathway at the time of transmission. *See also* switched virtual circuits (SVCs).

Note: page numbers in **boldface** type indicate key terms

A

- acceptable use policies, 461
 - access
 - See also* data access security; remote access
 - policy, 461
 - access control lists (ACLs), 308, 478
 - access control(s), 272
 - access points. *See* wireless access points (APs)
 - accounting software for small businesses, 518–519
 - account management, 358–359
 - See also* user accounts
 - account names, 396–397
 - ACLs (access control lists), 308, 478
 - ACPI (Advanced Configuration Power Management Interface), 318
 - acronyms, 49–50
 - Active Directory, 44, 358–359
 - group account in, 404
 - user account in, 398–399, 401
 - active partitions, 415
 - addressing IP. *See* Internet Protocol (IP) addressing
 - Address Resolution Protocol (ARP), 216–218
 - cache, 217
 - capturing packets, 219–220
 - command, 24, 25, 221
 - address space, 232
 - ad hoc topology, Wi-Fi configuration, 141–144
 - ADSL (Asymmetric Digital Subscriber Line), 148
 - American National Standard Institute (ANSI), 621
 - UTP cabling categories, 168–169
 - American Wire Gauge (AWG) standards, 638, 639
 - analog signal, 535, 537
 - ANSI (American National Standard Institute), 621
 - UTP cabling categories, 168–169
 - AP (access point), 77–79
 - See also* wireless access points
 - Apache Web server, 46
 - APIPA (automatic private IP addressing), 227
 - AppleTalk, 636, 644
 - environment, 639–640
 - Application layer, 209
 - role of, 225–226
 - Application layer, OSI reference model protocols, 268–269
 - applications
 - bare-metal virtualization, 376–378
 - hosted, 520, 556
 - hosted virtualization, 368–369
 - sharing in small-business networks, 505
 - source and destination, 223
 - application servers, 45
 - ARP (Address Resolution Protocol), 216–218
 - cache, 217
 - capturing packets, 219–220
 - command, 24, 25, 221
 - Asymmetric Digital Subscriber Line (ADSL), 148
 - Asynchronous Transfer Mode (ATM), 544–545
 - ATM (Asynchronous Transfer Mode), 544–545
 - attenuation, 165
 - auditing policy, 461, 462
 - authentication, 358, 466
 - IP Sec and, 472–473
 - logon hour and location restriction, 470–471
 - passwords, 396, 466–471
 - policies, 461, 466–467
 - authorization, 358
 - file and folder access, 472
 - implementing secure, 466
 - logon hour and location restriction, 470–471
 - automatic link aggregation, 318
 - automatic private IP addressing (APIPA), 227
 - auto-MDIX (medium dependent interface crossed), 291
 - auto-negotiate mode, 291
- ## B
- backbone cabling, 175, 176
 - backdoor, 483
 - backups, 442–446
 - disaster recovery and, 604–606
 - fault tolerance and, 442–443, 446
 - small-business networks, 523–524
 - types, 605–606
 - in Windows, 443–446

- bandwidth
 - cable selection criterion, 164–165, 196
 - hubs and, 64–65
 - switches and, 68
 - wireless APs and, 78–79
 - bandwidth sharing, 65
 - bare-metal virtualization, 367, 376–380
 - baseband, 122
 - baseline, 437–438
 - Baseline Security Services, 626
 - basic disk, 414
 - basic input/output system (BIOS), 8–9
 - Basic Rate Interface (BRI), 538–539
 - batch file, 354
 - beaconing, 637–638
 - binary arithmetic, 239–242
 - binary to decimal conversion, 241–242
 - bit patterns, 242
 - decimal to binary conversion, 240–241
 - subnet mask calculation, 243–246
 - BIOS (basic input/output system), 8–9
 - Setup utility, 9–11
 - BitLocker, 474
 - bits (binary digits), 5
 - patterns, 242, 244
 - blocking mode, 295
 - Bluetooth, 48
 - boot partition, 415
 - boot procedure, 9–11
 - BRI (Basic Rate Interface), 538–539
 - broadband, 122
 - cable modem technology, 145–147
 - DSL, 148
 - broadband optical telepoint networks, 191
 - broadcast domains, 92
 - broadcast frames, 295
 - broadcast storm, 295
 - budget as cable selection criterion, 196
 - bus(es), 7
 - PC, 314–317
 - bus mastering, 317
- C**
- cable, 166–188
 - See also* unshielded twisted-pair (UTP) cable
 - bandwidth rating, 164–165
 - baseband transmission, 122
 - basic and advanced testers, 601
 - broadband transmission, 122
 - coaxial, 122, 167, 630–631
 - crossover *vs.* straight-through, 179
 - fiber-optic, 122, 184–188
 - grade, 166
 - installation costs, 167
 - patch, 171, 179–181
 - selection criteria, 164–167
 - small-business networks, 513
 - structured, 172–176
 - termination, 177, 181–182
 - thinwire Ethernet (RG), 629–630
 - token ring environment, 638–639
 - troubleshooting, 602–603, 646
 - twisted-pair, 167–172
 - cable modems, 145–147
 - operation of, 147
 - cable plant, 166, 569
 - managing and installing, 172–176
 - upgrading, 197
 - cable segment, 165
 - cable testers, 601
 - CA (certification authority), 472–473
 - cache, in DNS query, 363–364
 - Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), 140–141
 - Carrier Sense Multiple Access with Collision Detection (CSMA/CD), 124–125
 - CATV cables, 145, 146
 - CCITT (Comité Consultatif International Téléphonique et Télégraphique), 622–623
 - CDE (Common Desktop Environment), 626
 - CEH (Certified Ethical Hacker) certification, 487
 - centralized computing, 358
 - centralized storage, 360–361
 - central office (CO), 548
 - central processing unit (CPU)
 - architecture, 381
 - monitoring utilization, 343, 344
 - multicore, 3–4
 - upgrading, 512
 - certification authority (CA), 472–473
 - Certified Ethical Hacker (CEH), 487
 - Channel Service Unit/Data Service Unit (CSU/DSU), 536, 540
 - checksum, 216, 224
 - CIDR (Classless Inter-Domain Routing), 236
 - CIFS (Common Internet File System), 354
 - CIR (Committed Information Rate), 543
 - circuit-switched WANs, 537
 - Citrix terminal server products, 378

- Citrix XenServer, 378, 380
- Classless Interdomain Routing (CIDR), 232, 236
- clear to send (CTS), 78
- client network software, 16
- client operating systems, 33, 361, 346–357
 - creating users in, 401–403
 - DHCP client, 346–347, 350–351
 - DNS client, 347–352
 - e-mail client, 356–357
 - file-sharing client, 352–354
 - groups in Windows, 406–407
 - HTTP client, 352
 - mapping drive letter, 354–355
 - shared printer, 355–356
- Client OS (Windows), 401–403, 406–407
- client(s), 32–34
- client/server computing, 32–34
 - applications, 365
 - problems with, 647
- client-to-gateway VPN mode, 517
- cloud computing, 556–558
- cloud storage, 361
- clustering, server, 365, 449–450
- CMOS (complementary metal oxide semiconductor), 9
- CMP (communication cable plenum), 166
- CMR (communication cable riser), 166
- coaxial cable, 122, 167
 - thickwire Ethernet, 630–631
- collision domains, 125–126
- collision(s), 65, 124–125
- Comité Consultatif International Téléphonique et Télégraphique (CCITT), 622–623
- command line, Windows, 334–337
- Committed Information Rate (CIR), 543
- Common Desktop Environment (CDE), 626
- Common Internet File System (CIFS), 354
- communication cable plenum (CMP), 166
- communication cable riser (CMR), 166
- communication protocols, e-mail, 45, 226
- communication servers, 45
- complementary metal oxide semiconductor (CMOS), 9
- computer bus, 7
- computer management, centralized, 358
- computer(s)
 - acronyms, 49–50
 - basic functions of, 2–3
 - boot procedure, 9–11
 - client, 33
 - input components, 3
 - output components, 4
 - overview of concepts, 2–14
 - processing components, 3–4
 - standalone upgraded to networked, 15–17
 - storage components, 4
- configuration
 - ad hoc Wi-Fi, 141–144
 - NIC drivers, 82–83
- connection hardware, cables, 166
- connectionless protocols, 214–215
- connection-oriented protocols, 214–215
- connectivity troubleshooting, 591–592
- connectors
 - fiber-optic, 186–187
 - RJ-11, 171
 - RJ-45, 172
- contact management software for small business, 519
- convergence, speed of, 307
- cooperative multitasking, 341
- copy backups, 605
- CPE (customer premises equipment), 547–548
- CPU (central processing unit)
 - architecture, 381
 - monitoring utilization, 345
 - multicore, 3–4
 - upgrading, 513
- cracker, 488
- CRC (Cyclical Redundancy Check), 126, 224
 - Data Link layer and, 273, 274
- credentials, 35. *See also* passwords
- CRM (customer relationship management), 519
- crossover cable, 179
- crosstalk, 165
- CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance), 140–141
- CSMA/CD (Carrier Sense Multiple Access with Collision Detection), 124–125, 127, 140
- CSU/DSU (Channel Service Unit/Data Service Unit), 536, 540
- customer equipment, WANs, 547
- customer premises equipment (CPE), 547
- customer relationship management (CRM), 519
- cut-through switching, 293
- Cyclical Redundancy Check (CRC), 126, 224
 - Data Link layer and, 273, 274

