

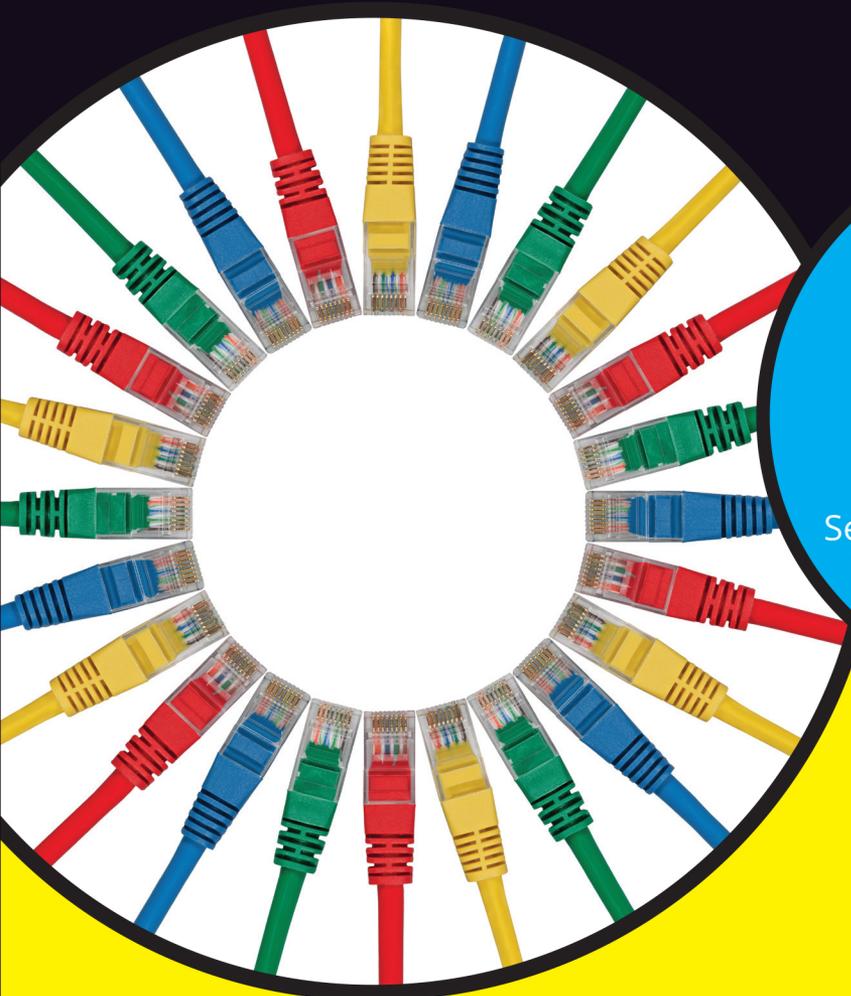
LEARNING MADE EASY



11th Edition

Networking

for
dummies[®]
A Wiley Brand



Build a wired or
wireless network

Secure and optimize
your network

Set up a server and manage
Windows® User Accounts

Doug Lowe

*Bestselling author of Java All-in-One
For Dummies*



Networking

for
dummies[®]
A Wiley Brand

11th Edition

by **Doug Lowe**

for
dummies[®]
A Wiley Brand

Networking For Dummies® , 11th Edition

Published by: **John Wiley & Sons, Inc.**, 111 River Street, Hoboken, NJ 07030-5774, www.wiley.com

Copyright © 2016 by John Wiley & Sons, Inc., Hoboken, New Jersey

Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc., and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHOR HAVE USED THEIR BEST EFFORTS IN PREPARING THIS BOOK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS BOOK AND SPECIFICALLY DISCLAIM ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES OR WRITTEN SALES MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A PROFESSIONAL WHERE APPROPRIATE. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For technical support, please visit www.wiley.com/techsupport.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2016938157

ISBN 978-1-119-25776-9 (pbk); ISBN 978-1-119-25777-6 (ebk); ISBN 978-1-119-25779-0 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents at a Glance

Introduction	1
Part 1: Getting Started with Networking	5
CHAPTER 1: Let's Network!	7
CHAPTER 2: Life on the Network	19
CHAPTER 3: More Ways to Use Your Network	39
Part 2: Setting Up a Network	55
CHAPTER 4: Planning a Network	57
CHAPTER 5: Dealing with TCP/IP	69
CHAPTER 6: Oh, What a Tangled Web We Weave: Cables, Switches, and Routers . . .	95
CHAPTER 7: Configuring Windows Clients	113
CHAPTER 8: Connecting Your Network to the Internet	123
CHAPTER 9: Setting Up a Wireless Network	131
CHAPTER 10: Virtual Networking	151
Part 3: Working with Servers	177
CHAPTER 11: Setting Up a Server	179
CHAPTER 12: Managing Windows User Accounts	191
CHAPTER 13: Managing Network Storage	207
CHAPTER 14: Managing Exchange Server 2016	223
CHAPTER 15: Creating an Intranet	237
Part 4: Managing and Protecting Your Network	251
CHAPTER 16: Welcome to Network Management	253
CHAPTER 17: Solving Network Problems	263
CHAPTER 18: Backing Up Your Data	281
CHAPTER 19: Securing Your Network	295
CHAPTER 20: Hardening Your Network	311
CHAPTER 21: Network Performance Anxiety	323
Part 5: More Ways to Network	335
CHAPTER 22: Life in Cloud City	337
CHAPTER 23: Managing Mobile Devices	347
CHAPTER 24: Connecting from Home	361

Part 6: Networking Beyond Windows	369
CHAPTER 25: Networking with Linux.	371
CHAPTER 26: Mac Networking	393
Part 7: The Part of Tens	403
CHAPTER 27: Ten Networking Commandments	405
CHAPTER 28: Ten Big Network Mistakes.	409
CHAPTER 29: Ten Things You Should Keep in Your Closet	417
Index	421

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	3
Beyond the Book	3
Where to Go from Here	4
PART 1: GETTING STARTED WITH NETWORKING	5
CHAPTER 1: Let's Network!	7
Defining a Network	8
Why Bother with a Network?	10
Sharing files	10
Sharing resources	11
Sharing programs	11
Servers and Clients	12
Dedicated Servers and Peers	13
What Makes a Network Tick?	14
It's Not a Personal Computer Anymore!	15
The Network Administrator	17
What Have They Got That You Don't Got?	18
CHAPTER 2: Life on the Network	19
Distinguishing between Local Resources and Network Resources	20
What's in a Name?	20
Logging On to the Network	22
Understanding Shared Folders	24
Four Good Uses for a Shared Folder	25
Store files that everybody needs	25
Store your own files	26
Make a temporary resting place for files on their way to other users	26
Back up your local hard drive	27
Oh, the Network Places You'll Go	27
Mapping Network Drives	29
Using a Network Printer	32
Adding a network printer	33
Printing to a network printer	35
Playing with the print queue	35
Logging Off the Network	37

CHAPTER 3: More Ways to Use Your Network	39
Sharing Your Stuff	39
Enabling File and Printer Sharing	40
Sharing a Folder	41
Using the Public Folder	43
Sharing a Printer	44
Using Microsoft Office on a Network	46
Accessing network files	46
Using workgroup templates	47
Networking an Access database	50
Working with Offline Files	51
PART 2: SETTING UP A NETWORK	55
CHAPTER 4: Planning a Network	57
Making a Network Plan	57
Being Purposeful	58
Taking Stock	59
What you need to know	59
Programs that gather information for you	61
To Dedicate or Not to Dedicate: That Is the Question	62
File servers	63
Print servers	63
Web servers	64
Mail servers	64
Database servers	64
Application servers	65
License servers	65
Choosing a Server Operating System	65
Planning the Infrastructure	65
Drawing Diagrams	66
CHAPTER 5: Dealing with TCP/IP	69
Understanding Binary	69
Counting by ones	70
Doing the logic thing	71
Introducing IP Addresses	72
Networks and hosts	72
The dotted-decimal dance	73
Classifying IP Addresses	73
Class A addresses	74
Class B addresses	75
Class C addresses	76

Subnetting	77
Subnets	77
Subnet masks	78
The great subnet roundup	80
Private and public addresses	80
Understanding Network Address Translation	81
Configuring Your Network for DHCP	82
Understanding DHCP	82
DHCP servers	83
Understanding scopes	84
Feeling excluded?	85
Reservations suggested	85
How long to lease?	86
Managing a Windows Server 2016 DHCP Server	87
Configuring a Windows DHCP Client	88
Using DNS	89
Domains and domain names	89
Fully qualified domain names	91
Working with the Windows DNS Server	92
Configuring a Windows DNS Client	93
CHAPTER 6: Oh, What a Tangled Web We Weave: Cables, Switches, and Routers	95
What Is Ethernet?	96
All about Cable	98
Cable categories	99
What's with the pairs?	101
To shield or not to shield	101
When to use plenum cable	101
Sometimes solid, sometimes stranded	102
Installation guidelines	102
The tools you need	104
Pinouts for twisted-pair cables	105
RJ-45 connectors	106
Crossover cables	107
Wall jacks and patch panels	108
Working with Switches	109
Daisy-Chaining Switches	110
Using a Router	111
CHAPTER 7: Configuring Windows Clients	113
Configuring Network Connections	113
Joining a Domain	120

CHAPTER 8:	Connecting Your Network to the Internet	123
	Connecting to the Internet.	124
	Connecting with cable or DSL	124
	Connecting with high-speed private lines	125
	Sharing an Internet connection.	126
	Securing Your Connection with a Firewall	127
	Using a firewall.	127
	The built-in Windows firewall	129
CHAPTER 9:	Setting Up a Wireless Network	131
	Diving into Wireless Networking.	132
	A Little High School Electronics	133
	Waves and frequencies.	133
	Wavelength and antennas	135
	Spectrums and the FCC	135
	Eight-Oh-Two-Dot-Eleventy Something?: Understanding Wireless Standards	137
	Home on the Range	138
	Using Wireless Network Adapters.	139
	Setting Wireless Access Points.	139
	Infrastructure mode	140
	Multifunction WAPs.	141
	Roaming Capabilities	142
	Wireless bridging.	142
	Ad-hoc networks	142
	Configuring a Wireless Access Point.	143
	Basic configuration options	143
	DHCP configuration.	144
	Connecting to a Wireless Network	145
	Paying Attention to Wireless Network Security.	147
CHAPTER 10:	Virtual Networking	151
	Understanding Virtualization.	151
	Understanding Hypervisors	153
	Understanding Virtual Disks	155
	Understanding Network Virtualization.	157
	Looking at the Benefits of Virtualization	158
	Introducing Hyper-V	160
	Understanding the Hyper-V hypervisor	160
	Understanding virtual disks	161
	Enabling Hyper-V.	162
	Getting Familiar with Hyper-V	163
	Creating a Virtual Switch	164
	Creating a Virtual Disk	166
	Creating a Virtual Machine.	170
	Installing an Operating System	174

PART 3: WORKING WITH SERVERS	177
CHAPTER 11: Setting Up a Server	179
Server Operating System Features	179
Network support	180
File-sharing services	180
Multitasking	180
Directory services	181
Security services	182
The Many Ways to Install a Network Operating System	183
Full install versus upgrade	183
Installing over the network	184
Gathering Your Stuff	184
A capable server computer	184
The server OS	185
Other software	186
A working Internet connection	186
A good book	186
Making Informed Decisions	186
Final Preparations	188
Installing a Server Operating System	188
Phase 1: Collecting Information	189
Phase 2: Installing Windows	189
Configuring Your Server	190
CHAPTER 12: Managing Windows User Accounts	191
Understanding Windows User Accounts	191
Local accounts versus domain accounts	192
User account properties	192
Creating a New User	193
Setting User Properties	196
Changing the user's contact information	197
Setting account options	197
Specifying logon hours	198
Restricting access to certain computers	199
Setting the user's profile information	200
Resetting User Passwords	201
Disabling and Enabling User Accounts	202
Deleting a User	202
Working with Groups	203
Creating a group	203
Adding a member to a group	204
Creating a Logon Script	206

CHAPTER 13: Managing Network Storage	207
Understanding Network Storage	207
File servers	207
Storage appliances	208
Understanding Permissions	208
Understanding Shares	210
Managing Your File Server	211
Using the New Share Wizard	212
Sharing a folder without the wizard	217
Granting permissions	219
CHAPTER 14: Managing Exchange Server 2016	223
Creating a Mailbox	224
Managing Mailboxes	226
Enabling mailbox features	226
Creating a forwarder	228
Setting mailbox storage limits	229
Configuring Outlook for Exchange	233
CHAPTER 15: Creating an Intranet	237
Defining an Intranet	237
Identifying Intranet Uses	238
Setting Up an Intranet	239
Setting Up an IIS Web Server	240
Understanding the Default Website	243
Creating Websites	246
PART 4: MANAGING AND PROTECTING YOUR NETWORK	251
CHAPTER 16: Welcome to Network Management	253
What a Network Administrator Does	254
Choosing the Part-Time Administrator	255
The Three “Ups” of Network Management	256
Managing Network Users	257
Acquiring Software Tools for Network Administrators	258
Building a Library	259
Pursuing Certification	260
Helpful Bluffs and Excuses	261
CHAPTER 17: Solving Network Problems	263
When Bad Things Happen to Good Computers	264
Fixing Dead Computers	265
Ways to Check a Network Connection	266
A Bunch of Error Messages Just Flew By!	267
Double-Checking Your Network Settings	268

Using a Windows Troubleshooter	268
Time to Experiment	270
Who's on First?	270
Restarting a Client Computer	271
Booting in Safe Mode	273
Using System Restore	273
Restarting Network Services	275
Restarting a Network Server	276
Looking at Event Logs	277
Documenting Your Trials and Tribulations	278
CHAPTER 18: Backing Up Your Data	281
Backing Up Your Data	281
Choosing Where to Back Up Your Data	282
Backing Up to Tape	282
Understanding Backup Software	283
Comparing Types of Backups	284
Normal backups	285
Copy backups	286
Daily backups	286
Incremental backups	287
Differential backups	288
Choosing between Local and Network Backups	288
Deciding How Many Sets of Backups to Keep	290
Verifying Tape Reliability	291
Keeping Backup Equipment Clean and Reliable	292
Setting Backup Security	293
CHAPTER 19: Securing Your Network	295
Do You Need Security?	296
Two Approaches to Security	297
Physical Security: Locking Your Doors	298
Securing User Accounts	299
Obfuscating your usernames	300
Using passwords wisely	300
Generating passwords For Dummies	301
Secure the Administrator account	302
Managing User Security	303
User accounts	303
Built-in accounts	304
User rights	305
Permissions (who gets what)	306
Group therapy	307
User profiles	308
Logon scripts	308
Securing Your Users	309

CHAPTER 20: Hardening Your Network	311
Firewalls	312
The Many Types of Firewalls	313
Packet filtering	313
Stateful packet inspection (SPI)	315
Circuit-level gateway	316
Application gateway	316
The Built-In Windows Firewall	317
Virus Protection	317
What is a virus?	317
Antivirus programs	319
Safe computing	320
Patching Things Up	320
CHAPTER 21: Network Performance Anxiety	323
Why Administrators Hate Performance Problems	324
What Exactly Is a Bottleneck?	325
The Five Most Common Network Bottlenecks	326
The hardware inside your servers	326
The server's configuration options	327
Servers that do too much	328
The network infrastructure	328
Malfunctioning components	329
Tune Your Network the Compulsive Way	329
Monitoring Network Performance	330
More Performance Tips	332
PART 5: MORE WAYS TO NETWORK	335
CHAPTER 22: Life in Cloud City	337
Introducing Cloud Computing	337
Looking at the Benefits of Cloud Computing	338
Detailing the Drawbacks of Cloud Computing	340
Examining Three Basic Kinds of Cloud Services	341
Applications	341
Platforms	342
Infrastructure	342
Public Clouds versus Private Clouds	343
Introducing Some of the Major Cloud Providers	344
Amazon	344
Google	344
Microsoft	345
Getting Into the Cloud	345