D

- daily backups, 605
- data access security, 465–487
 - authentication and authorization, 358, 466
 - encryption, 418–419, 472–474
 - firewalls, 477–482
 - spyware and spam protection, 484–485
 - troubleshooting, 648
 - VPNs, 474–477
 - wireless networks, 486–487
 - worm, virus, and rootkit protection, 483–484
- database server, 365
- Datacenter Edition, Windows, 382
- data circuit-terminating equipment (DCE), 548
- data collector set, 438, 440–441
- data frames. *See* frame(s)
- Data Link layer, OSI reference model, 273–274
- Data Over Cable Service Interface Specification (DOCSIS), 145, 147
- data protection, 461
 - checksum and, 224
- data segmentation, 223, 224–225
- data terminal equipment (DTE), 548
- DCE (data circuit-terminating equipment), 548
- DCE (Distributed Computing Environment), 626
- DDNS (Dynamic DNS) naming service, 229, 349
- DDoS (distributed denial-of-service) attacks, 492
- decimal
 - converting binary to, 241–242
 - converting to binary, 240–241
- dedicated bandwidth, 68
- deencapsulation, 268
- default gateway, 96
- default groups, 404–405
- default routes, 95
- default tolerance, server and, 364–365
- demarcation point, 548
- denial-of-service (DoS) attacks, 491, 492
- desktop computers in small-business networks, 512
- destination network, 303
- device drivers, NIC, 82–83
- DHCP (Dynamic Host Configuration Protocol), 46
 - protocol selection, 226–227
 - server, 361–362
 - working with DHCP client, 231, 346–347, 350–351
- Dial-Up Connection, in Windows 7, 553–555
- Dial-Up Networking (DUN) software, 551–552
- dial-up remote access, small-business networks, 525
- differential backups, 605
- differential signal, 183–184
- digital certificates, 472–473
- digital modem, 535, 538
- digital signal, 535, 536
- Digital Subscriber Line (DSL), 148
- directory servers, 44
- directory service, 36
- disaster recovery, 604–611
 - backing up data, 604–606
 - system repair and recovery, 606–608
 - System Restore point, 608–611
- disk-drives
 - RAID, 365, 381, 447–449, 512
 - securing data, 473–474
 - space organization, 332–333
- disk management
 - Linux disk utility, 425–427
 - Windows, 421–423
- disk mirroring, 448, 512
- disk quotas, 416–417
- disks, partitions in, 414–415
- disk striping with parity, 448–449, 512
- disk subsystem, 381
- disk systems, redundant, 447
- distance-vector protocols, 306
- Distributed Computing Environment (DCE), 626
- Distributed Management Environment (DME), 626
- distribution racks, 172, 173
- DME (Distributed Management Environment), 626
- DNS (Domain Name System), 46
 - client, 347–352
 - Command-line program Nslookup, 351–352
 - protocol selection, 226
 - server, 362–364
 - working with DNS tools, 229–231
 - zones, 363
- DOCIS (Data Over Cable Service Interface Specification), 145, 147
- documentation, 566–570
 - network changes and, 567
 - network security and, 568–569
 - as problem-solving resource, 584–589
 - solutions in problem-solving process, 585
 - troubleshooting and, 567–568
- document transfer, 43
- domain controllers, 36, 44
- domain local groups, 403
- domain model, 383
- Domain Name System (DNS), 46

- client, 347–352
- protocol selection, 226
- server, 362–364
- working with DNS tools, 229–231
- zone, 363
- domains, 36
 - broadcast, 92
 - collision, 125–126
- DoS (denial-of-service) attacks, 491, 492
- dotted decimal notation, 232
- drive mapping, 354–355
- Driver Rollback feature, 608
- drivers, 587
 - NIC, 17, 82–83
 - problems with, 646
- DSL (Digital Subscriber Line), 148
- DTE (data terminal equipment), 548
- dual attachment station (DAS), 641
- DUN (Dial-Up Networking) software, 551–552
- dynamic disk, 415
- Dynamic DNS (DDNS) naming service, 229, 349
- Dynamic Host Configuration Protocol (DHCP), 46
 - protocol selection, 226
 - server, 361–362
 - working with DHCP client, 231, 346–347, 350–351
- E**
- eavesdropping, network media and, 165–166
- EFS (Encrypting File System), 418, 473–474
- EIA (Electronic Industries Alliance), 623
 - UTP cabling categories, 168–169
- 802 standards, 262, 275–278, 624–625
 - 802.3ba standard, 131
 - Ethernet 802.2 frame type, 636
 - Ethernet 802.3 frame type, 127, 636
 - extensions to OSI reference model, 262, 277–278
 - 10 G 802.3ae standard, 130–131
 - 802.3i, 127
 - 802.11i standard, 487
 - Project 802, 275–276
 - specifications, 276–277
 - 802.3u, 128
 - 802.11 Wi-Fi, 137–138, 141, 149
 - electromagnetic interference (EMI), 165
 - differential signal and, 183–184
 - electronic eavesdropping, 165–166
 - Electronic Industries Alliance (EIA), 623
 - UTP cabling categories, 165
 - e-mail (electronic mail) servers, 356–357
 - communication protocols, 45, 226
 - hoax virus and, 484
 - spam, 485
 - EMI (electromagnetic interference), 165
 - differential signal and, 183–184
 - encapsulation, 32, 268
 - encoding, 164, 274
 - Encrypting File System (EFS), 418, 473–474
 - encryption, 418–419
 - IPSec, 472–473
 - end-to-end testing, 183
 - end user licensing agreement (EULA), 518
 - Enterprise Edition, Windows, 382
 - entrance facilities, 175–176
 - entrepreneurs, small business, 522
 - environmental considerations, cable selection criterion, 196
 - equipment
 - networking, 512–515
 - WANs, 547–548
 - equipment rooms, 175, 569–570
 - equipment sharing, small-business networks, 510–511
 - error handling, Ethernet, 126
 - ESX server, 378, 380
 - Ethernet, 62, 122–137, 545
 - coaxial cable, 122
 - collisions and collision domains, 125–126
 - determining and changing standard, 133–136
 - error handling, 126
 - 40 Gigabit, 131
 - frame types, 123–124, 635–636, 642
 - Gigabit (1000BaseT), 129
 - half-duplex *vs.* full-duplex communications, 127
 - media accessing, 123, 124–125
 - 100 Gigabit, 131
 - 100BaseFX, 128–129
 - 100BaseT4, 130
 - 100BaseTX, 128
 - 100VG-AnyLAN standard, 634–635
 - 1000BaseCX, 130
 - 1000BaseLX, 130
 - 1000BaseSX, 130
 - overview, 122–123
 - SNAP, 636
 - standards terminology, 127
 - 10Base2, 633–634
 - 10Base5, 631–633
 - 10GBaseT, 129
 - 10 Gigabit, 130–131
 - thickwire (thicknet), 630–634
 - thinwire (thinnet), 629–630, 633
 - viewing frame, 136–137
 - Ethernet 802.2 frame type, 636

- Ethernet 802.3 frame type, 127, 636
- Ethernet SubNetwork Address Protocol (SNAP), 636
- EtherTalk, 640
- EULA (end user licensing agreement), 518
- Event Viewer, 435–436
- exclusion, 362
- experience as problem-solving tool, 584–586
- extended LANs, 190, 194
- extended partition, 414–415
- extended star topology, 113–114
- F**
- failover cluster, 365
- FAQs (frequently asked questions), 586
- FAT (File Allocation Table), 415–416, 472
- fault tolerance, 318, 442–443
 - redundant disk systems, 447–449
 - redundant power and, 446–447
 - server and network, 364–365
- fault-tolerant storage, 512
- fax servers, 45
- fax service problems, 647
- FDDI. *See* Fiber Distributed Data Interface
- Federal Communications Commission (FCC), 192
- Fiber Distributed Data Interface (FDDI)
 - architecture, 640–641
 - physical ring topology and, 115
 - token-passing access method, 145, 640–641
- fiber-optic cable, 122, 184–188
 - cable types, 187–188
 - connectors, 186–187
 - installation, 187
- File Allocation Table (FAT), 415–416, 472
- file and print servers, 45
- file indexing system, 334
- file-sharing client, 352–354
- file sharing in small-business networks, 504–505
 - Windows, 428–431, 510
 - Windows homegroup, 508–510
- file systems, 12–13, 332
 - compression and encryption, 418–419
 - disk quotas, 416–417
 - FAT, 415–416
 - hierarchical, 333–334
 - Linux, 337–339, 427–428
 - NTFS permissions, 416, 419–421, 423–425, 472
 - secure access to, 334
 - shadow copies, 417–418
 - user profile, 408
 - Windows, 334–337, 421–423
- File Transfer Protocol (FTP), 46, 352
- Finger utility, 490
- firewalls, 477–482
 - IDS, 478
 - NAT, 479
 - routers, 478
 - small-business networks, 507–508, 524
 - testing with ShieldsUP!, 481–482
 - in Windows, 479–481
- flood
 - ping, 492
 - in switches, 290
- flow control, 225
- fragment-free switching, 293
- frame fragment, 293
- frame relay networks, 542–544
- frame(s)
 - broadcast, 295
 - creating, 280
 - Data Link Layer, 273
 - forwarding methods, 293
 - identifying TCP/IP layers in, 211–212
 - multicast, 294
 - packets and, 29–32
- frame types, Ethernet, 123–124, 635–636, 642
- frequently asked questions (FAQs), 586
- FTP (File Transfer Protocol), 46, 352
- full backup, 605
- full-duplex communications, 68, 127
- fully qualified domain name (FQDN), 228
- G**
- gatekeeper, 419
- gateway(s), 96
 - Terminal Services, 552
- gateway-to-gateway VPN mode, 517
- GBIC (gigabit interface converter), 513
- General Public License (GPL), 385
- Gigabit Ethernet, 129
- gigabit interface converter (GBIC), 513
- global groups, 403–404
- graphical user interface (GUI), 12
- group accounts, 396
 - See also* homegroup
 - default, 404–405
 - global, 403–404
 - Linux, 410–414
 - local, domain, 403
 - permissions, 421
 - special identity, 405

- universal, 404
- Windows, 403–407
- guest OS (operating system), 367
- Guest user, Windows, 397–398, 402–403
- GUI (graphical user interface), 12, 385
 - in Linux accounts, 412–414
- H**
- hackers, 488
- half-duplex communications, 68, 127
- hard drive, 365
 - fundamentals, 7–8
- hardware
 - interface to, 13–14
 - network server requirements, 45–47
 - PC, 5–11
 - token ring, 638
 - Windows, 46–47
- header, 32
- hertz (Hz), 190
- Hewlett-Packard (HP), 511
- hexadecimal notation, 81, 249
- hierarchical filing method, 333–334
- hoax virus, 484
- homegroup, 505–510
 - creating and joining, 505–506
 - sharing files with, 508–510
 - troubleshooting, 507–508
- hop count, 304
- hop, 303–304
- horizontal wiring, 174, 175
- host addresses, determining, 245
- host computer, 367
- host ID, IPv6, 249–250
- hosted applications, 520, 556–557
- hosted infrastructure, 557–558
 - hosted platforms, 557
 - hosted virtualization, 367–369
 - See also* virtualization
 - hot-swappable device, 364–365
 - HTTP (Hypertext Transfer Protocol), 225, 226
 - client, 352
 - hubs
 - indicator lights, 64–65, 75–77
 - network bandwidth and, 64–65
 - repeaters and, 62–66
 - Wireshark with, 72–73
 - Hypertext Transfer Protocol (HTTP), 225, 227
 - Hyper-V, 378, 379
 - hypervisor, 367, 378
 - Hz (Hertz), 190
- I**
- IAB (Internet Architecture Board), 623–624
- IBM, 636, 638
- ICANN (Internet Corporation for Assigned Names and Numbers), 624, 627
- ICMP (Internet Control Message Protocol), 218, 219–220
- IDSs (intrusion detection systems), 478
- IEEE. *See* Institute of Electrical and Electronics Engineers (IEEE)
- IESG (Internet Engineering Steering Group), 624
- IETF (Internet Engineering Task Force), 233, 623–624
- IMAP (Internet Message Access Protocol), 226, 356, 357
- incremental backups, 605
- indicator lights, 16, 69
 - hub, 65–66, 75–77
 - NIC, 81–82
- information gathering in problem-solving process, 579–580
- infrared (IR) wireless LAN networks, 191
- infrastructure, hosted, 557–558
- infrastructure as a service (IaaS), 557–558
- infrastructure mode, 138
- infrastructure services, 361–364
- installing, 380–385
 - cable, 167, 172–178
 - Linux, 384–385
 - NICs, 16
 - OSs, 380–385
 - protocol analyzers, 69–72
 - Windows servers (*See* installing Windows servers)
 - Wireshark, 69–72
- installing network services
 - Linux, 384–385
 - Windows, 380–384
- installing OSs, 380–385
 - hardware selection, 381–382
 - Linux, 384–385
 - preinstallation decisions, 383
- installing Windows servers, 380–384
 - planning for, 380–381
 - postinstallation tasks, 384
 - preinstallation decisions, 383–384
 - selecting hardware, 381–382
 - selecting right edition, 382
- Institute of Electrical and Electronics Engineers (IEEE), 275, 624–625
 - 802 standards (*See* 802 standards)
- Integrated Services Digital Network (ISDN), 537, 538–539
- interconnecting devices, 15
 - See also* hubs; routers; switch(es)
- interference, cable susceptibility, 165