CHAPTER 23: Managing Mobile Devices	347
The Many Types of Mobile Devices	348
Considering Security for Mobile Devices	349
Managing iOS Devices	350
Understanding the iPhone	350
Understanding the iPad	351
Integrating iOS devices with Exchange	351
Configuring an iOS device for Exchange email	353
Managing Android Devices	357
Looking at the Android OS	357
Perusing Android's core applications	358
Integrating Android with Exchange	359
CHAPTER 24: Connecting from Home	361
Using Outlook Web App	362
Using a Virtual Private Network	364
Looking at VPN security	365
Understanding VPN servers and clients	366
PART 6: NETWORKING BEYOND WINDOWS	369
CHAPTER 25: Networking with Linux	371
Comparing Linux with Windows	372
Choosing a Linux Distribution	374
Installing Linux	375
On Again, Off Again	377
Logging on	377
Logging off	378
Shutting down	378
Using GNOME	378
Getting to a Command Shell	379
Enabling the SUDO Command	380
Managing User Accounts	383
Network Configuration	384
Using the Network Configuration program	385
Restarting your network	386
Doing the Samba Dance	386
Understanding Samba	387
Installing Samba	388
Starting and stopping Samba	388
Using the Samba Server Configuration tool	389
CHAPTER 26: Mac Networking	393
Basic Mac Network Settings	394
Joining a Domain	397
Connecting to a Share	400

PART 7: THE PART OF TENS	403
CHAPTER 27: Ten Networking Commandments	405
I. Thou Shalt Back Up Thy Hard Drive Religiously.....	405
II. Thou Shalt Protect Thy Network from Infidels.....	406
III. Thou Shalt Keepeth Thy Network Drive Pure and Cleanse It of Old Files.....	406
IV. Thou Shalt Not Tinker with Thine Network Configuration Unless Thou Knowest What Thou Art Doing.....	406
V. Thou Shalt Not Covet Thy Neighbor's Network.....	407
VI. Thou Shalt Schedule Downtime before Working upon Thy Network.....	407
VII. Thou Shalt Keep an Adequate Supply of Spare Parts.....	407
VIII. Thou Shalt Not Steal Thy Neighbor's Program without a License...	408
IX. Thou Shalt Train Thy Users in the Ways of the Network.....	408
X. Thou Shalt Write Down Thy Network Configuration upon Tablets of Stone.....	408
CHAPTER 28: Ten Big Network Mistakes	409
Skimping on Hardware.....	409
Turning Off or Restarting a Server Computer While Users Are Logged On.....	410
Deleting Important Files on the Server.....	411
Copying a File from the Server, Changing It, and Then Copying It Back.....	411
Sending Something to the Printer Again Just Because It Didn't Print the First Time.....	412
Assuming That the Server Is Safely Backed Up.....	412
Connecting to the Internet without Considering Security Issues...	412
Plugging In a Wireless Access Point without Asking.....	413
Thinking You Can't Work Just Because the Network Is Down.....	413
Running Out of Space on a Server.....	414
Always Blaming the Network.....	415
CHAPTER 29: Ten Things You Should Keep in Your Closet	417
Duct Tape.....	417
Tools.....	418
Patch Cables.....	418
Cable Ties.....	418
Twinkies.....	418
Replacement Parts.....	419
Cheap Network Switches.....	419
The Complete Documentation of the Network on Tablets of Stone...	420
The Network Manuals and Disks.....	420
Ten Copies of This Book.....	420
INDEX	421

Introduction

Welcome to the eleventh edition of *Networking For Dummies*, the book that's written especially for people who have this nagging feeling in the back of their minds that they should network their computers but haven't a clue about how to start or where to begin.

Do you often copy a spreadsheet to a flash drive just so you can give it to someone else in your office? Are you frustrated because you can't use the fancy color laser printer that's on the financial secretary's computer? Do you wait in line to use the computer that has the customer database? You need a network!

Or maybe you already have a network, but you have just one problem: Someone promised that a network would make your life easier, but it's instead turned your computing life upside down. Just when you had this computer thing figured out, someone popped into your office, hooked up a cable, and said, "Happy networking!" Makes you want to scream.

Regardless, you've found the right book. Help is here, within these humble pages.

This book talks about networks in everyday (and often irreverent) terms. The language is friendly; you don't need a graduate education to get through it. And the occasional potshot helps unseat the hallowed and sacred traditions of networking, bringing just a bit of fun to an otherwise dry subject. The goal is to bring the lofty precepts of networking down to earth, where you can touch them and squeeze them and say, "What's the big deal? I can do this!"

About This Book

This isn't the kind of book you pick up and read from start to finish, as if it were a cheap novel. If I ever see you reading it at the beach, I'll kick sand in your face. This book is more like a reference, the kind of book you can pick up, turn to just about any page, and start reading. Each chapter covers a specific aspect of networking, such as printing from the network, hooking up network cables, or setting up security so that bad guys can't break in. Just turn to the chapter you're interested in and start reading.

Each chapter is divided into self-contained chunks, all related to the major theme of the chapter. For example, the chapter on hooking up the network cable contains nuggets like these:

- »» What is Ethernet?
- »» All about cable
- »» To shield or not to shield
- »» Wall jacks and patch panels
- »» Switches

You don't have to memorize anything in this book. It's a need-to-know book: You pick it up when you need to know something. Need to know what 100BaseT is? Pick up the book. Need to know how to create good passwords? Pick up the book. Otherwise, put it down and get on with your life.

Feel free to skip the sidebars that appear throughout the book; these shaded gray boxes contain interesting info that isn't essential to your understanding of the subject at hand. The same goes for any text I mark with the Technical Stuff icon.

If you need to type something, you see the text you need to type like this: **Type this stuff**. In this example, you type **Type this stuff** at the keyboard and then press Enter. An explanation usually follows, just in case you're scratching your head and grunting, "Huh?"

Within this book, you may note that some web addresses break across two lines of text. If you're reading this book in print and want to visit one of these web pages, simply key in the web address exactly as it's noted in the text, pretending as though the line break doesn't exist. If you're reading this as an e-book, you've got it easy — just click the web address to be taken directly to the web page.

Foolish Assumptions

I'm making only two assumptions about who you are: You're someone who works with a PC, and you either have a network or you're thinking about getting one. I hope that you know (and are on speaking terms with) someone who knows more about computers than you do. My goal is to decrease your reliance on that person, but don't throw away his phone number yet.

Is this book useful for Macintosh users? Absolutely. Although the bulk of this book is devoted to showing you how to link Windows-based computers to form a network, you can find information about how to network Macintosh computers as well.

Windows 10? Gotcha covered. You'll find plenty of information about how to network with the latest and greatest Microsoft desktop operating system.

Windows Server 2016? No worries. You'll find plenty of information about the newest version of Microsoft's server operating system.

Icons Used in This Book

Those nifty little pictures in the margin aren't there just to pretty up the place. They also have practical functions.



TECHNICAL
STUFF

Hold it — technical details lurk just around the corner. Read on only if you have a pocket protector.



TIP

Pay special attention to this icon; it lets you know that some particularly useful tidbit is at hand — perhaps a shortcut or a little-used command that pays off big.



REMEMBER

Did I tell you about the memory course I took?



WARNING

Danger, Will Robinson! This icon highlights information that may help you avoid disaster.

Beyond the Book

In addition to the material in the print or e-book you're reading right now, this product also comes with some access-anywhere goodies on the web. Check out the free Cheat Sheet for links to useful websites for networking information, private

IP address ranges for networks, and more. To get this Cheat Sheet, simply go to www.dummies.com and type **Networking For Dummies Cheat Sheet** in the Search box.

Where to Go from Here

Yes, you can get there from here. With this book in hand, you're ready to plow right through the rugged networking terrain. Browse through the Table of Contents and decide where you want to start. Be bold! Be courageous! Be adventurous! Above all, have fun!

1

Getting Started with Networking

IN THIS PART . . .

Find out what a network is and what you can do with one.

Compare server and client computers.

Access network resources such as shared storage and network printers.

Use Microsoft Office and other software on a network.

IN THIS CHAPTER

Getting a handle on networks

Considering why networking is useful (and is everywhere)

Telling the difference between servers and clients

Looking under the hood at the network operating system

Asking “How does it work when a network works if a network works for me?” (Say what?)

Assessing how networks change computing life

Identifying (and offering sympathy to) the network administrator

Comparing servers to clients: What have they got that you don't got?

Chapter 1

Let's Network!

Computer networks get a bad rap in the movies. In the 1980s, the *Terminator* movies featured Skynet, a computer network that becomes self-aware (a computer network of the future), takes over the planet, builds deadly terminator robots, and sends them back through time to kill everyone unfortunate enough to have the name Sarah Connor. In the *Matrix* movies, a vast and powerful computer network enslaves humans and keeps them trapped in a simulation of the real world. And in the 2015 blockbuster *Spectre*, James Bond goes rogue (again) to prevent the Evil Genius Ernst Blofeld from taking over the world (again) by linking the computer systems of all the world's intelligence agencies together to form a single all-powerful evil network that spies on everybody.

Fear not. These bad networks exist only in the dreams of science fiction writers. Real-world networks are much more calm and predictable. Although sophisticated networks do seem to know a lot about you, they don't think for themselves and they don't evolve into self-awareness. And although they can gather a sometimes disturbing amount of information about you, they aren't trying to kill you, even if your name is Sarah Connor.

Now that you're over your fear of networks, you're ready to breeze through this chapter. It's a gentle, even superficial, introduction to computer networks, with a slant toward the concepts that can help you use a computer that's attached to a network. This chapter goes easy on the details; the detailed and boring stuff comes later.

Defining a Network

A *network* is nothing more than two or more computers connected by a cable or by a wireless radio connection so that they can exchange information.

Of course, computers can exchange information in ways other than networks. Most of us have used what computer nerds call the *sneakernet*. That's where you copy a file to a flash drive or other portable storage device and then walk the data over to someone else's computer. (The term *sneakernet* is typical of computer nerds' feeble attempts at humor.)

The whole problem with the sneakernet is that it's slow, and it wears a trail in your carpet. One day, some penny-pinching computer geeks discovered that connecting computers with cables was cheaper than replacing the carpet every six months. Thus, the modern computer network was born.

You can create a simple computer network by hooking together all the computers in your office with cables and using the computer's *network interface* (an electronic circuit that resides inside your computer and has a special jack on the computer's backside). Then you tweak a few simple settings in the computer's operating system (OS) software, and *voilà!* You have a working network. That's all there is to it.

If you don't want to mess with cables, you can create a wireless network instead. In a wireless network, the computers use wireless network adapters that communicate via radio signals. All modern laptop computers have built-in wireless network adapters, as do most desktop computers. (If yours doesn't, you can purchase a separate wireless network adapter that plugs into one of the computer's USB ports.)

Figure 1-1 shows a typical network with four computers. You can see that all four computers are connected by a network cable to a central network device: the *switch*. You can also see that Ward's computer has a fancy laser printer attached to it. Because of the network, June, Wally, and the Beaver can also use this laser printer. (Also, you can see that the Beaver stuck yesterday's bubble gum to the back of his computer. Although the bubble gum isn't recommended, it shouldn't adversely affect the network.)

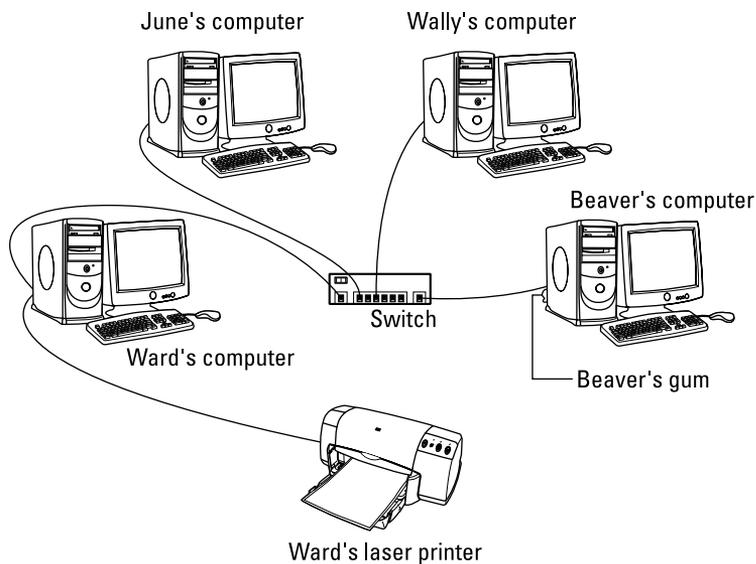


FIGURE 1-1:
A typical network.

Computer networking has its own strange vocabulary. Although you don't have to know every esoteric networking term, it helps to be acquainted with a few of the basic buzzwords:

» **LAN:** Networks are often called LANs, short for *local area network*.

LAN is the first *TLA* — or *three-letter acronym* — of this book. You don't really need to remember it or any of the many TLAs that follow. In fact, the only three-letter acronym you need to remember is TLA. You might guess that the acronym for *four-letter acronym* is *FLA*. Wrong! A four-letter acronym is an *ETLA*, which stands for *extended three-letter acronym*. After all, it just wouldn't be right if the acronym for *four-letter acronym* had only three letters.

» **On the network:** Every computer connected to the network is said to be "on the network." The technical term (which you can forget) for a computer that's on the network is a *node*.



- » **Online, offline:** When a computer is turned on and can access the network, the computer is *online*. When a computer can't access the network, it's *offline*. A computer can be offline for several reasons. The computer can be turned off, the user may have disabled the network connection, the computer may be broken, the cable that connects it to the network can be unplugged, or a wad of gum can be jammed into the disk drive.
- » **Up, down:** When a computer is turned on and working properly, it's *up*. When a computer is turned off, broken, or being serviced, it's *down*. Turning off a computer is sometimes called *taking it down*. Turning it back on is sometimes called *bringing it up*.
- » **Local, remote:** A resource such as a disk drive is *local* if it resides in your computer. It's *remote* if it resides in another computer somewhere else on your network.
- » **Internet:** The *Internet* is a huge amalgamation of computer networks strewn about the entire planet. Networking the computers in your home or office so that they can share information with one another and connecting your computer to the worldwide Internet are two separate but related tasks.

Why Bother with a Network?

Frankly, computer networks are a bit of a pain to set up. So why bother? Because the benefits of having a network outweigh the difficulties of setting up one.

You don't have to be a PhD to understand the benefits of networking. In fact, you learned everything you need to know in kindergarten: Networks are all about sharing. Specifically, networks are about sharing three things: files, resources, and programs.

Sharing files

Networks enable you to share information with other computers on the network. Depending on how you set up your network, you can share files with your network friends in several different ways. You can send a file from your computer directly to a friend's computer by attaching the file to an email message and then mailing it. Or you can let your friend access your computer over the network so that your friend can retrieve the file directly from your hard drive. Yet another method is to copy the file to a disk on another computer and then tell your friend where you put the file so that your friend can retrieve it later. One way or the other, the data travels to your friend's computer over the network cable and not on a CD or DVD or flash drive, as it would in a sneakernet.

Sharing resources

You can set up certain computer resources — such as hard drives or printers — so that all computers on the network can access them. For example, the laser printer attached to Ward’s computer in Figure 1-1 is a *shared resource*, which means that anyone on the network can use it. Without the network, June, Wally, and the Beaver would have to buy their own laser printers.

Hard drives can be shared resources, too. In fact, you must set up a hard drive as a shared resource to share files with other users. Suppose that Wally wants to share a file with the Beaver, and a shared hard drive has been set up on June’s computer. All Wally has to do is copy his file to the shared hard drive in June’s computer and tell the Beaver where he put it. Then, when the Beaver gets around to it, he can copy the file from June’s computer to his own (unless, of course, that hooligan Eddie Haskell deletes the file first).



TIP

You can share other resources, too, such as an Internet connection. In fact, sharing an Internet connection is one of the main reasons why many networks are created.