- International Organization for Standardization (ISO), 262, 625
 - Internet
 - See also* World Wide Web
 - access technologies, 145–149
(*See also* cable modems)
 - small-business access, 515–516
 - use policies, 461
 - WANs *versus* VPNs, 546
 - Internet Architecture Board (IAB), 623–624
 - Internet Control Message Protocol (ICMP), 218, 219–220
 - Internet Corporation for Assigned Names and Numbers (ICANN), 624, 627
 - Internet Engineering Steering Group (IESG), 624
 - Internet Engineering Task Force (IETF), 233, 623–624
 - Internet Message Access Protocol (IMAP), 226, 356–357
 - Internet Network Information Center (InterNIC), 624
 - Internet Protocol (IP) addresses, 232–251
 - address classes, 232–233
 - ARP, 216–218
 - binary arithmetic, 239–242
 - CIDR, 236
 - configuring, 211
 - defined and verified, 213–214
 - DHCP, 361–362
 - Ipconfig, 22, 23, 24
 - IPv4, 211, 215–216
 - IPv6, 216, 248–251
 - MAC addresses and, 214
 - NAT, 234–235
 - private addresses, 233
 - routers and, 93
 - scope, 361–362
 - subnet masks, 236–239, 243–246
 - supernetting, 247
 - Windows Server 2008, 383
 - Internet Protocol version 4 (IPv4), 211, 215–216
 - Internet Protocol version 6 (IPv6), 216, 248–251
 - Internet Research Task Force (IRTF), 624
 - Internet Service Provider (ISP), 190
 - Internetworking devices
 - documentation, 570, 588
 - physical security, 465
 - Internetwork layer, 209–210, 213
 - packet routing through, 214
 - protocols at, 215–219
 - Internetwork Packet Exchange (IPX), 643
 - Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX), 642–643
 - internetwork(s)
 - See also* WANs
 - LANs *versus*, 26
 - InterNIC (Internet Network Information Center), 624
 - intrusion detection systems (IDSs), 478
 - IP addresses. *See* Internet Protocol (IP) addresses
 - IP prefix, 236
 - IP Security (IPSec), 218–219
 - encryption and, 472–473
 - L2TP and, 475
 - NICs and, 318
 - IPv4 (Internet Protocol version 4), 211, 215–216
 - IPv6 (Internet Protocol version 6), 248–251
 - addresses, 216, 249
 - host ID, 249–250
 - subnetting with, 250–251
 - IPX (Internetwork Packet Exchange), 643
 - IPX Routing Information Protocol (IPX RIP), 643
 - IPX/SPX (Internetwork Packet Exchange/Sequenced Packet Exchange), 642–643
 - IrDA devices, 191
 - IRTF (Internet Research Task Force), 624
 - ISDN (Integrated Services Digital Network), 537, 538–539
 - ISO (International Organization for Standardization), 625
 - IT staffing, documentation and, 568
- ## K
- Kerberos authentication, 472
 - kernel, of operating system, 341–342
 - knowledge base, as problem-solving tool, 586
- ## L
- LANs. *See* local area networks
 - laser-based LAN technologies, 191
 - Last Known Good Configuration, 607
 - last mile (local loop), 548–549
 - layer(s)
 - Internetwork, 209–210, 213, 214, 215–219
 - network communication, 18–19
 - OSI reference model, 265, 266–274, 279–280
 - TCP/IP, 208–212
 - viewing network software, 21–22
 - Layer 2 Tunneling Protocol (L2TP), 475
 - LC (locking connection) connectors, 186
 - LDAP (Lightweight Directory Access Protocol), 36, 44, 359
 - leased line networks, 539–541, 543–544
 - least privileges principle, 420
 - Lightweight Directory Access Protocol (LDAP), 36, 44, 359

- line-of-sight networks, 191
 - Link Control Protocol (LCP), 551
 - link light, 16. *See also* indicator lights
 - link-state algorithms, 307
 - link-state protocols, 307
 - Linux
 - disk naming in, 415
 - disk utility, 425–427
 - displaying processes, 344–345
 - encrypting files, 473, 474
 - file sharing, 354, 434, 510
 - file system, 337–339, 427–428
 - groups, 410–414
 - installing servers, 384–385
 - logon restrictions, 44, 471
 - NIS logon service, 36, 359
 - passwords, 409–410, 467, 470–471
 - printer sharing in, 434–435
 - Samba, 354, 434
 - security policy, 359
 - small-business networks, 520–522
 - user accounts, 409–414
 - VPNs, 476
 - LLC (Logical Link Control), 278
 - load-balancing cluster, 365
 - Local Area Connection Properties dialog box, 21, 22
 - local area networks (LANs), 19–21, 26
 - bandwidth and, 65
 - extended, 190
 - internetworks *versus*, 26–27
 - media selection criteria, 196–197
 - routers and, 88–92
 - virtual, 296–299
 - WANs and, 534
 - wireless (*See* wireless LANs)
 - local loop, 548–549
 - local profile, 408
 - LocalTalk, 640
 - locking connection (LC) connectors, 186
 - lockout policies, 468–469
 - See also* passwords
 - logical bus topology, implemented as physical star, 118
 - Logical Link Control (LLC) sublayer, 278
 - logical ring topology, implemented as physical star, 119
 - logical topology, 110, 117–120
 - logon restrictions, 470–471
 - loopback address, 233
 - L2TP (Layer 2 Tunneling Protocol), 475
- M**
- MAC addresses, 19, 25
 - ARP and, 216–218
 - Data Link layer and, 273–274
 - Ethernet and, 123
 - filtering feature, 486–487
 - IP addresses and, 214
 - NICs and, 80–81
 - packet forwarding and, 302–303
 - switching table and, 66–68, 292–293
 - viewing, 278–279
 - mail servers, 45, 226, 356–357
 - malware, 483
 - managed switches, 294, 513
 - management
 - See also* network management
 - small-business network, 524–525
 - management information bases (MIBs), 601, 612
 - mandatory (user) profile, 409
 - MAN (metropolitan area networks), 28, 29
 - maximum transmission unit (MTU), 270
 - MDI crossed (MDI-X) devices, 179
 - mechanical transfer registered jack (MT-RJ), 187
 - media access, token ring, 144–145
 - Media Access Control (MAC), 124, 278
 - See also* MAC addresses
 - medium dependent interface (MDI) devices, 179
 - medium interface connectors (MIC), 186
 - memory management, 12
 - mesh topology, 116–117
 - metrics, 304
 - metropolitan area networks (MAN), 28
 - MIBs (management information bases), 601, 612
 - MIC (medium interface connectors), 186
 - Microsoft
 - See also* Windows *entries*
 - Hyper-V, 378, 379, 686–688
 - Network Monitor, 597, 598–600
 - redirector, 352
 - Security Essentials, 485–486
 - Server Core, 46
 - Small-Business Server, 519–520
 - Virtual PC, 369, 373–374, 652–659
 - microwave networking
 - satellite, 195–196
 - terrestrial, 195
 - mobile computing, 190
 - modems, 535–536
 - analog, 535
 - cable, 145–147
 - digital, 535, 538
 - V.92, 537–538

- monitoring network performance, 435–442
 - advanced tools, 601–602
 - CPU utilization, 343, 344
 - data collector set, 438, 440–441
 - Event Viewer, 435–436
 - Network Monitor, 437–440, 594
 - Performance monitor, 437–441
 - Task Manager, 342–344
- motherboard, 5–7
- Motif Toolkit API, 626
- MPLS (Multiprotocol Label Switching) networks, 545–546
- MSAUs (multistation access units), 638
- MT-RJ (mechanical transfer registered jack), 187
- MTU (maximum transmission unit), 270
- multicast processing, 294
- multimode fiber (MMF) cables, 187
- multiplexing (“muxing”), 540
- multiport repeaters, 63–64
 - See also* hubs
- multiprocessing, 341
- Multiprotocol Label Switching (MPLS) networks, 545–546
- multistation access units (MSAUs), 638
- multitasking, 13, 340–341
- multithreaded application, 341
- “muxing” (multiplexing), 540
- N**
- NADN (nearest active downstream neighbor), 637
- name server, 20, 383
- naming services
 - See also* DNS (Domain Name System)
 - NetBIOS names, 644
- narrowband radio, 192
 - LAN technologies, 192–193
- NAS (network-attached storage), 360, 510
- NAT (Network Address Translation), 234–235, 370, 479
- NAUN (nearest active upstream neighbor), 637
- NCP (NetWare Core Protocol), 643
- NDIS (Network Device Interface Specification), 642
- nearest active downstream neighbor (NADN), 637
- neighbors, 306, 637
- NetBEUI (NetBIOS Extended User Interface), 643–644
- NetBIOS (Network Basic Input/Output System), 643
 - names, 644
- NetInfo, 492–494
- Netstat Program, 221–222
- NetWare, 642
- NetWare Core Protocol (NCP), 643
- Network access layer, 210, 213
- network accounts, problems with, 647
- Network Address Translation (NAT), 234–235, 370, 479
- network appliance, 360
- network applications. *See* applications
- network architecture
 - ATM, 544–545
 - Ethernet (*See* Ethernet)
 - FDDI, 640–641
 - token ring, 636–639
- network-attached storage (NAS), 360, 510
- network backbone, 115
- network bandwidth, 64–65
 - cable selection criteria and, 164–165, 196
- switches and, 68
- wireless APs and, 78–79
- Network Basic Input/Output System (NetBIOS), 643, 644
- network client software, 16
- network communication
 - See also* local area networks (LANs)
 - components, 14–15
 - fundamentals of, 14–25
 - layers of process, 18–19
 - steps of, 17–18
 - testing network addresses, 22–25
 - viewing layers in, 21–22
- Network Control Protocol (NCP), 552
- network description, 569
- Network Device Interface Specification (NDIS), 642
- network diagrams, 588, 589
- Network File System (NFS), 354, 428, 626
- Network Information Service (NIS), 36, 359
- networking acronyms, 49–50
- network interface cards (NICs), 15, 79–88
 - advanced features, 317–318
 - basics, 79–80
 - configuration, 82–83
 - driver software, 17, 82–83
 - examining properties of, 86–88
 - as gatekeeper, 81
 - indicator lights, 81–82
 - MAC addresses and, 80–81
 - PC bus options, 314–317
 - selecting, 82
 - troubleshooting, 602–603, 646
 - USB, 316–317
 - wireless, 84–86

- Network layer, OSI reference model, 272
 - network management
 - groups (Linux), 410–414
 - groups (Windows), 403–407
 - small-business networks, 524–525
 - user accounts, 396–414
 - network media, 15
 - See also* cable
 - access, 123, 124–125
 - access, token ring, 144–145
 - criteria for choosing, 164–167
 - LAN, selection criteria, 196–197
 - technology and, 121–122
 - network models, 34–43
 - Network Monitor (Microsoft), 597, 598–600
 - network monitor(s), 437–440, 594
 - See also* monitoring network performance
 - network operating systems (NOSs), 332–366
 - See also* file systems
 - client operating systems, 346–357
 - fundamentals, 332–345
 - overview, 345–366
 - server operating system, 357–366
 - network operations problems, troubleshooting, 646–647
 - network performance
 - establishing baseline, 437–438
 - monitoring (*See* monitoring network performance)
 - troubleshooting, 603–604
 - network protocols, 16–17
 - See also* protocols
 - network security. *See* security; security attacks
 - network servers, 44–47
 - application servers, 45
 - communication servers, 45
 - directory servers, 44
 - fax servers, 45
 - file and print servers, 45
 - hardware requirements, 45–47
 - mail servers, 45
 - Web services, 46
 - network server software, 16
 - network technologies and media, 121–122
 - network unreachable, 95
 - newsgroups, consulting, 587
 - New Technology file system (NTFS), 416
 - permissions, 419–421, 423–425, 472
 - next hop, 303–304
 - NFS (Network File System), 354, 428, 626
 - NICs. *See* network interface cards (NICs)
 - NIS (Network Information Service), 36, 359
 - NOSs. *See* network operating systems
 - Nslookup command-line program, 351–352
 - NTFS (New Technology file system), 416
 - encryption, 418
 - permissions, 419–421, 423–425, 472
 - NWLink, 642
- O**
- Object Management Group (OMG), 625–626
 - ODI (Open Data-link Interface), 642
 - OMG (Object Management Group), 625–626
 - onboard co-processors, 318
 - 100BaseFX Ethernet, 128–129
 - 100BaseT4 standard, 130
 - 100BaseTX Ethernet, 128
 - 1000BaseCX standard, 130
 - 1000BaseLX standard, 130
 - 1000BaseSX standard, 130
 - 1000BaseT Ethernet, 129
 - 100VG-AnyLAN standard, 634–635
 - online periodicals, 587
 - online support services, as problem-solving tools, 587
 - Open Data-link Interface (ODI), 642
 - The Open Group (TOG), 626–627
 - open security policies, 463
 - open source software, 520
 - Open Systems Interconnection reference model. *See* OSI reference model
 - operating systems (OS)
 - See also* client operating systems; server operating systems
 - fundamentals, 332–345
 - hardware and, 12–14
 - installing, 380–385
 - kernel, 341–342
 - virtualization (*See* virtualization)
 - OSI reference model, 262–275
 - hands-on troubleshooting with, 593–594
 - layers, 265, 268–274, 279–280
 - peer communication in, 266–267
 - protocols, 268
 - step-by-step troubleshooting method, 574–576
 - structure, 264–268
 - summary of, 274–275
 - TCP/IP model *versus*, 265–266, 275
- P**
- packet assembler/disassemblers (PADs), 542
 - packet filtering, 308
 - firewalls, 478

- packet forwarding, 302–303
- packet(s), 290
 - frames and, 29–32
 - Network layer, 272
 - routed through Internetwork layer, 214
- packet-switching networks
 - frame relay, 542–544
 - virtual circuits, 541–542
 - WANs, 541
 - X.25 specification, 542
- PADs (packet assembler/disassemblers), 542
- PAM (Pluggable Authentication Modules), 467
- partitions, in file systems, 414–415
- passwords, 34–35, 358
 - cracking, 491
 - dos and don'ts, 467–468
 - Linux environment, 409–410, 467, 470–471
 - rules for creating, 396–397
 - shadow, 467
 - small-business networks, 506, 523
 - Windows environment, 399–400, 466–469
- patch cables, 171, 179–181
- patch panels, 172, 176
- PAT (Port Address Translation), 235
- PCI Express (PCIe) buses, 315, 316
- PCI (Peripheral Component Interconnect) buses, 314
- PCI-X (Peripheral Component Interconnect-Extended) buses, 314
- PCMCIA cards, 315–316
- PCM (pulse code modulation), 538
- PC (Peripheral Component) buses, 314–317
- PDAs (personal digital assistants), 47, 48
- PDU (protocol data unit), 268
 - Network layer, 272
 - Transport layer, 271
- peer communication, 266–267
- peer-to-peer networking, 34–36, 504
 - advantages and disadvantages, 35–36
 - exploring, 38–40
 - server-based networks *versus*, 36–37
- penetration tester, 487
- Performance Monitor, 437–441
- Performance tab, Task Manager, 344
- Peripheral Component Interconnect (PCI) buses, 314
- permanent virtual circuits (PVCs), 542
- permissions
 - See also* file sharing
 - NTFS, 419–421, 423–425, 472
- personal area networks, wireless (WPANs), 48–49
- personal computer (PC) hardware, 5–11
- physical bus topology, 110–112
 - limitations of, 112
 - signal bounce, 112
 - signal propagation, 111–112
- Physical layer, OSI reference model, 274, 576
- physical ring topology, 115
- physical security, 463–465
 - best practices, 463–464
 - internetworking devices, 465
 - servers, 464–465
- physical star topology, 112–115
 - logical bus implemented as, 118
 - logical ring implemented as, 119
 - network building, 120–121
- physical topology, 110–117
 - See also* topology
- ping, 22, 24, 25, 590–592
- ping floods, 492
- ping scanner, 488, 489
- plain old telephone service (POTS), 537–538, 539
- platform as a service (PaaS), 557
- Pluggable Authentication Modules (PAM), 467
- Point-to-Point Protocol (PPP), 551–552
- point-to-point topology, 116
- Point-to-Point Tunneling Protocol (PPTP), 475
- POP3 (Post Office Protocol version 3), 45, 226, 356–357
- Port Address Translation (PAT), 235
- port forwarding, 516, 517
- port numbers, 223, 271
- ports
 - switch, 290–291
 - uplink, 66, 75–77
- port scanner, 488–489, 490–491
- port security, 299
- Post Office Protocol version 3 (POP3), 45, 226, 356–357
- POST (power-on self test), 8
- POTS (plain old telephone service), 537–538, 539
- power conditioning, 447
- power fluctuations, troubleshooting, 603
- power-on self test (POST), 8
- power supply
 - redundant, 446–447
 - uninterruptible (UPS), 447, 464
- PPP (Point-to-Point Protocol), 551–552
- PPTP (Point-to-Point Tunneling Protocol), 475
- Preboot Execution Environment (PXE), 318