Sharing programs

Rather than keep separate copies of programs on each person’s computer, putting programs on a drive that everyone shares is sometimes best. For example, if ten computer users all use a particular program, you can purchase and install ten copies of the program, one for each computer. Or you can purchase a ten-user license for the program and then install just one copy of the program on a shared drive. Each of the ten users can then access the program from the shared hard drive.

In most cases, however, running a shared copy of a program over the network is unacceptably slow. A more common way of using a network to share programs is to copy the program’s installation disks or CDs to a shared network drive. Then you can use that copy to install a separate copy of the program on each user’s local hard drive. For example, Microsoft Office enables you to do this if you purchase a license from Microsoft for each computer on which you install Office.

The advantage of installing Office from a shared network drive is that you don’t have to lug around the installation disks or CDs to each user’s computer. And the system administrator can customize the network installation so that the software is installed the same way on each user’s computer. (However, these benefits are significant only for larger networks. If your network has fewer than about ten computers, you’re probably better off installing the program separately on each computer directly from the installation disks or CDs.)



WARNING

Remember that purchasing a single-user copy of a program and then putting it on a shared network drive — so that everyone on the network can access it — is illegal. If five people use the program, you need to either purchase five copies of the program or purchase a network license that specifically allows five or more users.



TIP

That being said, many software manufacturers sell their software with a concurrent usage license, which means that you can install the software on as many computers as you want, but only a certain number of people can use the software at any given time. Usually, special licensing software that runs on one of the network's server computers keeps track of how many people are currently using the software. This type of license is frequently used with more specialized (and expensive) software, such as accounting systems or computer drafting systems.

Another benefit of networking is that networks enable computer users to communicate with one another over the network. The most obvious way networks allow computer users to communicate is by passing messages back and forth, using email or instant-messaging programs. Networks also offer other ways to communicate: For example, you can hold online meetings over the network. Network users who have inexpensive video cameras (webcams) attached to their computers can have videoconferences. You can even play a friendly game of Hearts over a network — during your lunch break, of course.

Servers and Clients

The network computer that contains the hard drives, printers, and other resources that are shared with other network computers is a *server*. This term comes up repeatedly, so you have to remember it. Write it on the back of your left hand.

Any computer that's not a server is a *client*. You have to remember this term, too. Write it on the back of your right hand.

Only two kinds of computers are on a network: servers and clients. Look at your left hand and then look at your right hand. Don't wash your hands until you memorize these terms.

The distinction between servers and clients in a network has parallels in sociology — in effect, a sort of class distinction between the “haves” and “have-nots” of computer resources:

- » Usually, the most powerful and expensive computers in a network are the servers. There's a good technical reason: All users on the network share the server's resources.

- » The cheaper and less-powerful computers in a network are the clients. *Clients* are the computers used by individual users for everyday work. Because clients' resources don't have to be shared, they don't have to be as fancy.
- » Most networks have more clients than servers. For example, a network with ten clients can probably get by with one server.
- » In many networks, a clean line of demarcation exists between servers and clients. In other words, a computer functions as either a server or a client, not both. For the sake of an efficient network, a server can't become a client, nor can a client become a server.
- » Other (usually smaller) networks can be more evenhanded by allowing any computer in the network to be a server and allowing any computer to be both server and client at the same time.

Dedicated Servers and Peers

In some networks, a server computer is a server computer and nothing else. It's dedicated to the sole task of providing shared resources, such as hard drives and printers, to be accessed by the network client computers. This type of server is a *dedicated server* because it can perform no other task than network services.

Some smaller networks take an alternative approach by enabling any computer on the network to function as both a client and a server. Thus, any computer can share its printers and hard drives with other computers on the network. And while a computer is working as a server, you can still use that same computer for other functions, such as word processing. This type of network is a *peer-to-peer network* because all the computers are thought of as *peers*, or equals.

Here are some points to ponder concerning the differences between dedicated server networks and peer-to-peer networks while you're walking the dog tomorrow morning:

- » Peer-to-peer networking features are built into Windows. Thus, if your computer runs Windows, you don't have to buy any additional software to turn your computer into a server. All you have to do is enable the Windows server features.
- » The network server features that are built into desktop versions of Windows (such as Windows 7 and 8) aren't particularly efficient because these versions of Windows weren't designed primarily to be network servers.



REMEMBER

If you dedicate a computer to the task of being a full-time server, use a special server operating system rather than the standard Windows desktop operating system. A *server operating system* is specially designed to handle networking functions efficiently.

- The most commonly used server operating systems are the server versions of Windows.

As of this writing, the current server version of Windows is Windows Server 2016. However, many companies still use the previous version (Windows Server 2012), and a few even use its predecessor, Windows Server 2008.

- Another popular server operating system is *Linux*. Linux is popular because it is free. However, it requires a more expertise to set up than Windows Server.
- » Many networks are both peer-to-peer *and* dedicated-server networks at the same time. These networks have
- At least one server computer that runs a server operating system such as Windows Server 2016
 - *Client* computers that use the server features of Windows to share their resources with the network
- » Besides being dedicated, your servers should also be sincere.



TIP

What Makes a Network Tick?

To use a network, you don't really have to know much about how it works. Still, you may feel a little bit better about using the network if you realize that it doesn't work by voodoo. A network may seem like magic, but it isn't. The following list describes the inner workings of a typical network:

- » **Network interface:** Inside any computer attached to a network is a special electronic circuit called the *network interface*. The network interface has either an external jack into which you can plug a network cable — or, in the case of a wireless network interface, an antenna.
- » **Network cable:** The network cable physically connects the computers. It plugs into the network interface card (NIC) on the back of your computer.

The type of network cable most commonly used is twisted-pair cable, so named because it consists of several pairs of wires twisted together in a certain way. Twisted-pair cable superficially resembles telephone cable. However,

appearances can be deceiving. Most phone systems are wired using a lower grade of cable that doesn't work for networks.

For the complete lowdown on networking cables, see Chapter 6.

Network cable isn't necessary when wireless networking is used. For more information about wireless networking, see Chapter 9.



TIP

» **Network switch:** Networks built with twisted-pair cabling require one or more switches. A *switch* is a box with a bunch of cable connectors. Each computer on the network is connected by cable to the switch. The switch, in turn, connects all the computers to each other.

In the early days of twisted-pair networking, devices known as *hubs* were used rather than switches. The term *hub* is sometimes used to refer to switches, but true hubs went out of style sometime around the turn of the century.

In networks with just a few computers, the network switch is often combined with another networking device called a *router*. A router is used to connect two networks. Typically, a router is used to connect your network to the Internet. By combining a router and a switch in a single box, you can easily connect several computers to the Internet and to each other.

» **Network software:** Of course, the software makes the network work. To make any network work, a whole bunch of software has to be set up just right. For peer-to-peer networking with Windows, you have to play with the Control Panel to get networking to work. And a network operating system such as Windows Server 2016 requires a substantial amount of tweaking to get it to work just right.



TECHNICAL
STUFF



TECHNICAL
STUFF

It's Not a Personal Computer Anymore!

If I had to choose one point that I want you to remember from this chapter more than anything else, it's this: After you hook up your personal computer (PC) to a network, it's not a "personal" computer anymore. You're now part of a network of computers, and in a way, you've given up one of the key concepts that made PCs so successful in the first place: independence.

I got my start in computers back in the days when mainframe computers ruled the roost. *Mainframe computers* are big, complex machines that used to fill entire rooms and had to be cooled with chilled water. My first computer was a water-cooled Binford Hex Core Model 2000. Argh, argh, argh. (I'm not making up the part about the water. A plumber was often required to install a mainframe computer. In fact, the really big ones were cooled by liquid nitrogen. I *am* making up the part about the Binford Hex Core 2000.)

Mainframe computers required staffs of programmers and operators in white lab coats just to keep them going. The mainframes had to be carefully managed. A whole bureaucracy grew up around managing them.

Mainframe computers used to be the dominant computers in the workplace. Personal computers changed all that: They took the computing power out of the big computer room and put it on the user's desktop, where it belongs. PCs severed the tie to the centralized control of the mainframe computer. With a PC, a user could look at the computer and say, "This is mine — all mine!" Mainframes still exist, but they're not nearly as popular as they once were.

But networks have changed everything all over again. In a way, it's a change back to the mainframe-computer way of thinking: central location, distributed resources. True, the network isn't housed in the basement and doesn't have to be installed by a plumber. But you can no longer think of "your" PC as your own. You're part of a network — and like the mainframe, the network has to be carefully managed.

Here are several ways in which a network robs you of your independence:

- » **You can't just indiscriminately delete files from the network.** They may not be yours.
- » **You're forced to be concerned about network security.** For example, a server computer has to know who you are before it allows you to access its files. So you have to know your user ID and password to access the network. This precaution prevents some 15-year-old kid from hacking his way into your office network by using its Internet connection and stealing all your computer games.
- » **You may have to wait for shared resources.** Just because Wally sends something to Ward's printer doesn't mean that it immediately starts to print. The Beav may have sent a two-hour print job before that. Wally just has to wait.
- » **You may have to wait for access to documents.** You may try to retrieve an Excel spreadsheet file from a network drive, only to discover that someone else is using it. Like Wally, you just have to wait.
- » **You don't have unlimited storage space.** If you copy a 100GB video file to a server's drive, you may get calls later from angry co-workers complaining that no room is left on the server's drive for their important files.
- » **Your files can become infected from viruses given to you by someone over the network.** You may then accidentally infect other network users.
- » **You have to be careful about saving sensitive files on the server.** If you write an angry note about your boss and save it on the server's hard drive, your boss may find the memo and read it.

- » **The server computer must be up and running at all times.** For example, if you turn Ward's computer into a server computer, Ward can't turn his computer off when he's out of the office. If he does, you can't access the files stored on his computer.
- » **If your computer is a server, you can't just turn it off when you're finished using it.** Someone else may be accessing a file on your hard drive or printing on your printer.

The Network Administrator

Because so much can go wrong — even with a simple network — designating one person as network administrator is important. This way, someone is responsible for making sure that the network doesn't fall apart or get out of control.

The network administrator doesn't have to be a technical genius. In fact, some of the best network administrators are complete idiots when it comes to technical stuff. What's important is that the administrator is organized. That person's job is to make sure that plenty of space is available on the file server, that the file server is backed up regularly, that new employees can access the network, among other tasks.

The network administrator's job also includes solving basic problems that the users themselves can't solve — and knowing when to call in an expert when something really bad happens. It's a tough job, but somebody's got to do it. Here are a few tips that might help:

- » Part 4 of this book is devoted entirely to the hapless network administrator. So if you're nominated, read the chapters in that part. If you're lucky enough that someone *else* is nominated, celebrate by buying her a copy of this book.
- » In small companies, picking the network administrator by drawing straws is common. The person who draws the shortest straw loses and becomes administrator.
- » Of course, the network administrator can't be a *complete* technical idiot. I was lying about that. (For those of you in Congress, the word is *testifying*.) I exaggerated to make the point that organizational skills are more important than technical skills. The network administrator needs to know how to do various maintenance tasks. Although this knowledge requires at least a little technical know-how, the organizational skills are more important.

What Have They Got That You Don't Got?

With all this technical stuff to worry about, you may begin to wonder whether you're smart enough to use your computer after it's attached to the network. Let me assure you that you are. If you're smart enough to buy this book because you know that you need a network, you're more than smart enough to use the network after it's put in. You're also smart enough to install and manage a network yourself. It isn't rocket science.

I know people who use networks all the time. They're no smarter than you are, but they do have one thing that you don't have: a certificate. And so, by the powers vested in me by the International Society for the Computer Impaired, I present you with the certificate in Figure 1-2, confirming that you've earned the coveted title Certified Network Dummy, better known as CND. This title is considered much more prestigious in certain circles than the more stodgy CNE or MCSE badges worn by real network experts.

Congratulations, and go in peace.

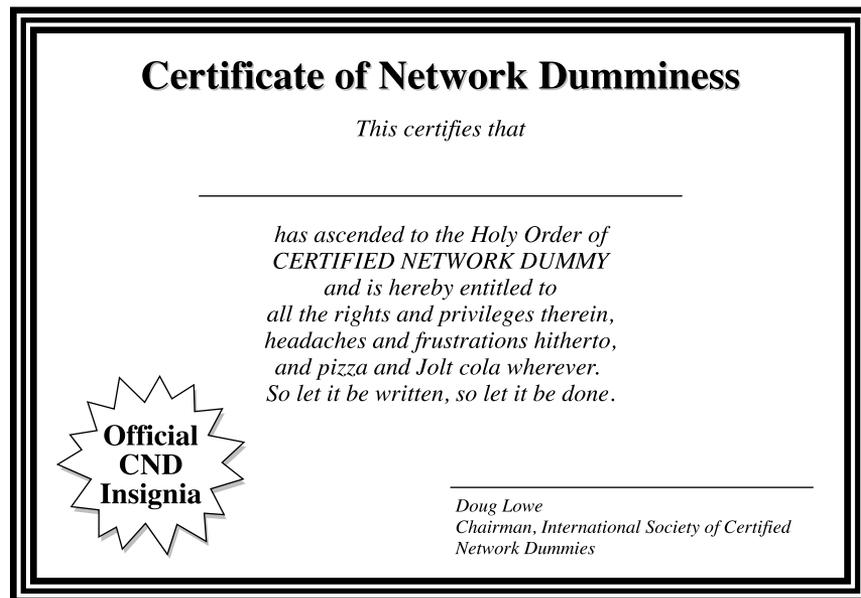


FIGURE 1-2:
Your official CND
certificate.

IN THIS CHAPTER

Using local resources and network resources

Playing the name game

Logging on to a network

Using shared folders

Going places with networks

Mapping your network drives

Using a network printer

Logging off the network

Chapter 2

Life on the Network

After you hook up your PC to a network, it's not an island anymore, separated from the rest of the world like some kind of isolationist fanatic waving a "Don't tread on me" flag. The network connection changes your PC forever. Now your computer is part of a system, connected to other computers on the network. You have to worry about annoying network details, such as using local and shared resources, logging on and accessing network drives, using network printers, logging off, and who knows what else.

Oh, bother.

This chapter brings you up to speed on what living with a computer network is like. Unfortunately, this chapter gets a little technical at times, so you may need your pocket protector.

Distinguishing between Local Resources and Network Resources

In case you don't catch this statement in Chapter 1, one of the most important differences between using an isolated computer and using a network computer lies in the distinction between local resources and network resources. *Local resources* are items — such as hard drives, printers, and CD or DVD drives — that are connected directly to your computer. You can use local resources whether you're connected to the network or not. *Network resources*, on the other hand, are the hard drives, printers, optical drives, and other devices that are connected to the network's server computers. You can use network resources only after your computer is connected to the network.

Whenever you use a computer network, you need to know which resources are local resources (belong to you) and which are network resources (belong to the network). In most networks, your C: drive is a local drive, as is your My Documents folder. If a printer is sitting next to your PC, it's probably a local printer. You can do anything you want with these resources without affecting the network or other users on the network (as long as the local resources aren't shared on the network). Keep these points in mind:

- » You can't tell just by looking at a resource whether it's a local resource or a network resource. The printer that sits right next to your computer is probably your local printer, but then again, it may be a network printer. The same statement is true for hard drives: The hard drive in your PC is probably your own, but it (or part of it) may be shared on the network, thus enabling other users to access it.
- » Because dedicated network servers are full of resources, you may say that they're not only dedicated (and sincere), but also resourceful. (Groan. Sorry. This is yet another in a tireless series of bad computer-nerd puns.)

What's in a Name?

Just about everything on a computer network has a name: The computers themselves have names, the people who use the computers have names, the hard drives and printers that can be shared on the network have names, and the network itself has a name. Knowing all the names used on your network isn't essential, but you do need to know some of them.

Here are some additional details about network names:

- » **Every person who can use the network has a username (sometimes called a *user ID*).** You need to know your username to log on to the network. You also need to know the usernames of your buddies, especially if you want to steal their files or send them nasty notes.