- preemptive multitasking, 340
 - Presentation layer, OSI reference model, 269
 - preshared keys, 472
 - preventative measures, 584
 - primary partition, 414
 - Primary Rate Interface (PRI), 539
 - printer sharing, 355–356, 428, 432–435
 - in small-business networks, 510–511
 - printing, problems with, 647
 - PRI (Primary Rate Interface), 539
 - privacy policy, 461
 - private networks, virtual. *See* virtual private networks (VPNs)
 - problem-solving process, 576–584
 - considering possible causes, 580–581
 - determining definition and scope, 578–579
 - devising preventive measures, 584
 - devising solutions, 581–582
 - documenting solutions, 584
 - gathering information, 579–580
 - implementing solutions, 582
 - solve by example, 573–574
 - testing solutions, 582
 - problem-solving resources
 - drivers, 587
 - experience, 584–586
 - knowledge base, 586
 - network documentation, 587–589
 - newsgroups, 587
 - online periodicals, 587
 - online support services, 587
 - TDR, 600–601
 - updates, 587
 - World Wide Web, 586–587
 - processes and services, 340–341
 - Processes tab, Task Manager, 344
 - Project 802, 275–276
 - See also* 802 standards
 - protocol analyzers, 490, 597–598
 - See also* Wireshark
 - downloading and installing, 69–72
 - protocol data unit (PDU), 268
 - Network layer, 272
 - Transport layer, 271
 - protocol(s)
 - connectionless *versus* connection-oriented, 214–215
 - network, 16–17
 - OSI reference model, 268
 - routable *versus* nonroutable, 642–644
 - protocol suites, 642–644
 - AppleTalk, 644
 - IPX/SPX, 642–643
 - NetBEUI, 643–644
 - NetBIOS, 643–644
 - provider equipment, WANs, 548–549
 - public data networks (PDNs), 542
 - public switched telephone network (PSTN), 536
 - pulse code modulation (PCM), 538
 - PVCs (permanent virtual circuits), 542
 - PXE (Preboot Execution Environment), 318
- Q**
- Quality of Service (QoS), 248, 318
- R**
- radio
 - narrowband, 192–193
 - radio frequency interference (RFI), 165
 - spread-spectrum, 193
 - RAID (redundant array of independent disks), 365, 381
 - fault tolerance and, 447–449, 512
 - random access memory (RAM)
 - buffering, 317
 - fundamentals, 8
 - short-term storage, 4–5
 - Windows 2008, 381
 - recovery. *See* disaster recovery
 - redirector, 352
 - redundant array of independent disks (RAID), 365, 381, 447–449
 - redundant power supply, 446–447
 - reference models
 - OSI (*See* OSI reference model)
 - role of, 263–264
 - reflection (signal bounce), 112
 - reflective wireless networks, 191
 - registered jack 45 (RJ-45) connectors, 171
 - remote access, 365, 549–555
 - dial-up, 551–552, 553–555
 - RRAS, 45, 475–476, 549
 - small-business networks, 471, 516–517
 - Telnet, 525
 - VPN, 517, 550–551, 555
 - Web networking, 552–553
 - Remote Monitoring (RMON), 602
 - repeaters
 - hubs and, 62–66
 - multiport, 63–64
 - replacement method, troubleshooting, 574
 - request to send (RTS), 78
 - reservation, 362

- resources
 - online/electronic materials, 649–650
 - printed materials, 649
 - problem-solving, 584–589
 - RFI (radio frequency interference), 165
 - ring topology, 115, 119
 - See also* token ring networks
 - RIP (Routing Information Protocol), 306
 - RIPv2 (Routing Information Protocol version 2), 306
 - RJ-11 connectors, 171
 - RJ-45 (registered jack 45) jack, 172
 - RJ-45 (registered jack 45) plug, 171, 178, 180–181
 - RMON (Remote Monitoring), 601–602
 - roaming profile, 408–409
 - root hints, 363
 - rootkits, 483–484
 - routers, 88–100, 300–310
 - broadcast domains and, 92
 - communicating over, 96–98
 - default gateway and, 96
 - default routes, 95
 - as firewalls, 478, 524
 - interfaces, 302–303
 - LANS and, 88–92
 - port forwarding, 516, 517
 - routing protocols, 305–308
 - small-business networks, 89
 - static routes, 304, 307–308
 - switches *versus*, 90, 91–92
 - tables (*See* routing tables)
 - Trace Route program and, 99–100, 592–593
 - VLANs and, 298–299
 - WANs and, 536
 - Routing and Remote Access Service (RRAS), 45, 475–476, 549
 - Routing Information Protocol (RIP), 306
 - Routing Information Protocol version 2 (RIPv2), 306
 - routing protocols, 305–308
 - routing tables, 93–95, 303–305
 - supernetting and, 247–248
 - viewing and changing, 308–310
 - RRAS (Routing and Remote Access Service), 45, 475–476, 549
- S**
- sales software for small business, 518–519
 - Samba (Linux), 354, 434
 - SAN (storage area network), 47, 361
 - SAP (Service Advertising Protocol), 643
 - SAS (single attachment station), 641
 - SATA drives, 381
 - satellite microwave, 195–196
 - satellite technologies, 148–149
 - scanners, 511
 - ping scanner, 488, 489
 - port scanner, 488–489, 490–491
 - scatter infrared networks, 191
 - SCSI (Small Computer System Interface), 365, 381
 - SC (straight connection) connectors, 186
 - SDSL (Symmetric Digital Subscriber Line), 148
 - search engine, in problem-solving, 586
 - secure access, to files, 334
 - Secure Communication Services, 626
 - Secure Shell (SSH), 525
 - Secure Socket Tunneling Protocol (SSTP), 475
 - security
 - attacks (*See* security attacks)
 - data access (*See* data access security)
 - documentation and, 568–569
 - IPSec, 218–219, 318, 472–473, 475
 - physical, 463–465
 - policy management, 359–360, 460–463
 - small-business networks, 515, 523–524
 - switch port, 299
 - troubleshooting, 647–648
 - wireless networks, 139–140, 312–313
 - security attacks, 487–494
 - crackers, 488
 - disabling network resources, 491–492
 - discovering network resources, 488–491
 - gaining access to network resources, 491
 - hackers, 488
 - NetInfo, 492–494
 - penetration testers, 487
 - Security Essentials, Microsoft, 485–486
 - security policies, 460–463
 - common elements, 463
 - development, 460–461
 - elements, 461
 - levels of security, 462
 - open, 463
 - segmentation, data, 223, 224–225
 - segment (unit of information), 209
 - Sequential Packet eXchange (SPX), 643
 - Serial Attached SCSI (SAS) disks, 365
 - server-based networks
 - advantages and disadvantages, 36–37
 - peer-to-peer networks *versus*, 36–37
 - server cluster, 365, 449–450
 - server/domain-based model, 36–37
 - Server Message Block (SMB), 353–354, 428, 644

- server operating systems, 33, 357–366
 - additional features, 365–366
 - centralized storage, 360–361
 - centralized user account and computer management, 358–360
 - infrastructure services, 361–364
 - role of, 357–358
 - server and network fault tolerance, 364–365
- server(s)
 - See also* client/server computing; network servers; server operating system
 - additional features, 365–366
 - clients and, 32–34
 - clustering, 365, 449–450
 - communication, 45
 - consolidation of, 377
 - documentation of, 570
 - installing, 380–385
 - network default tolerance and, 364–365
 - physical security, 464–465
 - small-business networks, 511–512
- Service Advertising Protocol (SAP), 643
- Service Lookup Protocol (SLP), 643
- services, processes and, 340–341
- service set identifiers (SSIDs), 84, 313
- Session layer, OSI reference model, 270, 644
- shadow copies, 417–418
- shadow passwords, 467
- shared adapter memory, 317
- shared folder, creating, 41–43
- shared printer, 355–356, 428
 - in Linux, 434–435
 - in small-business networks, 510–511
 - in Windows, 432–434
- shared system memory, 317
- sharing
 - See also* file sharing in small-business networks
 - permissions, 419–420, 429
- shielded twisted-pair (STP) cable, 168, 170
- ShieldsUP!, 481–482
- signal, differential, 183–184
- signal bounce, 112
- signal propagation, 111–112
- Simple Mail Transfer Protocol (SMTP), 45, 226, 356–357
- Simple Network Management Protocol (SNMP), 601–602
- Simple Server Monitor, 595–597
- single attachment station (SAS), 641
- single-mode fiber (SMF) cables, 187
- single UNIX specification, 626
- SLP (Service Lookup Protocol), 643
- small-business networks, 504–525
 - accounting software, 518
 - communication, 515–517
 - data and application sharing in, 504–505
 - dial-up remote access, 525
 - equipping, 511–517
 - file sharing, 428–431, 504–505, 508–510
 - homegroups, 505–510
 - identifying requirements, 517–522
 - Internet access, 515–516
 - Linux as alternative to Windows, 520–522
 - management, 524–525
 - peer-to-peer networks, 35
 - remote access in, 471, 515–517
 - routers, 89
 - sales and contact management software, 518–519
 - security, 515, 523–524
 - servers and desktops, 511–512
 - sharing equipment, 510–511
 - support, 522–525
 - VPN remote access, 517, 525
 - wireless connection, 514–515
- Small-Business Server, 519
- smart multistation access units (SMSAUs), 638
- SMA (subminiature type A) connectors, 186
- SMB (Server Message Block), 354, 428, 644
- SMSAUs (smart multistation access units), 638
- SMTP (Simple Mail Transfer Protocol), 45, 226, 356–357
- snapshot, 367
- SNAP (SubNetwork Address Protocol), 636
- SNMP (Simple Network Management Protocol), 601–602
- social engineering, 484
- software
 - See also* networking operating systems (NOSs); operating systems
 - client network, 16
 - firewall, 477
 - small-business applications, 518–519
 - viewing layers, 21–22
 - virtualization, 367–380
- software agents, 601
- software as a service (SaaS), 556
- solve-by-example method, troubleshooting, 573–574
- SONET (Synchronous Optical Network), 540–541
- spam, 484–485

- span, as cable selection criterion, 196
 - Spanning Tree Protocol (STP), 294–296
 - special identity groups, 405
 - special interest groups (SIG), 619–620
 - specialized networks, 47–48
 - SPI (stateful packet inspection), 478
 - spoofed address, 491
 - spread-spectrum LAN technologies, 193–194
 - spread-spectrum radio, 193–194
 - SPX (Sequential Packet Exchange), 642
 - spyware, 484–485
 - SQL (Structured Query Language), 626–627
 - SSIDs (service set identifiers), 84, 312–313
 - SSTP (Secure Socket Tunneling Protocol), 475
 - stand-alone computer, upgrading, 15–16
 - Standard Edition, Windows, 382
 - standards, 619–627
 - See also*
 - 10Base2, 633–634
 - 10Base5, 631–633
 - 10BaseT, 127–128
 - documentation and, 568
 - Ethernet, 127, 133–136
 - important bodies, 620–627
 - process for making, 619–620
 - standards bodies, 620–627
 - V-series, 622–623
 - X-series, 623
 - standards compliance, documentation and, 568
 - standards-making process, 619–620
 - startup repair, 606–607
 - star topology, 112–114
 - stateful packet inspection (SPI), 478
 - static route, 304, 307–308
 - storage, centralized, 360–361
 - storage area network (SAN), 47, 361
 - storage device, network-attached, 360
 - storage management, 414–428
 - See also* file systems
 - disk quotas, 416–417
 - encryption, 418–419
 - file compression, 418
 - shadow copies, 417–418
 - volumes and partitions, 414–415
 - store-and-forward switching, 293
 - STP (shielded twisted-pair) cable, 168, 170
 - STP (Spanning Tree Protocol), 294–296
 - straight connection (SC) connectors, 186
 - straight-through cable, 179
 - straight tip (ST) connectors, 186
 - structured cabling, 173–176
 - backbone cabling, 175, 176
 - entrance facilities, 175–176
 - equipment rooms, 175
 - horizontal wiring, 174, 175
 - telecommunications closet, 174–175
 - work area, 174
 - Structured Query Language (SQL), 626–627
 - ST (straight tip) connectors, 186
 - subminiature type A (SMA) connectors, 186
 - subnet mask, 236–239
 - calculating, 243–246
 - subnetting, 236
 - with IPv6, 250–251
 - SubNetwork Address Protocol (SNAP), 636
 - supernetting, 247
 - surge protection, 447
 - SVCs (switched virtual circuits), 542
 - switched virtual circuits (SVCs), 542
 - switch(es), 66–69, 290–300
 - advanced features, 294–299
 - basic operation, 66–68
 - frame forwarding methods, 293
 - indicator lights, 69, 75–77
 - logical functioning of, 119
 - port modes of operation, 291
 - routers *versus*, 90, 91–92
 - small-business networks, 513
 - spanning tree protocol, 294–296
 - VLANs, 296–299
 - Wireshark with, 74–75
 - switching loop, 294–295
 - observing, 299–300
 - switching table, 66–68, 292–293
 - Symmetric Digital Subscriber Line (SDSL), 148
 - Synchronous Optical Network (SONET), 540–541
 - SYN (synchronization) segment, 224, 491
 - system partition, 415
 - system reliability and performance, 435–442
 - system repair/recovery, 606–608
 - System Restore utility, 607–608
 - creating, 608–611
- ## T
- tape backups, 605–606
 - Task Manager, 342–344
 - T-carrier lines, 539–540
 - T1 lines, 539–540, 543
 - T3 lines, 539
 - TCO (total cost of ownership), 520–521

- TCP/IP. *See* Transmission Control Protocol/Internet Protocol
- TCP (Transmission Control Protocol), 222–225
 - three-way handshake and, 224
- TC (telecommunications closet), 174, 175, 569–570
- TDRs (time-domain reflectometers), 600–601
- technical support, 568
 - problem-solving and, 585–586
- Telecommunication Industries Association (TIA), 168
 - UTP cabling categories, 168–169
- telecommunications closet (TC), 174, 175
 - documentation of, 569–570
- Telnet, small-business networks, 525
- 100BaseTX Ethernet, 128–129
- 10Base2 standard (Ethernet), 633–634
- 10Base5 standard (Ethernet), 631–633
- 10BaseT Ethernet, 127–128
- 10GBaseT Ethernet, 129
- 10 Gbps IEEE 802.3ae standard, 130–131
- terminal adapter, 539
- Terminal Services Gateway, 552–553
- termination, cable, 176–177, 181–183
- terminators, 112
- terrestrial microwave, 195
- testability, of cable installation, 167
- testing, in virtual network, 377
- testing solutions, in problem-solving process, 582–583
- thicknet (thickwire Ethernet), 630–634
- thinnet (thin wire Ethernet), 629–630, 633
- thread, 341
- three-way handshake, 224
- TIA (Telecommunication Industries Association), 168
 - UTP cabling categories, 168–169
- time-division multiplexing (TDM), 540
- time-domain reflectometers (TDRs), 600–601
- timestamp, 304
- TOG (The Open Group), 626–627
- token-passing access method, 144–145, 640–641
- token ring networks, 144, 636–639
 - beaconing, 637–638
 - cabling in, 638–639
 - function, 637
 - hardware components, 638
 - token-passing access method, 144–145, 640–641
- Token Talk, 640
- topologies
 - ad hoc, 141–144
 - extended star, 113–114
 - logical, 110, 117–120
 - logical bus, 118–119
 - logical ring, 119
 - mesh, 116, 117
 - physical, 110–117
 - physical bus, 110–112
 - physical star (*See* physical star topology)
 - point-to-point, 116
 - ring, 115, 119
- total cost of ownership (TCO), 520–521
- TP cable. *See* twisted-pair (TP) cable
- TPM (Trusted Platform Module), 473
- Trace Route, 99–100, 592–593
- traffic monitoring, with Wireshark, 594–595
- trailer, 32
- Transmission Control Protocol/Internet Protocol (TCP/IP), 208–212, 383
 - identifying layers in frame, 211–212
 - layered architecture, 208–210
 - OSI model *versus*, 265–266, 275
 - transport layer in, 222
 - viewing layers in Windows, 210–211
- Transmission Control Protocol (TCP), 222–225
 - transparent mode, EFS encryption, 473
- Transport layer, 209, 222–225, 270
 - OSI reference model, 270–271
- trial-and-error approach, to troubleshooting, 571–573
- Trojan program, 483
- troubleshooting, 571–576
 - See also* problem-solving process
 - cabling and related components, 602–603, 646
 - documentation and, 567–568
 - driver problems, 646
 - general questions for, 645–648
 - homegroups, 508
 - network operations problems, 646–647
 - network tools, 589–602
 - OSI model, 574–576, 593–594
 - permissions and, 420
 - poor network performance, 603–604
 - power fluctuations, 603
 - printing and fax service problems, 647
 - replacement method, 574
 - solving by example approach, 573–574
 - step-by-step approach with OSI model, 574–576
 - trial-and-error approach, 571–573
 - upgrades, 603