You can find more information about usernames and logging on in the section “Logging On to the Network,” later in this chapter.



WARNING

- » **Letting folks on the network use their first names as their usernames is tempting but not a good idea.** Even in a small office, you eventually run into a conflict. (And what about Mrs. McCave — made famous by Dr. Seuss — who had 23 children and named them all Dave?)

Create a consistent way of creating usernames. For example, you may use your first name plus the first two letters of your last name. Then Wally’s username is `wallyc1`, and Beaver’s is `beaverc1`. Or you may use the first letter of your first name followed by your complete last name. Then Wally’s username is `w1eaver`, and Beaver’s is `bc1eaver`. (In most networks, capitalization doesn’t matter in usernames. Thus, `bc1eaver` is the same as `BC1eaver`.)



TIP

- » **Every computer on the network must have a unique computer name.**

You don’t have to know the names of all the computers on the network, but it helps if you know your own computer’s name and the names of any server computers you need to access.

The computer’s name is sometimes the same as the username of the person who uses the computer, but that’s usually a bad idea because in many companies, people come and go more often than computers. Sometimes the names indicate the physical location of the computer, such as `office-12` or `back-room`. Server computers often have names that reflect the group that uses the server most, like `acctng-server` or `cad-server`.

Some network nerds like to assign techie-sounding names, like `BL3K5-87a`. And some like to use names from science fiction movies; HAL, Co1ossus, M5, and Data come to mind. Cute names like Herbie aren’t allowed. (However, Tigger and Pooh are entirely acceptable — recommended, in fact. Networks are what Tiggers like the best.)

Usually, the sensible approach to computer naming is to use names that have numbers, such as `computer001` or `computer002`.



TIP



REMEMBER

- » **Network resources, such as shared disk folders and printers, have names.** For example, a network server may have two printers, named `laser` and `inkjet` (to indicate the type of printer), and two shared disk folders, named `AccountingData` and `MarketingData`.



TIP

» **Server-based networks have a username for the network administrator.**

If you log on using the administrator's username, you can do anything you want: add new users, define new network resources, change Wally's password, anything. The administrator's username is usually something clever such as Administrator.



REMEMBER

» **The network itself has a name.**

The Windows world has two basic types of networks:

- *Domain networks* are the norm for large corporate environments that have dedicated servers with IT staff to maintain them.
- *Workgroup networks* are more common in homes or in small offices that don't have dedicated servers or IT staff.

A domain network is known by — you guessed it — a *domain name*. And a workgroup network is identified by — drum roll, please — a *workgroup name*. Regardless of which type of network you use, you need to know this name to gain access to the network.

Logging On to the Network

To use network resources, you must connect your computer to the network, and you must go through the supersecret process of logging on, which is how you let the network know who you are so that it can decide whether you're one of the good guys.

Logging on is a little bit like cashing a check. You must have two forms of identification:

» **Your username:** The name by which the network knows you.

Your username is usually some variation of your real name, like *Beav* for the Beaver.

Everyone who uses the network must have a username.

» **Your password:** A secret word that only you and the network know. If you type the correct password, the network believes that you are who you say you are.

Every user has a different password, and the password should be a secret.



REMEMBER

In the early days of computer networking, you had to type a logon command at a stark MS-DOS prompt and then supply your user ID and password. Nowadays, the glory of Windows is that you get to log on to the network through a special network logon screen. Figure 2-1 shows the Windows 10 version of this dialog box.

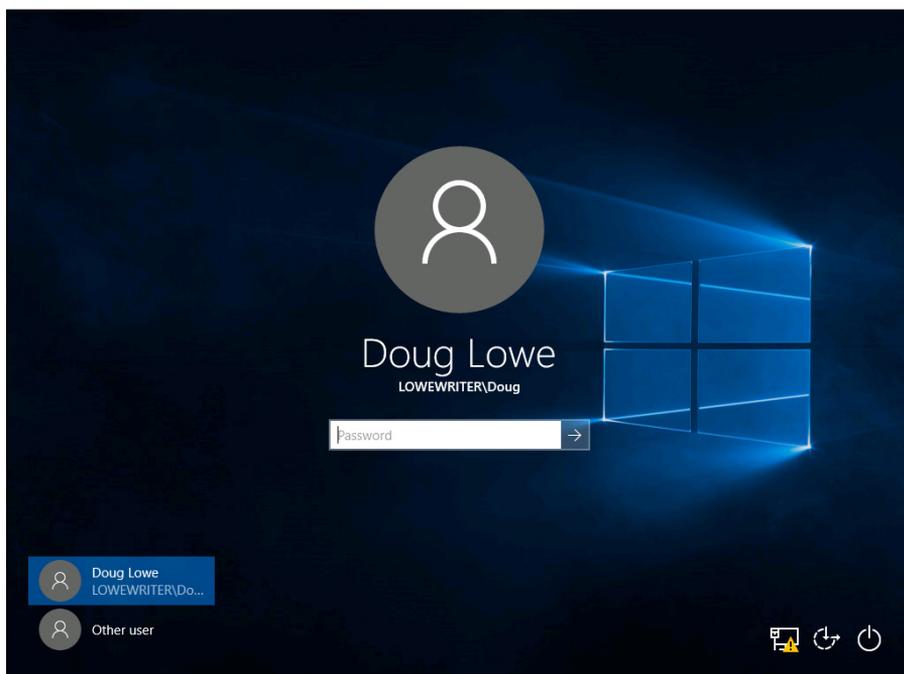


FIGURE 2-1:
Logging in to a
Windows 10
system.



TIP

Here are some more logon points to ponder:

- »» The terms *user ID* and *logon name* are sometimes used instead of *username*. They all mean the same thing.
- »» As long as we're talking about words that mean the same thing, *log in* and *log on* mean the same thing, as do (respectively) *log out* and *log off* as ways of saying, "I'm outta here." Although you see both out there in the world, this book uses *log on* and *log off* throughout — and if there's any exception, the book says why and grouches about it a bit.
- »» As far as the network's concerned, you and your computer aren't the same thing. Your username refers to you, not to your computer. That's why you have a username and your computer has a computer name. You can log on to the network by using your username from any computer that's attached to the network. Other users can log on at your computer by using their own usernames.

When others log on at your computer by using their own usernames, they can't access any of your network files that are protected by your password. However, they *can* access any local files that you haven't protected. Be careful which people you allow to use your computer.

- » If you're logging on to a domain network on a Windows computer, you must type the domain name before your username, separated from it by a backslash. For example:

```
lowewriter\dlowe
```

Here, the domain name is `lowewriter`, and the username is `dlowe`.

Note that Windows remembers the domain and username from your last login, so ordinarily all you have to enter is your password. To log on to a different domain or as a different user, you must click **Switch User**. Then you can click the **Other User** icon and enter a different domain name and username, along with the password for the user you want to log on as.

- » On an older Windows XP system, the logon dialog box has a field in which you can enter the domain name you want to log on to.
- » Your computer may be set up so that it logs you on automatically whenever you turn it on. In that case, you don't have to type your username and password. This setup makes the task of logging on more convenient but takes the sport out of it. And it's a terrible idea if you're the least bit worried about bad guys getting into your network or personal files.
- » Guard your password with your life. I'd tell you mine, but then I'd have to shoot you.

Understanding Shared Folders

Long ago, in the days Before Network (B.N.), your computer probably had just one hard drive, known as the C: drive. Maybe it had two — C: and D:. The second drive might be another hard disk, or possibly a CD-ROM or DVD-ROM drive. Even to this day, the descendants of those drives are physically located inside your PC. They're your *local drives*.

Now that you're on a network, however, you may have access to drives that aren't located inside your PC but are located instead in one of the other computers on the network. These network drives can be located on a dedicated server computer or, in the case of a peer-to-peer network, on another client computer.

In some cases, you can access an entire network drive over the network. But in most cases, you can't access the entire drive. Instead, you can access only certain folders on the network drives. Either way, the shared drives or folders are known in Windows terminology as *shared folders*.

Here's where it gets confusing: The most common way to access a shared folder is to assign a drive letter to it. Suppose that a server has a shared folder named Marketing. You can assign drive letter M to this shared folder. Then you access the Marketing folder as drive M:. The M: drive is then called a *network drive* because it uses the network to access data in a shared folder. Assigning a drive letter to a shared folder is *mapping a drive*.

Shared folders can be set up with restrictions on how you can use them. For example, you may be granted full access to some shared folders so that you can copy files to or from them, delete files on them, or create or remove folders on them. On other shared folders, your access may be limited in certain ways. For example, you may be able to copy files to or from the shared folder but not delete files, edit files, or create new folders. You may also be asked to enter a password before you can access a protected folder. The amount of disk space you're allowed to use on a shared folder may also be limited. For more information about file-sharing restrictions, see Chapter 13.



TIP

In addition to accessing shared folders that reside on other people's computers, you can designate your computer as a server to enable other network users to access folders that you share. To find out how to share folders on your computer with other network users, see Chapter 3.

Four Good Uses for a Shared Folder

After you know which shared network folders are available, you may wonder what you're supposed to do with them. This section describes four good uses for a network folder.

Store files that everybody needs

A shared network folder is a good place to store files that more than one user needs to access. Without a network, you have to store a copy of the file on everyone's computer, and you have to worry about keeping the copies synchronized (which you can't do, no matter how hard you try). Or you can keep the file on a disk and pass it around. Or you can keep the file on one computer and play Musical

Chairs; whenever someone needs to use the file, he goes to the computer that contains the file.

On a network, you can keep one copy of the file in a shared folder on the network, and everyone can access it.

Store your own files

You can also use a shared network folder as an extension of your own hard drive storage. For example, if you filled up all the free space on your hard drive with pictures, sounds, and movies that you downloaded from the Internet, but the network server has billions and billions of gigabytes of free space, you have all the drive space you need. Just store your files on the network drive!

Here are a few guidelines for storing files on network drives:

- » **Using the network drive for your own files works best if the network drive is set up for private storage that other users can't access.** That way, you don't have to worry about the nosy guy down in Accounting who likes to poke around in other people's files.
- » **Don't overuse the network drive.** Remember that other users have probably filled up their own hard drives, so they want to use the space on the network drive too.
- » **Before you store personal files on a network drive, make sure that you have permission.** A note from your mom will do.
- » **On domain networks, a drive (typically, drive H:) is commonly mapped to a user's home folder.** The *home folder* is a network folder that's unique for each user. You can think of it as a network version of My Documents. If your network is set up with a home folder, use it rather than My Documents for any important work-related files. That's because the home folder is usually included in the network's daily backup schedule. By contrast, most networks do *not* back up data you store in My Documents.

Make a temporary resting place for files on their way to other users

“Hey, Wally, could you send me a copy of last month's baseball stats?”

“Sure, Beav.” But how? If the baseball stats file resides on Wally’s local drive, how does Wally send a copy of the file to Beaver’s computer? Wally can do it by copying the file to a network drive. Then Beaver can copy the file to his local hard drive.

Here are some tips to keep in mind when you use a network drive to exchange files with other network users:

- » **Remember to delete files that you saved to the network drive after they’re picked up!** Otherwise, the network drive quickly fills up with unnecessary files.
- » **Create a folder on the network drive specifically intended for holding files en route to other users.** I like to name this folder PITSTOP.



TIP

In many cases, it’s easier to send files to other network users by email than by using a network folder. Just send a message to the other network user and attach the file you want to share. The advantage of sending a file by email is that you don’t have to worry about details like where to leave the file on the server and who’s responsible for deleting the file.

Back up your local hard drive

If enough drive space is available on the file server, you can use it to store backup copies of the files on your hard drive. Just copy the files that you want to back up to a shared network folder.

Obviously, if you copy *all* your data files to the network drive — and everybody else follows suit — it can fill up quickly. Check with the network manager before you start storing backup copies of your files on the server. The manager may have already set up a special network drive that’s designed just for backups. And if you’re lucky, your network manager may be able to set up an automatic backup schedule for your important data so that you don’t have to remember to back it up manually.

I hope that your network administrator also routinely backs up the contents of the network server’s disk to tape. (Yes, *tape* — see Chapter 20 for details.) That way, if something happens to the network server, the data can be recovered from the backup tapes.

Oh, the Network Places You’ll Go

Windows enables you to access network resources, such as shared folders, by browsing the network. In Windows 7, choose Network from the Start menu. In Windows 8, 8.1, and 10, open Windows Explorer (click File Explorer on the

taskbar) and then click Network. Figure 2-2 shows the Windows 10 version of the network browser.

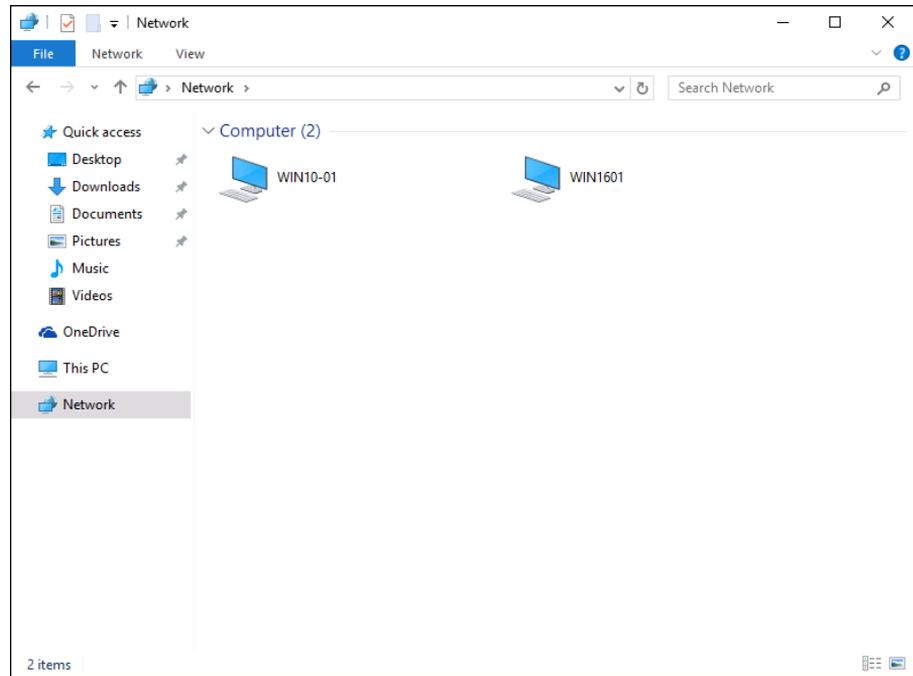


FIGURE 2-2:
Browsing the
network in
Windows 10.

The network shown in Figure 2-2 consists of just two computers: the desktop client computer running Windows 10 (WIN10-01) and a server computer running Windows Server 2016 (WIN1601). In an actual network, you would obviously see more than just two computers.

You can open a computer by double-clicking its icon to reveal a list of shared resources available on the computer. For example, Figure 2-3 shows the resources shared by the WIN1601 computer.

You can also browse the network from any Windows application program. For example, you may be working with Microsoft Word and want to open a document file that's stored in a shared folder on your network. All you have to do is use the Open command to bring up the dialog box, and then choose Network in the Navigation pane to view the available network devices.