- troubleshooting tools, 589–602
 - advanced monitoring tools, 601–602
 - cable testers, 601
 - network monitors, 594–600
 - ping and Trace Route for, 590–593
 - protocol analyzers, 597–598
 - time-domain reflectometer, 600–601
 - TrueCrypt, 474
 - trunk port, 297–298
 - Trusted Platform Module (TPM), 473–474
 - twisted-pair (TP) cable
 - See also* unshielded twisted-pair (UTP) cable
 - plant components, 171–172
 - shielded, 168, 170
- ## U
- UDP (User Datagram Protocol), 222, 271
 - uninterruptible power supply (UPS), 447, 464
 - universal groups, 404
 - Universal Serial Bus (USB), 316–317
 - unshielded twisted pair (UTP) cable, 121–122
 - categories, 122, 167–170
 - installing, 176–178
 - limitations of, 184
 - managing, 172–176
 - terminating, 181–182
 - updates, as problem-solving tools, 587
 - upgrades
 - stand-alone to networked computer, 15–17
 - troubleshooting, 603
 - uplink port, 66, 75–77
 - UPS (uninterruptible power supply), 447, 464
 - USB (Universal Serial Bus), 316–317
 - user accounts
 - account management, 358–359
 - Linux, 409–414
 - managing, 42, 396–414
 - user profile, 407–409
 - Windows, 397–409
 - user authorization. *See* authorization
 - User Datagram Protocol (UDP), 222–224
 - user interface, 12
 - user profile, 407–409
 - UTP (unshielded twisted-pair) cable, 121–122
 - categories, 122, 167–170
 - installing, 176–178
 - limitations of, 184
 - managing, 172–176
 - terminating, 181–183
- ## V
- vampire tap, 631
 - VirtualBox, 369, 374–375
 - virtual circuits, 541–542
 - virtualization, 365–380
 - bare-metal, 367, 376–380
 - hosted, 367–369
 - virtualization software, 651
 - See also* Virtual PC; VMware
 - Microsoft Hyper-V, 378, 379, 686–688
 - summary, 375–376
 - virtual local area networks (VLANs), 296–299
 - virtual machine (VM), 367, 654–657, 664
 - See also* VMware
 - configuring hardware for, 670–671
 - Virtual PC (Microsoft), 369, 373–374, 652–659
 - configuring networking and hardware, 657–659
 - downloading and installing, 653, 654
 - Guest OS, 654–657
 - host key, 659
 - requirements for, 652
 - virtual machine, 654–657
 - virtual private networks (VPNs)
 - benefits, 476–477
 - in Linux and other environments, 476
 - securing communication with, 474–477
 - small-business network remote access, 517, 524–525
 - WANs and, 546–547
 - in Windows, 475–476, 550–551
 - virus protection, 483, 484
 - VLANs (virtual local area networks), 296–299
 - VMware, 369–376
 - Player, 369, 372–373
 - summary, 375–376
 - VirtualBox, 369, 374–375
 - vSphere, 378, 379–380
 - Workstation, 369–372, 685–686
 - VMware Server, 672–684
 - configuring hardware options, 681–682
 - configuring network options, 680–681
 - creating virtual machine and guest OS, 676–680
 - downloading and installing, 674–676
 - Guest OSs supported, 672
 - host OSs supported, 673
 - installing VMware tools, 682–684
 - requirements for, 673
 - volume mount points, 417
 - volumes, 414
 - VPNs. *See* virtual private networks (VPNs)

V-series standards, 622–623

vSphere, 378, 379–380

W

WANs (wide area networks),
28, 29

circuit-switched, 537

connection methods, 536–548

CSU/DSU, 536

customer equipment, 547

fundamentals, 534–536

Internet *vs.* VPNs, 546

modems, 535–536

packet-switched, 541

provider equipment, 548

routers, 536

troubleshooting, 648

virtual circuits, 541–542

WAN devices, 534–536

wardrivers, 486, 524

W3C (World Wide Web Consortium),
627

Web-based networks, remote access,
552–553

Web servers, 46

WEP (Wired Equivalent Privacy), 487

Whois query, 489–490

wide area networks. *See* WANs (wide
area networks)

Wi-Fi Protected Access (WPA)
standard, 487

Wi-Fi (Wireless Fidelity), 137–144

802.11, 137–138

access method and operation,
140–141

communication channels, 138–139

configuring ad hoc, 141–144

modes, 138

security, 139–140

WiMAX (Worldwide Interoperability
for Microwave Access), 149

Windows

Active Directory, 44, 358–359

Backup and Restore, 443–446

DHCP management console, 362

dial-up connection, 553–555

disk management, 421–423

Event Viewer in, 435–436

file shares, 428–431, 472, 510

file system, 334–337

firewall, 479–481

groups (Client OS), 401–403,
406–407

hardware requirements for, 46–47

HomeGroup Network, 505–511

logon restrictions, 470–471

Microsoft Hyper-V, 686–688

Microsoft Virtual PC, 373–374

NIC drivers in, 82–83

passwords, 399–400, 466–469

Performance Monitor, 437–440

Remote Assistance, 525

repair and recovery in, 606–608

Resource Manager, 442

Server 2008 installation, 380–384

shared printers in, 428, 432–434

Small-Business Server, 519

system repair/recovery, 606–608

Task Manager, 342–344

user accounts, 397–409

viewing TCP/IP layers in, 210–212

Virtual PC, 373–374

VPNs, 475, 555

Web servers, 46

Windows Internet Naming Service
(WINS), protocol selection, 363

Wired Equivalent Privacy (WEP), 487

wired networking, 164

See also cable

small-business network, 512

wireless access points (APs), 310–314

advanced settings, 313–314

bandwidth and, 78–79

basic operation, 77–78

basic settings, 310–312

connecting to, 85–86

security options, 312–313

small-business networks, 514, 515

wireless bridges, 116, 194

Wireless Fidelity. *See* Wi-Fi (Wireless
Fidelity)

wireless LANs, 30, 189–194

access points, 77–79, 85–86

components, 190

extended technologies, 190, 194

infrared technologies, 191

laser-based technologies, 191

narrowband radio technologies,
192–193

spread-spectrum technologies,
193–194

transmission, 190–191

wireless networking, 188–197

benefits of, 188–189

extended LAN technologies, 190, 194

LAN components, 190

LAN transmission, 190–191

microwave, 195–196

security, 486–487, 524

small-business networks,
514–515, 524–525

spread-spectrum LAN technologies,
193–194

types of networks, 189–190

- wireless NICs (network interface cards), 84–86
 - wireless personal area networks (WPANs), 47–48
 - Wireshark, 69–75, 598
 - downloading and installing, 69–72
 - with hub, 72–73
 - with switch, 74–75
 - as traffic monitor, 594–595
 - wiring, access to, 464
 - work area, 174, 175
 - workgroup model, 34–36, 383
 - working groups, standards and, 619–620
 - workstations, documentation of, 570
 - Worldwide Interoperability for Microwave Access (WiMAX), 149
 - World Wide Web Consortium (W3C), 627
 - World Wide Web (WWW), 46
 - See also* Internet
 - HTTP protocol, 226
 - as problem-solving tool, 586–588
 - worm, virus, and rootkit protection, 483–484
 - worms, 483
 - WPANs (wireless personal area networks), 47–48
 - WPA (Wi-Fi Protected Access) standard, 487
 - WSRM (Windows System Resource Manager), 442
 - WWW (World Wide Web)
 - See also* Internet
 - HTTP protocol, 226
 - as problem-solving tool, 586–588
- ## X
- XenServer (Citrix), 378, 380
 - X.25 networks, 542
 - X-series standards, 623
 - X-Window System, 626

This page intentionally left blank

This book is intended to be sold with a CD-ROM. If this book does not contain a CD-ROM, you are not getting the full value of your purchase.

If the disk/CD in this book is missing or if the package containing them has been opened, this book is not returnable. By opening and breaking the seal on this package, you are agreeing to be bound by the following agreement:

The software included with this product may be copyrighted, in which case all rights are reserved by the respective copyright holder. You are licensed to use software copyrighted by the Publisher and its licensor on a single computer. You may copy and/or modify the software as needed to facilitate your use of it on a single computer. Making copies of the software for any other purpose is a violation of the United States copyright laws.

This software is sold as is without warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. Neither the publisher nor its dealers or distributors assume any liability for any alleged or actual damages arising from the use of this program. (Some states do not allow for the excusing of implied warranties, so the exclusion may not apply to you.)