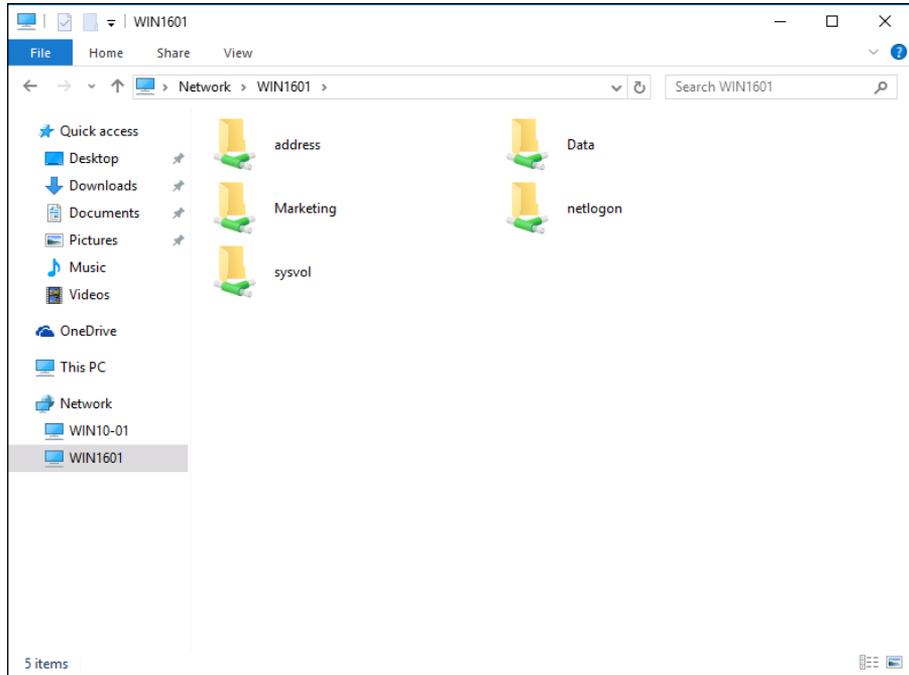


FIGURE 2-3:
The resources
available on a
server computer.

Mapping Network Drives



TIP

If you often access a particular shared folder, you may want to use the special trick known as mapping to access the shared folder more efficiently. *Mapping* assigns a drive letter to a shared folder. Then you can use the drive letter to access the shared folder as though it were a local drive. In this way, you can access the shared folder from any Windows program without having to browse the network.

For example, you can map a shared folder named *Data* on the server named *Win1601 Files* to drive *K:* on your computer. Then, to access files stored in the shared *Data* folder, you look on drive *K:*

To map a shared folder to a drive letter, follow these steps:

1. Open File Explorer.

- *Windows 7:* Choose Start ⇨ Computer.
- *Windows 8, 8.1, and 10:* Open the desktop and click the File Explorer icon on the taskbar, and then click Computer in the Location list on the left side of the screen.

2. Open the Map Network Drive dialog box.

- *Windows 7:* Access this dialog by clicking the Map Network Drive button located on the toolbar.
- *Windows 8 and 8.1:* Click Map Network Drive on the ribbon.
- *Windows 10:* Click the Computer tab, and then click Map Network Drive.

Figure 2-4 shows the Map Network Drive dialog box for Windows 10. The dialog box for earlier versions of Windows is similar.

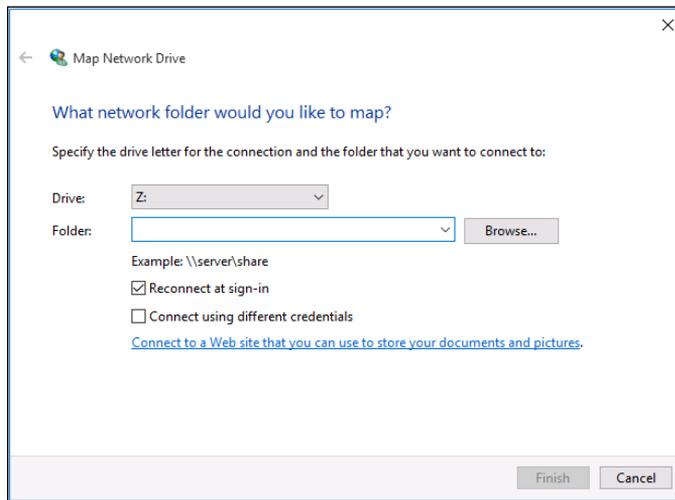


FIGURE 2-4:
The Map Network Drive dialog box.

3. (Optional) Change the drive letter in the Drive drop-down list.

You probably don't have to change the drive letter that Windows selects (in Figure 2-4, drive Z:). If you're picky, though, you can select the drive letter from the Drive drop-down list.

4. Click the Browse button.

This step summons the dialog box shown in Figure 2-5.

5. Use the Browse for Folder dialog box to find and select the shared folder you want to use.

You can navigate to any shared folder on any computer in the network.

6. Click OK.

The Browse for Folder dialog box is dismissed, and you return to the Map Network Drive dialog box (refer to Figure 2-4).

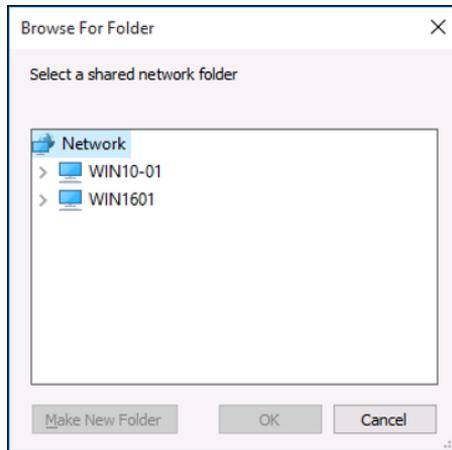


FIGURE 2-5:
Browsing for the
folder to map.

7. (Optional) If you want this network drive to be automatically mapped each time you log on to the network, select the Reconnect at Sign-in check box.

If you leave the Reconnect at Sign-in check box deselected, the drive letter is available only until you shut down Windows or log out of the network. If you select this option, the network drive reconnects automatically each time you log on to the network.



TIP

8. Click OK.

You return to the This PC folder, as shown in Figure 2-6. Here, you can see the newly mapped network drive.

Your network administrator may have already set up your computer with one or more mapped network drives. If so, you can ask her to tell you which network drives have been mapped. Or you can just open the This PC folder and have a look.

Here are a few additional tips:

- » **Assigning a drive letter to a network drive is called *mapping the drive*, or *linking the drive*, by network nerds.** “Drive Q: is mapped to a network drive,” they say.
- » **Network drive letters don’t have to be assigned the same way for every computer on the network.** For example, a network drive that’s assigned drive letter M on your computer may be assigned drive letter Z on someone else’s computer. In that case, your drive M: and the other computer’s drive Z: refer to the same data. This arrangement can be confusing. If your network is set up this way, put pepper in your network administrator’s coffee.

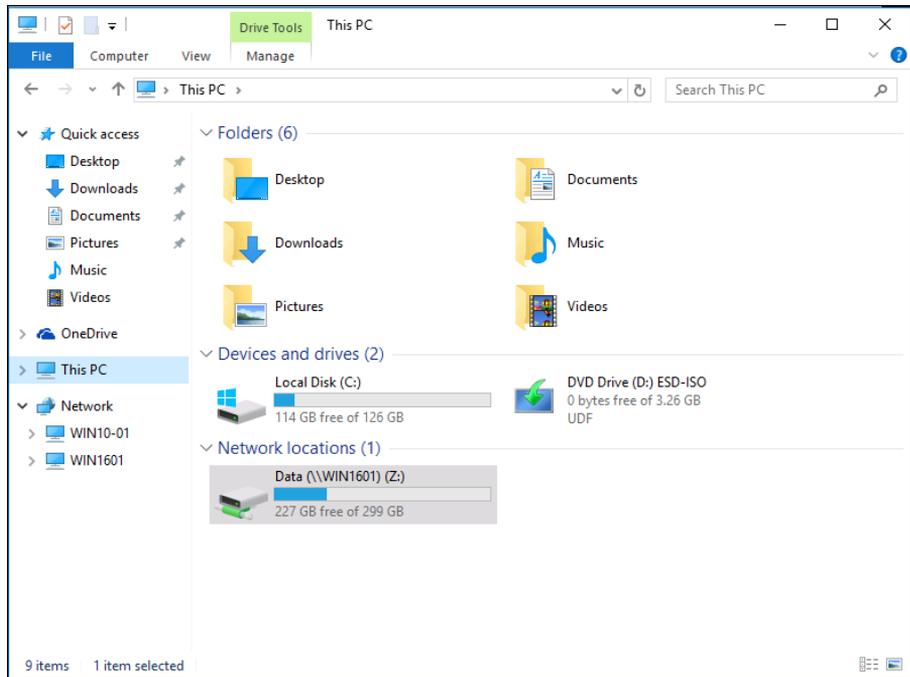


FIGURE 2-6:
The This PC folder shows a mapped network drive.

- » **Accessing a shared network folder through a mapped network drive is much faster than accessing the same folder by browsing the network.** Windows has to browse the entire network to list all available computers whenever you browse the network. By contrast, Windows doesn't have to browse the network to access a mapped network drive.
- » **If you select the Reconnect at Sign-in option for a mapped drive (refer to Figure 2-4), you receive a warning message if the drive isn't available when you log on.** In most cases, the problem is that the server computer isn't turned on. Sometimes, however, this message is caused by a broken network connection. For more information about fixing network problems such as this one, see Chapter 19.

Using a Network Printer

Using a network printer is much like using a network hard drive: You can print to a network printer from any Windows program by choosing the Print command to call up a Print dialog box from any program and choosing a network printer from the list of available printers.

Keep in mind, however, that printing on a network printer isn't exactly the same as printing on a local printer; you have to take turns. When you print on a local printer, you're the only one using it. When you print to a network printer, however, you are (in effect) standing in line behind other network users, waiting to share the printer. This line complicates the situation in several ways:

- » **If several users print to the network printer at the same time, the network has to keep the print jobs separate from one another.** If it didn't, the result would be a jumbled mess, with your 268-page report getting mixed in with the payroll checks. That would be bad. Fortunately, the network takes care of this situation by using the fancy *print spooling* feature.
- » **Network printing works on a first-come, first-served basis.** Invariably, when I get in line at the hardware store, the person in front of me is trying to buy something that doesn't have a product code on it. I end up standing there for hours waiting for someone in Plumbing to pick up the phone for a price check. Network printing can be like that. If someone sends a two-hour print job to the printer before you send your half-page memo, you have to wait.
- » **You may have access to a local printer and several network printers.** Before you were forced to use the network, your computer probably had just one printer attached to it. You may want to print some documents on your cheap (oops, I mean *local*) inkjet printer but use the network laser printer for important stuff. To do that, you have to find out how to use your programs' functions for switching printers.

Adding a network printer

Before you can print to a network printer, you have to configure your computer to access the network printer that you want to use. From the Start menu, open the Control Panel and then double-click the Printers icon. If your computer is already configured to work with a network printer, an icon for the network printer appears in the Printers folder. You can tell a network printer from a local printer by the shape of the printer icon. Network printer icons have a pipe attached to the bottom of the printer.

If you don't have a network printer configured for your computer, you can add one by using the Add Printer Wizard. Just follow these steps:

1. **Open the Control Panel.**
 - *Windows 7 or earlier:* Choose Start ⇨ Control Panel.
 - *Windows 8 and later:* Press the Windows key, type **Control**, and then click the Control Panel icon.

2. **Click Devices and Printers.**
3. **Click the Add a Printer button on the toolbar.**

This step starts the Add Printer Wizard, as shown in Figure 2-7.

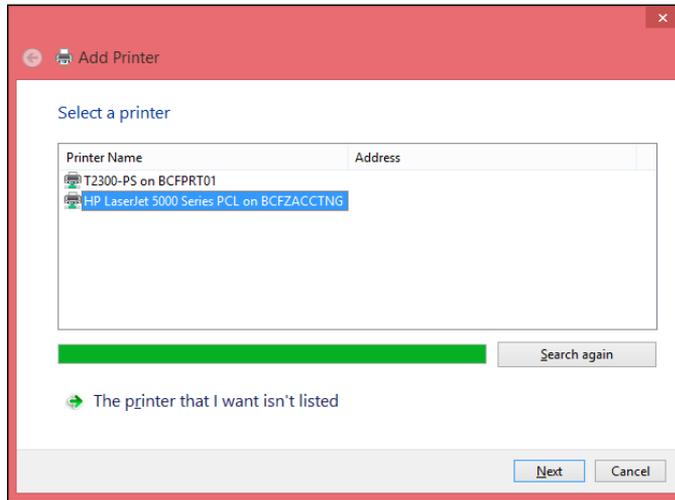


FIGURE 2-7:
The Add Printer Wizard asks you to pick a printer.



TIP

4. **Click the printer you want to use.**

If you can't find the printer you want to use, ask your network administrator for the printer's *UNC path*, which is the name used to identify the printer on the network, or its IP address. Then click The Printer That I Want Isn't Listed and enter the UNC or IP address for the printer when prompted.

5. **Click Next to add the printer.**

The wizard copies to your computer the correct printer driver for the network printer. (You may be prompted to confirm that you want to add the driver. If so, click Install Driver to proceed.)

The Add Printer Wizard displays a screen that shows the printer's name and asks whether you want to designate the printer as your default printer.

6. **(Optional) Designate the printer as your default printer.**

7. **Click Next to continue.**

A final confirmation dialog box is displayed.

8. **Click Finish.**

You're done!



TIP

Many network printers, especially newer ones, are connected directly to the network by using a built-in Ethernet card. Setting up these printers can be tricky. You may need to ask the network administrator for help in setting up this type of printer. (Some printers that are connected directly to the network have their own web addresses, such as `Printer.CleaverFamily.com`. If that's the case, you can often set up the printer in a click or two: Use your browser to go to the printer's web page and then click a link that enables you to install the printer.)

Printing to a network printer

After you install the network printer in Windows, printing to the network printer is a snap. You can print to the network printer from any Windows program by using the Print command to summon the Print dialog box, which is usually found on the File menu. For example, Figure 2-8 shows the Print dialog box for WordPad (the free text-editing program that comes with Windows). The available printers are listed near the top of this dialog box. Choose the network printer from this list and then click OK to print your document. That's all there is to it!

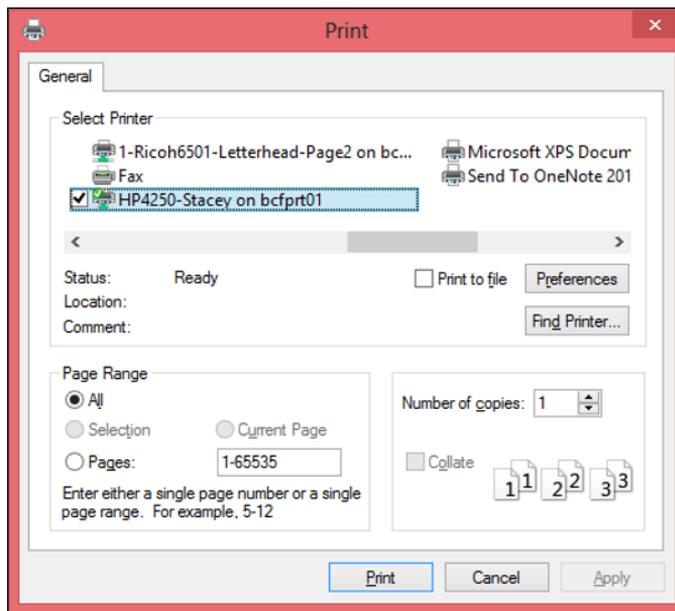


FIGURE 2-8:
A typical Print dialog box.

Playing with the print queue

After you send your document to a network printer, you usually don't have to worry about it. You just go to the network printer, and voilà! Your printed document is waiting for you.

That's what happens in the ideal world. In the real world, where you and I live, all sorts of things can happen to your print job between the time you send it to the network printer and the time it prints:

- » You discover that someone else already sent a 50 trillion-page report ahead of you that isn't expected to finish printing until the national debt is paid off.
- » The price of a framis valve suddenly goes up by \$2, rendering foolish the recommendations you made in your report.
- » Your boss calls and tells you that his brother-in-law will be attending the meeting, so won't you please print an extra copy of the proposal for him? Oh, and a photocopy won't do. Originals only, please.
- » You decide to take lunch, so you don't want the output to print until you get back.

Fortunately, your print job isn't totally beyond your control just because you already sent it to the network printer. You can easily change the status of jobs that you already sent. You can change the order in which jobs print, hold a job so that it doesn't print until you say so, or cancel a job.

You can probably make your network print jobs do other tricks, too: shake hands, roll over, and play dead. But the basic tricks — hold, cancel, and change the print order — are enough to get you started.

To play with the printer queue, open the Control Panel by choosing Start ⇨ Control Panel in Windows 7 or earlier; or press the Windows key, type **Control**, and then click the Control Panel icon. Then click Devices and Printers and double-click the icon for the printer that you want to manage. A window similar to the one shown in Figure 2-9 appears. You can see that just one document has been sent to the printer.

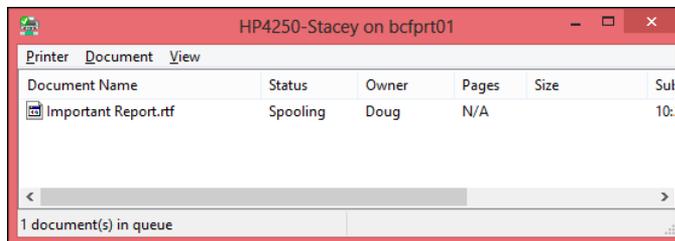


FIGURE 2-9:
Managing a print queue.

To manipulate the print jobs that appear in the print queue or in the printer itself, use these tricks:

- » **To temporarily stop a job from printing:** Select the job and choose Document ⇨ Pause Printing. Choose the same command again to release the job from its state of frustration and print it out, already.
- » **To delete a print job:** Select the job and choose Document ⇨ Cancel Printing.
- » **To stop the printer:** Choose Printer ⇨ Pause Printing. To resume, choose the command again.
- » **To delete all print jobs:** Choose Printer ⇨ Purge Print Documents.
- » **To cut to the front of the line:** Drag to the top of the list the print job that you want to print.

All these tips apply to your print jobs only. Unfortunately, you can't capriciously delete other people's print jobs.

The best thing about Windows printer management is that it shelters you from the details of working with different network operating systems. Whether you print on a NetWare printer, a Windows 2003 network printer, or a shared Windows printer, the Printer window icon manages all print jobs in the same way.

Logging Off the Network

After you finish using the network, log off. Logging off the network makes the network drives and printers unavailable. Your computer is still physically connected to the network (unless you cut the network cable with pruning shears; it's a bad idea — don't do it!), but the network and its resources are unavailable to you.

Here are a few other tips to keep in mind when you log off:

- » After you turn off your computer, you're automatically logged off the network. After you start your computer, you have to log on again.



REMEMBER

Logging off the network is a good idea if you're going to leave your computer unattended for a while. As long as your computer is logged in to the network, anyone can use it to access the network. And because unauthorized users can access it under your user ID, you get the blame for any damage they do.

» In Windows, you can log off the network by clicking the Start button and choosing the Log Off command. This process logs you off the network without restarting Windows:

- *In Windows 7:* Click Start and then click the right-facing arrow that appears next to the little padlock icon.
- *In Windows 8 and later:* Press Ctrl+Alt+Del and then choose Sign Out.

IN THIS CHAPTER

Transforming your computer into a network server

Sharing folders with network users

Working in the Public folder

Sharing your printer

Using Office on a network

Working with files offline

Chapter 3

More Ways to Use Your Network

Chapter 2 introduces you to the basics of using a network: logging on, accessing data on shared network folders, printing, and logging off. In this chapter, I go beyond these basics. You find out how to turn your computer into a server that shares its own files and printers, how to use one of the most popular network computer applications — email — and how to work with Office on a network.

Sharing Your Stuff

As you probably know, networks consist of two types of computers: client computers and server computers. In the economy of computer networks, *client computers* are the consumers — the ones that use network resources, such as shared printers and disk drives. *Servers* are the providers — the ones that offer their own printers and hard drives to the network so that the client computers can use them.

This chapter shows you how to turn your humble Windows client computer into a server computer so that other computers on your network can use your printer and any folders that you decide you want to share. In effect, your computer functions as both a client and a server at the same time. A couple of examples show how:

- » It's a **client** when you send a print job to a network printer or when you access a file stored on another server's hard drive.
- » It's a **server** when someone else sends a print job to your printer or accesses a file stored on your computer's hard drive.

Enabling File and Printer Sharing

Before you can share your files or your printer with other network users, you must set up a Windows File and Printer Sharing feature. Without this feature installed, your computer can be a network client but not a server.

If you're lucky, the File and Printer Sharing feature is already set up on your computer. To find out, open Windows Explorer and right-click Desktop in the Navigation pane. If the menu includes a Share With command, File and Printer Sharing is already set up, so you can skip the rest of this section. If you can't find a Share With command, follow these steps:

- 1. Click the Start button, type Network and Sharing Center, and press Enter.**

This step opens the Network and Sharing Center.

- 2. Click Change Advanced Sharing Settings.**

The Advanced Sharing Settings page is displayed.

- 3. Click the down arrow next to the network you want to enable file and printer sharing for.**

- *For a home computer:* Click the down arrow next to Home or Work (Windows 7) or Private (Windows 8 and later).
- *For a computer in a public location:* Click the down arrow next to Guest or Public.
- *For a computer connected to a domain network:* Click the down arrow next to Domain.



WARNING

Figure 3-1 shows the settings for a Domain network. The settings for a Home, Guest, or Public computer are the same.

Do *not* enable file or printer sharing for the Public network. Enabling file or printer sharing on a public network exposes your computer's data to other users on the same public network.

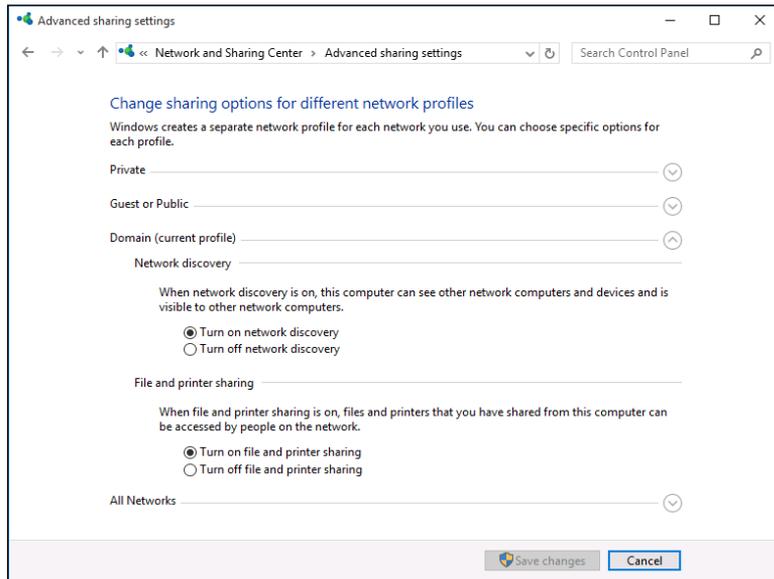


FIGURE 3-1: Enabling file and printer sharing.

4. **Select the Turn on File and Printer Sharing option.**
5. **Click the Save Changes button.**

This action saves your changes and closes the Advanced Sharing Settings page.

Sharing a Folder

To enable other network users to access files that reside on your hard drive, you must designate a folder on the drive as a *shared* folder. Note that you can also share an entire drive, if you so desire. If you share an entire drive, other network users can access all the files and folders on the drive. If you share a folder, network users can access only those files that reside in the folder you share. (If the folder you share contains other folders, network users can access files in those folders, too.)



WARNING

Don't share an entire hard drive unless you want to grant *everyone on the network* the freedom to sneak a peek at every file on your hard drive. Instead, you should share just the folder or folders containing the specific documents that you want others to be able to access. For example, if you store all your Word documents in the My Documents folder, you can share your My Documents folder so that other network users can access your Word documents.

To share a folder on a desktop version of Windows, follow these steps:

1. Open File Explorer.

- *Windows 7:* Choose Start ⇨ Computer.
- *Windows 8 and later:* Open the desktop and click the File Explorer icon on the taskbar; then click Computer in the Location list on the left side of the screen.

2. Navigate to the folder you want to share.

3. Right-click the folder you want to share and choose Properties.

The Properties dialog box appears.

4. Click the Sharing tab and then click the Share button.

The File Sharing dialog box appears, as shown in Figure 3-2.

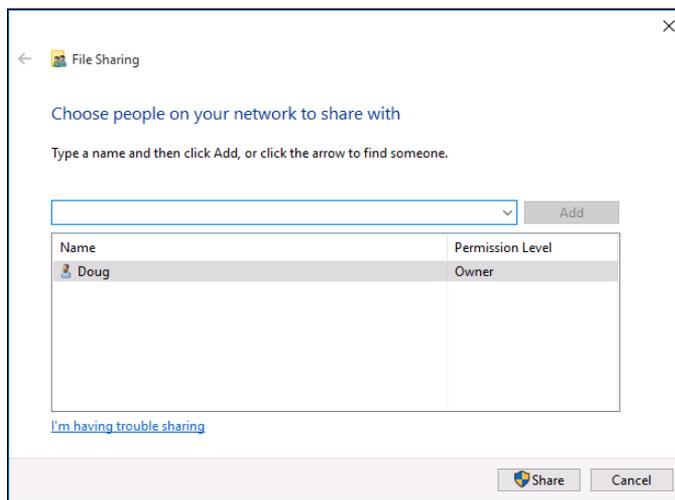


FIGURE 3-2:
The File Sharing dialog box.

5. Click the arrow in the drop-down list, choose Everyone, and then click Add.

This action designates that anyone on your network can access the shared folder.

If you prefer, you can limit access to just certain users. To do so, select each person you want to grant access to and then click Add.

6. Select the level of access you want to grant each user.

You can use the drop-down list in the Permission Level column to choose from three levels of access:

- *Reader:* A reader can open files but can't modify or create new files or folders.
- *Contributor:* A contributor can add files to the share but can change or delete only her own files.
- *Owner:* An owner has full access to the shared folder. He or she can create, change, or delete any file in the folder.

7. Click Share.

A confirmation dialog box appears to confirm that the folder has been shared.

Using the Public Folder

Windows includes an alternative method of sharing files on the network: the Public folder. The *Public folder* is simply a folder that's designated for public access. Files you save in this folder can be accessed by other users on the network and by any user who logs on to your computer.

Before you can use the Public folder, you must enable it. Just follow the steps listed in the section "Enabling File and Printer Sharing" earlier in this chapter, but choose the Turn on Sharing option in the All Networks (Windows 10) or Public Sharing Settings (Windows 8 and 8.1) section.

After you enable Public folder sharing, you can access the Public folder on your own computer in Windows 7 by choosing Start ⇨ Computer, expanding the Libraries item in the left pane, and then expanding the Documents, Music, Pictures, or Videos items. In Windows 8 and later, open the desktop, click the File Explorer icon on the taskbar, expand the Libraries item in the left pane, and then expand the Documents, Music, Pictures, or Videos items.

Figure 3-3 shows an example of a Public folder.

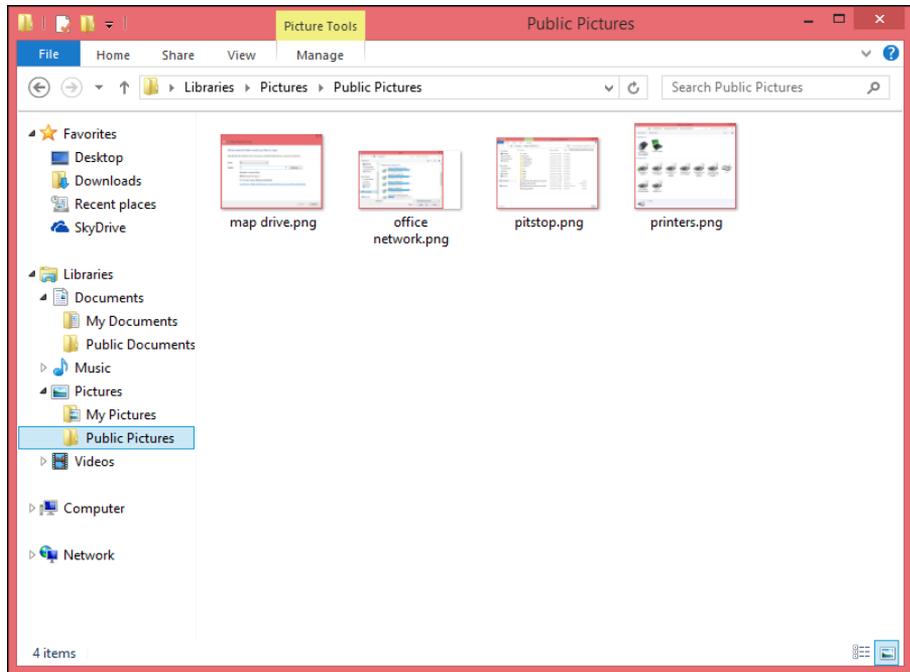


FIGURE 3-3:
The Public folder.

As you can see, the Public folder includes several predefined subfolders designed for sharing documents, downloaded files, music, pictures, and videos. You can use these subfolders if you want, or you can create your own subfolders to help organize the data in your Public folder.



TIP

To access the Public folder of another computer, use the techniques that I describe in Chapter 2 to either browse to the Public folder or map it to a network drive.

Sharing a Printer

Sharing a printer is much more traumatic than sharing a hard drive. When you share a hard drive, other network users access your files from time to time. When they do, you hear your drive click a few times, and your computer may hesitate for a half-second or so. The interruptions caused by other users accessing your drive are sometimes noticeable but rarely annoying.

When you share a printer, you get to see Murphy's Law in action: Your co-worker down the hall is liable to send a 140-page report to your printer just moments before you try to print a 1-page memo that has to be on the boss's desk in two

minutes. The printer may run out of paper — or worse, jam — during someone else’s print job — and you’re expected to attend to the problem.

Although these interruptions can be annoying, sharing your printer makes a lot of sense in some situations. If you have the only decent printer in your office or workgroup, everyone will bug you to let them use it anyway. You may as well share the printer on the network. At least this way, they won’t line up at your door to ask you to print their documents for them.

To share a printer, follow these steps:

1. Open the Control Panel.

- *Windows 7 or earlier:* Choose Start ⇨ Control Panel.
- *Windows 8 and later:* Press the Windows key, type **Control**, and then click the Control Panel icon.

2. Click Devices and Printers.

3. Right-click the printer that you want to share and choose Printer Properties.

The Properties dialog box for the printer appears.

4. Click the Sharing tab.

The Sharing tab appears, as shown in Figure 3-4. Notice that the options for sharing the printer are disabled.

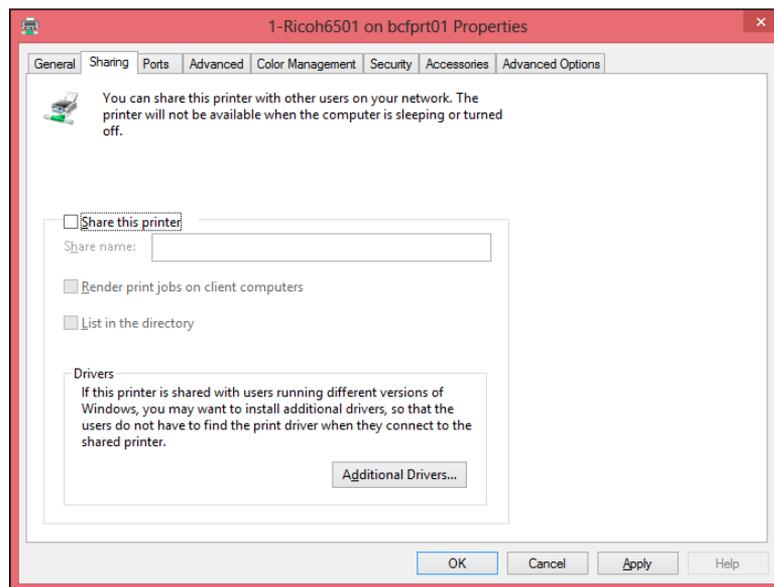


FIGURE 3-4: Sharing a printer.

5. **Select the Share This Printer option.**
6. **(Optional) Change the share name if you don't like the name suggested by Windows.**

Because other computers will use the share name to identify the shared printer, pick a descriptive name.

7. **Click OK.**

You return to the Printers folder. The icon for the printer is modified to indicate that it has been shared.

To take your shared printer off the network so that other network users can't access it, follow Steps 1–4 in the preceding set of steps. Deselect the Share This Printer check box and then click OK.

Using Microsoft Office on a Network

Microsoft Office is far and away the most popular suite of application programs used on personal computers, and it includes the most common types of application programs used in an office: a word processing program (Word), a spreadsheet program (Excel), a presentation program (PowerPoint), and an excellent email program (Outlook). Depending on the version of Office you purchase, you may also get a database program (Access), a desktop publishing program (Publisher), a set of Ginsu knives (KnifePoint), and a slicer and dicer (ActiveSalsa).



TIP

To get the most from using Office on a network, you should purchase the Microsoft Office Resource Kit. The Office Resource Kit, also known as *ORK*, contains information about installing and using Office on a network and comes with a CD that has valuable tools. If you don't want to purchase the ORK, you can view it online and download the ORK tools from the Microsoft TechNet website (<http://technet.microsoft.com/en-US/>). Nanoo-nanoo, Earthling.

Accessing network files

Opening a file that resides on a network drive is almost as easy as opening a file on a local drive. All Office programs use File⇄Open to summon the Open dialog box, as shown in its Excel incarnation in Figure 3-5. (The Open dialog box is nearly identical in other Office programs.)

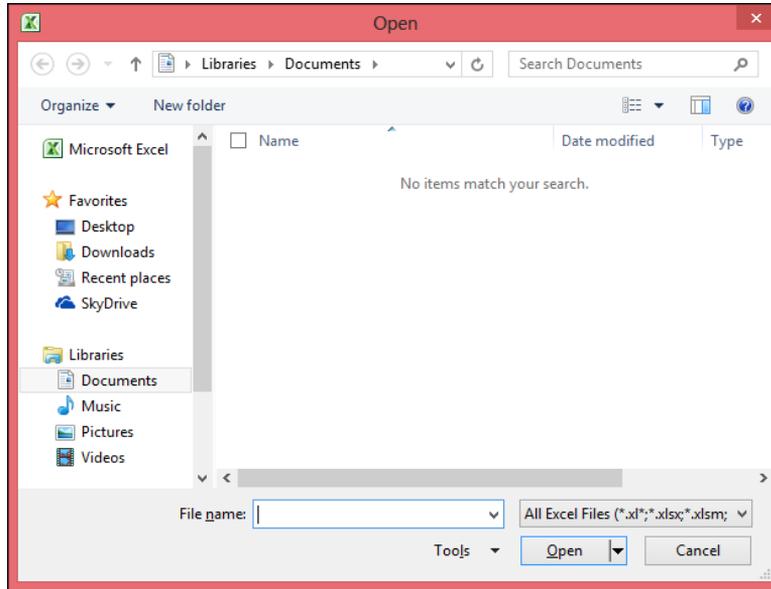


FIGURE 3-5:
The Open dialog
box in Excel.

To access a file that resides on a network volume that's mapped to a drive letter, all you have to do is use the drop-down list at the top of the dialog box to select the network drive.

You can map a network drive directly from the Open dialog box by navigating to the folder you want to map, right-clicking the folder, and choosing Map Network Drive.



TIP

If you try to open a file that another network user has opened already, Office tells you that the file is already in use and offers to let you open a read-only version of the file. You can read and edit the read-only version, but Office doesn't let you overwrite the existing version of the file. Instead, you have to use the Save As command to save your changes to a new file.

Using workgroup templates

Although an occasional sacrifice to the Office gods may make your computing life a bit easier, a template isn't a place of worship. Rather, a *template* is a special type of document file that holds formatting information, boilerplate text, and other customized settings that you can use as the basis for new documents.

Three Office programs — Word, Excel, and PowerPoint — enable you to specify a template whenever you create a new document. When you create a new document

in Word, Excel, or PowerPoint by choosing File⇨New, you see a dialog box that lets you choose a template for the new document.

Office comes with a set of templates for the most common types of documents. These templates are grouped under the various tabs that appear across the top of the New dialog box.

In addition to the templates that come with Office, you can create your own templates in Word, Excel, and PowerPoint. Creating your own templates is especially useful if you want to establish a consistent look for documents prepared by your network users. For example, you can create a Letter template that includes your company's letterhead or a Proposal template that includes a company logo.

Office enables you to store templates in two locations. Where you put them depends on what you want to do with them:

- » **The User Templates folder on each user's local disk drive:** If a particular user needs a specialized template, put it here.
- » **The Workgroup Templates folder on a shared network drive:** If you have templates that you want to make available to all network users on the network server, put them here. This arrangement still allows each user to create templates that aren't available to other network users.

When you use both a User Templates folder and a Workgroup Templates folder, Office combines the templates from both folders and lists them in alphabetical order in the New dialog box. For example, the User Templates folder may contain templates named Blank Document and Web Page, and the Workgroup Templates folder may contain a template named Company Letterhead. In this case, three templates appear in the New dialog box, in this order: Blank Document, Company Letterhead, and Web Page.

To set the location of the User Templates and Workgroup Templates folders, follow these steps in Microsoft Word:

- 1. Click the Office button and then click Word Options.**
The Word Options dialog box opens.
- 2. Click the Advanced tab.**
The Advanced options appear.
- 3. Scroll down to the General section and then click the File Locations button.**
The File Locations dialog box appears, as shown in Figure 3-6.

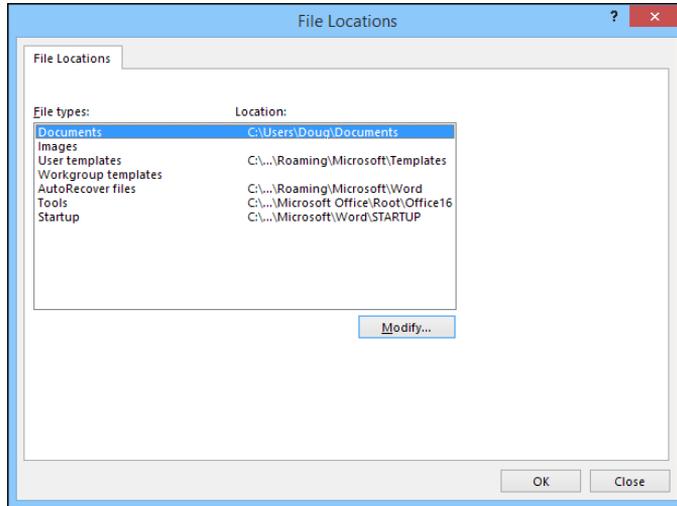


FIGURE 3-6: Setting the file locations.

4. Double-click the Workgroup Templates item.

This step opens a dialog box that lets you browse to the location of your template files.

5. Browse to the template files and then click OK.

You return to the File Locations dialog box.

6. Click OK to dismiss the File Locations dialog box.

You return to the Word Options dialog box.

7. Click OK again.

The Word Options dialog box is dismissed.



TIP

Although the User Templates and Workgroup Templates settings affect Word, Excel, and PowerPoint, you can change these settings only from Word. The Options dialog boxes in Excel and PowerPoint don't show the User Templates or Workgroup Templates options.

When you install Office, the standard templates that come with Office are copied into a folder on the computer's local disk drive, and the User Templates option is set to this folder. The Workgroup Templates option is left blank. You can set the Workgroup Templates folder to a shared network folder by clicking Network Templates, clicking the Modify button, and specifying a shared network folder that contains your workgroup templates.

Networking an Access database

If you want to share a Microsoft Access database among several network users, be aware of a few special considerations. Here are the most important ones:

- » When you share a database, more than one user may try to access the same record at the same time. This situation can lead to problems if two or more users try to update the record. To handle this potential traffic snarl, Access locks the record so that only one user at a time can update it. Access uses one of three methods to lock records:
 - *Edited Record*: This method locks a record whenever a user begins to edit a record. For example, if a user retrieves a record in a form that allows the record to be updated, Access locks the record while the user edits it so that other users can't edit the record until the first record is finished.
 - *No Locks*: This method doesn't really mean that the record isn't locked. Instead, No Locks means that the record isn't locked until a user writes a change to the database. This method can be confusing to users because it enables one user to overwrite changes made by another user.
 - *All Records*: All Records locks an entire table whenever a user edits any record in the table.
- » Access lets you split a database so that the forms, queries, and reports are stored on each user's local disk drive, but the data itself is stored on a network drive. This feature can make the database run more efficiently on a network, but it's a little more difficult to set up. (To split a database, choose Tools ⇨ Database Utilities ⇨ Database Splitter.)
- » Access includes built-in security features that you should use if you share an Access database from a Windows client computer. If you store the database on a domain server, you can use the server's security features to protect the database.
- » Access automatically refreshes forms and datasheets every 60 seconds. That way, if one user opens a form or datasheet and another user changes the data a few seconds later, the first user sees the changes within one minute. If 60 seconds is too long (or too short) an interval, you can change the refresh rate by using the Advanced tab in the Options dialog box.

Working with Offline Files

Desktop computers are by nature stationary beasts. As a result, they're almost always connected to their networks. Notebook computers, however, are more transitory. If you have a notebook computer, you're likely to tote it around from place to place. If you have a network at work, you probably connect to the network when you're at work. But then you take the notebook computer home for the weekend, and you aren't connected to your network.

Of course, your boss wants you to spend your weekends working, so you need a way to access your important network files while you're away from the office and disconnected from the network. That's where the offline files feature comes in. It lets you access your network files even while you're disconnected from the network.

It sounds like magic, but it isn't really. Imagine how you'd work away from the network without this feature. You simply copy the files you need to work on to your notebook computer's local hard disk. Then, when you take the computer home, you work on the local copies. When you get back to the office, you connect to the network and copy the modified files back to the network server.

That's essentially how the offline files feature works, except that Windows does all the copying automatically. Windows also uses symlinks and mirrors to make it look like the copies are actually on the network even though you're not connected to the network. For example, if you map a drive (drive M:, for example) and make it available offline, you can still access the offline copies of the file on the M: drive. That's because Windows knows that when you aren't connected to the network, it should redirect drive M: to its local copy of the drive M: files.

The main complication of working with offline files, of course, is what happens when two or more users want to access the same offline files. Windows can attempt to straighten that mess out for you, but it doesn't do a great job of it. Your best bet is to not use the offline files feature with network resources that other users may want available offline, too. In other words, it's okay to make your home drive available offline because that drive is accessible only to you. I don't recommend making shared network resources available offline, though, unless they're read-only resources that don't contain files you intend to modify.

Before you can use offline files, you must first enable the Offline Files feature. To do that, open the Control Panel, double-click the Sync Center icon, and click Manage Offline Files. This brings up the Offline Files dialog box, shown in Figure 3-7. Next, click Enable Offline Files and then click OK.

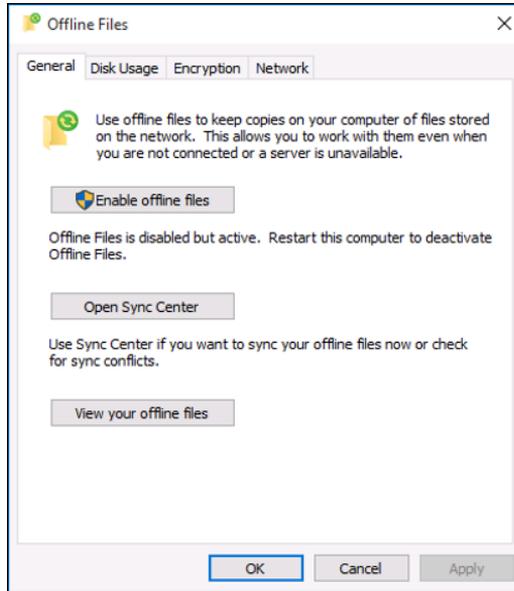


FIGURE 3-7:
Enabling offline files.

After you've enabled offline files, using the offline files feature is easy: Just open the Computer folder, right-click the mapped network drive you want to make available offline, and choose Always Available Offline.

If you don't want to designate an entire mapped drive for offline access, you can designate individual folders within a mapped drive by using the same technique: Right-click the folder and then choose Always Available Offline.

When you first designate a drive or folder as available offline, Windows copies all the files on the drive or folder to local storage. Depending on how many files are involved, this process can take a while, so plan accordingly.

After you designate a drive as available offline, Windows takes care of the rest. Each time you log on to or out of the network, Windows synchronizes your offline files. Windows compares the time stamp on each file on both the server and the local copy and then copies any files that have changed.

Here are a few other thoughts to consider about offline files:

- » If you want, you can force Windows to synchronize your offline files by right-clicking the drive or folder and choosing Sync.
- » Make sure that no files in the folder are currently open at the time you set the Make Available Offline option. If any files are open, you'll receive an error

message. You have to close the open files before you can designate the folder for offline access.

- » The Properties dialog box for mapped drives includes an Offline Files tab, as shown in Figure 3-8.

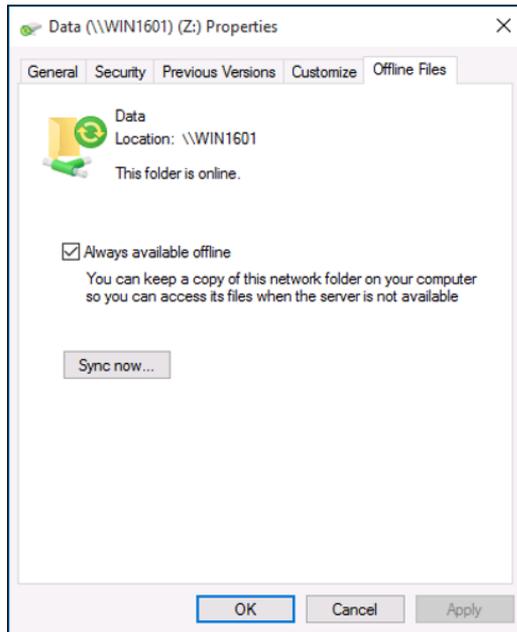


FIGURE 3-8:
Offline file properties.

- » Employers love the offline files feature because it encourages their employees to work at home during evenings and weekends. In fact, every time you use the offline files feature to work at home, your boss sends Bill Gates a nickel. That's how he got so rich.



Setting Up a Network

IN THIS PART . . .

Create a network plan.

Understand and use TCP/IP.

Work with cables, network adapters, switches, and other important networking components.

Configure Windows computers for networking.

Get connected to the Internet.

Use wireless devices in your network.

Virtualize your servers.

Making a network plan**Taking stock of your computer stock****Making sure that you know why you need a network****Identifying basic network decisions that you can't avoid**

Chapter 4

Planning a Network

Okay, so you're convinced that you need to network your computers. What now? Do you stop by Computers-R-Us on the way to work, install the network before drinking your morning coffee, and expect the network to be fully operational by noon?

I don't think so.

Networking your computers is just like any other worthwhile endeavor: Doing it right requires a bit of planning. This chapter helps you to think through your network before you start spending money. It shows you how to come up with a networking plan that's every bit as good as the plan that a network consultant would charge thousands of dollars for. See? This book is already saving you money!

Making a Network Plan

Before you begin any networking project, whether a new network installation or an upgrade of an existing network, start with a detailed plan. If you make technical decisions too quickly before studying all the issues that affect the project, you'll regret it. You'll discover too late that a key application won't run over the network, the network has unacceptably slow performance, or key components of the network don't work together.

Here are some general thoughts to keep in mind while you create your network plan:

- » **Don't rush the plan.** The costliest networking mistakes are the ones that you make before you install the network. Think things through and consider alternatives.
- » **Write down the network plan.** The plan doesn't have to be a fancy, 500-page document. If you want to make it look good, pick up a small three-ring binder. This binder will be big enough to hold your network plan with room to spare.
- » **Ask someone else to read your network plan before you buy anything.** Preferably, ask someone who knows more about computers than you do.
- » **Keep the plan up to date.** If you add to the network, dig up the plan, dust it off, and update it.



TIP

“The best laid schemes of mice and men gang aft a-gley, and leave us naught but grief and pain for promised joy.” Robert Burns lived a few hundred years before computer networks, but his famous words ring true. A network plan isn't chiseled in stone. If you discover that something doesn't work how you thought it would, that's okay. Just change your plan.

Being Purposeful

One of the first steps in planning your network is making sure that you understand why you want the network in the first place. Here are some of the more common reasons for creating or upgrading a network, all of them quite valid:

- » Everyone in the office needs access to the Internet. Probably the most common reason for setting up a small network is to share an Internet connection. And even in larger networks, shared Internet access is one of the primary benefits of the network.
- » My co-worker and I exchange files using flash drives just about every day. With a network, it would be easier to trade files.
- » I don't want to buy everyone a color laser printer when I know the one we have now just sits there taking up space most of the day. So wouldn't investing in a network be better than buying a color laser printer for every computer?
- » Business is so good that one person typing in orders eight hours each day can't keep up. If the sales and accounting data existed on a network server, I could hire another person to help, and I won't have to pay overtime to either person.

- » My sister-in-law just upgraded the network at her office, and I don't want her to think that I'm behind the times.
- » My existing network performs like it's made of kite string and tin cans. I should have upgraded it five years ago to speed up access to shared files, provide better security, and easier management.

Make sure that you identify all the reasons why you think you need a network and then write them down. Don't worry about winning the Pulitzer Prize for your stunning prose. Just make sure that you write down what you expect a network to do for you.

If you were making a 500-page networking proposal, you'd place the description of why a network is needed in a tabbed section labeled "Justification." In your network binder, file the description under "Purpose."



TIP

As you consider the reasons why you need a network, you may conclude that you don't need a network after all. That's okay. You can always use the binder for your stamp collection.

Taking Stock

One of the most challenging parts of planning a network is figuring out how to work with the computers that you already have. In other words, how do you get from here to there? Before you can plan how to get "there," you have to know where "here" is. In other words, you have to take a thorough inventory of your current computers.

What you need to know

You need to know the following information about each of your computers. Don't sweat it right now if some of these terms don't make sense. They're all just pieces of the puzzle.

- » **The processor type and, if possible, its clock speed:** It would be nice if each of your computers had a shiny new Core i7 eight-core processor. In most cases, though, you find a mixture of computers: some new, some old, some borrowed, some blue. You may even find a few archaic Pentium computers.

You can't usually tell what kind of processor that a computer has just by looking at the computer's case. Most computers, however, display the processor type when you turn them on or reboot them. If the information on the startup

screen scrolls too quickly for you to read it, try pressing Pause to freeze the information. After you finish reading it, press Pause again so that your computer can continue booting.

- » **The size of the hard drive and the arrangement of its partitions:** To find out the size of your computer's hard drive in Windows 10, open the File Explorer (found in the desktop taskbar), and then right-click the drive icon and choose the Properties command from the shortcut menu that appears. (The procedure for earlier versions of Windows is similar.) Figure 4-1 shows the Properties dialog box for a 126GB disk drive that has about 115GB of free space.

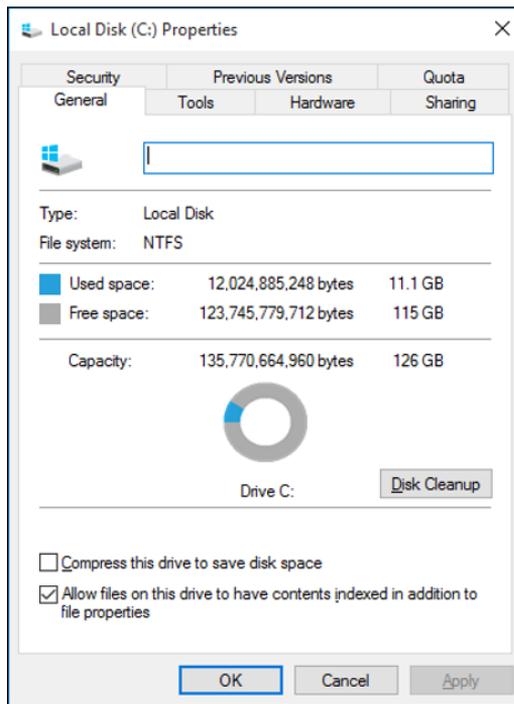


FIGURE 4-1: The Properties dialog box for a disk drive.

If your computer has more than one hard drive, Windows lists an icon for each drive in the Computer window. Jot down the size and amount of free space available on each drive.

- » **The amount of memory:** To find this information in Windows 10, open the File Explorer, right-click This PC, and choose the Properties command. The amount of memory on your computer is shown in the dialog box that appears. For example, Figure 4-2 shows the System Properties dialog box for a computer with 8GB of RAM.

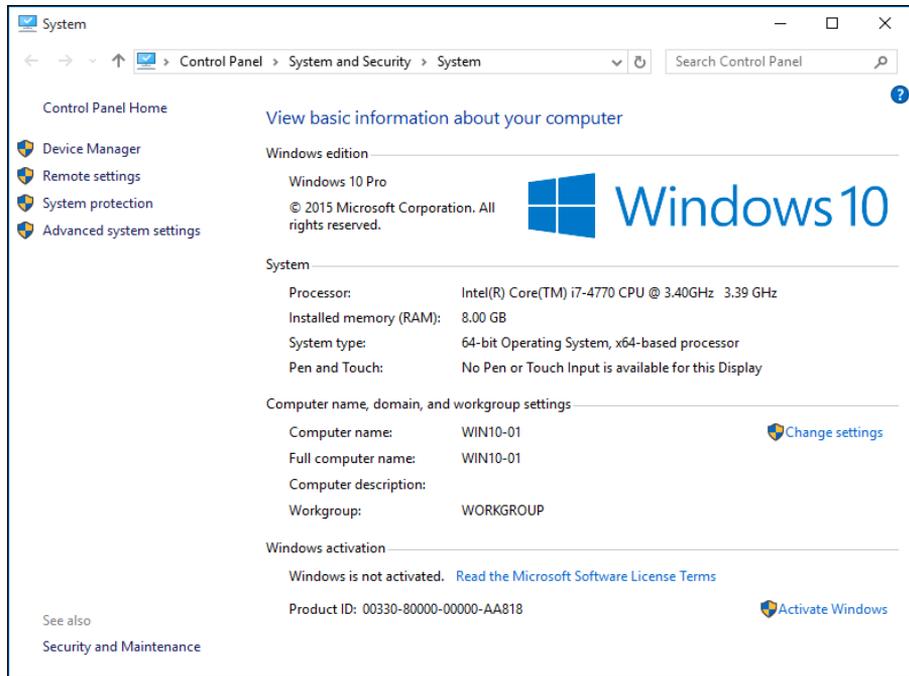


FIGURE 4-2:
The Properties
page for a
computer with
8GB of RAM.

- » **The operating system version:** This you can also deduce from the System Properties dialog box. For example, the Properties page shown in Figure 4-2 indicates that the computer is running Windows 10 Pro.
- » **What kind of printer, if any, is attached to the computer:** Usually, you can tell just by looking at the printer. You can also tell by double-clicking the Devices and Printers icon in Control Panel.
- » **Any other devices connected to the computer:** A DVD or Blu-ray drive? Scanner? External disk or tape drive? Video camera? Battle droid? Hot tub?
- » **What software is used on the computer:** Microsoft Office? AutoCAD? QuickBooks? Make a complete list and include version numbers.
- » **Does the computer have wireless capability?** Nearly all laptops do. Most desktops do not, but you can always add an inexpensive USB wireless adapter if you want your network to be entirely wireless.

Programs that gather information for you

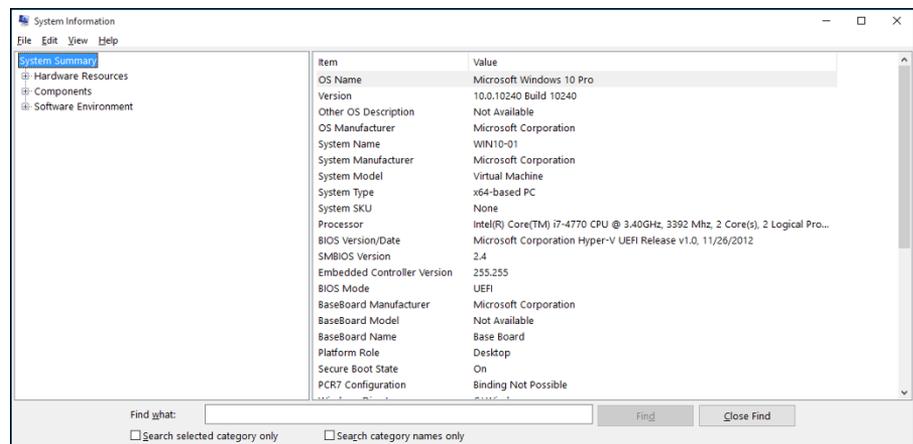
Gathering information about your computers is a lot of work if you have more than a few computers to network. Fortunately, several software programs are available that can automatically gather the information for you. These programs

inspect various aspects of a computer, such as the CPU type and speed, amount of RAM, and the size of the computer's hard drives. Then they show the information on the screen and give you the option of saving the information to a hard drive file or printing it.

Windows comes with just such a program, called Microsoft System Information. Microsoft System Information gathers and prints information about your computer. To start Microsoft System Information in Windows 10, right-click the Start button and choose Run, then type `msinfo32` and press Enter.

When you fire up Microsoft System Information, you see a window similar to the one shown in Figure 4-3. Initially, Microsoft System Information displays basic information about your computer, such as your version of Microsoft Windows, the processor type, the amount of memory on the computer, and so on. You can obtain more detailed information by clicking Hardware Resources, Components, or other categories in the left side of the window.

FIGURE 4-3:
Let the System Information program gather the data you need.



To Dedicate or Not to Dedicate: That Is the Question

One of the most basic questions that a network plan must answer is whether the network will have one or more dedicated servers or rely completely on peer-to-peer networking. If the only reason for purchasing your network is to share a printer and exchange an occasional file, you may not need a dedicated server computer. In that case, you can create a peer-to-peer network by using the computers that you already have. However, all but the smallest networks will benefit from having a separate, dedicated server computer.

- » **Using a dedicated server computer makes the network faster, easier to work with, and more reliable.** Consider what happens, though, when the user of a server computer that doubles as a workstation decides to turn off the computer, not realizing that someone else is accessing files on his hard drive.
- » **You don't necessarily have to use your biggest and fastest computer as your server computer.** I've seen networks where the slowest computer on the network is the server. This advice is especially true when the server is mostly used to share a printer or to store a small number of shared files. So if you need to buy a computer for your network, consider promoting one of your older computers to be the server and using the new computer as a client.

Assuming that your network will require one or more dedicated servers, you should next consider what types of servers the network will need. In some cases, a single server computer can fill one or more of these roles. Whenever possible, limit each server computer to a single server function.

File servers

File servers provide centralized disk storage that can be conveniently shared by client computers on the network. The most common task of a file server is to store shared files and programs. For example, members of a small workgroup can use disk space on a file server to store their Microsoft Office documents.

File servers must ensure that two users don't try to update the same file at the same time. The file servers do this by *locking* a file while a user updates the file so that other users can't access the file until the first user finishes. For document files (for example, word processing or spreadsheet files), the whole file is locked. For database files, the lock can be applied just to the portion of the file that contains the record or records being updated.

Print servers

Sharing printers is one of the main reasons that many small networks exist. Although it isn't necessary, a server computer can be dedicated for use as a *print server*, whose sole purpose is to collect information being sent to a shared printer by client computers and print it in an orderly fashion.

- » A single computer may double as both a file server and a print server, but performance is better if you use separate print and file server computers.

- » With inexpensive inkjet printers running about \$100 each, just giving each user his or her own printer is tempting. However, you get what you pay for. Instead of buying \$100 printers for 15 users, you may be better off buying one high-speed \$1,500 laser printer and sharing it. The \$1,500 laser printer will be much faster, will probably produce better-looking output, and will be less expensive to operate.

Web servers

A *web server* is a server computer that runs software that enables the computer to host an Internet website. The two most popular web server programs are Microsoft's IIS (Internet Information Services) and Apache, an open source web server managed by the Apache Software Foundation.

Mail servers

A *mail server* is a server that handles the network's email needs. It is configured with email server software, such as Microsoft Exchange Server. Exchange Server is designed to work with Microsoft Outlook, the email client software that comes with Microsoft Office.

Most mail servers actually do much more than just send and receive email. For example, here are some of the features that Exchange Server offers beyond simple email:

- » Collaboration features that simplify the management of collaborative projects.
- » Audio and video conferencing.
- » Chat rooms and instant messaging (IM) services.
- » Microsoft Exchange Forms Designer, which lets you develop customized forms for applications, such as vacation requests or purchase orders.

Database servers

A *database server* is a server computer that runs database software, such as Microsoft's SQL Server 2014. Database servers are usually used along with customized business applications, such as accounting or marketing systems.

Application servers

An *application server* is a server computer that runs a specific application. For example, you might use an accounting application that requires its own server. In that case, you'll need to dedicate a server to the accounting application.

License servers

Some organizations use software that requires licenses that are distributed from a centralized license server. For example, engineering firms often use computer-aided design (CAD) software such as AutoCAD that requires a license server. In that case, you'll need to set up a server to handle the licensing function.

Choosing a Server Operating System

If you determine that your network will require one or more dedicated servers, the next step is to determine what network operating system those servers should use. If possible, all the servers should use the same network operating system (NOS) so that you don't find yourself supporting different operating systems.

Although you can choose from many network operating systems, from a practical point of view, your choices are limited to the following:

- » Windows Server 2016 or 2012
- » Linux or another version of Unix

For more information, see Chapter 11.

Planning the Infrastructure

You also need to plan the details of how you will connect the computers in the network. This task includes determining which network topology the network will use, what type of cable will be used, where the cable will be routed, and what other devices (such as repeaters, bridges, hubs, switches, and routers) will be needed.

Although you have many cabling options to choose from, you'll probably use Cat 5e or better UTP (unshielded twisted pair) for most — if not all — of the desktop

client computers on the network. However, you have many decisions to make beyond this basic choice:

- » Will you use inexpensive consumer-grade network switches such as those you can buy at a consumer electronics store or an office supply store, or will you want professional-grade switches, which are more expensive but provide advanced management features?
- » Where will you place the switch — on a desktop somewhere within the group or in a central wiring closet?
- » How many client computers and other devices will you place on each switch, and how many switches will you need to support all of these computers and other devices?
- » If you need more than one switch, what type of cabling will you use to connect the switches to one another?

For more information about network cabling, see Chapter 6.



TIP

If you're installing new network cable, don't scrimp on the cable itself. Because installing network cable is a labor-intensive task, the cost of the cable itself is a small part of the total cable installation cost. And if you spend a little extra to install higher-grade cable now, you won't have to replace the cable in a few years when it's time to upgrade the network.

Drawing Diagrams

One of the most helpful techniques for creating a network plan is to draw a picture of it. The diagram can be a detailed floor plan, showing the actual location of each network component. This type of diagram is sometimes called a “physical map.” If you prefer, the diagram can be a *logical map*, which is more abstract and Picasso-like. Any time you change the network layout, update the diagram. Also include a detailed description of the change, the date that the change was made, and the reason for the change.

You can diagram very small networks on the back of a napkin, but if the network has more than a few computers, you'll want to use a drawing program to help you create the diagram. You can use a professional drawing program such as Microsoft Visio, but you can also find simpler and less expensive online drawing tools to help you document your network. Figure 4-4 shows a network diagram drawn with an inexpensive online tool called Lucidchart (www.lucidchart.com).

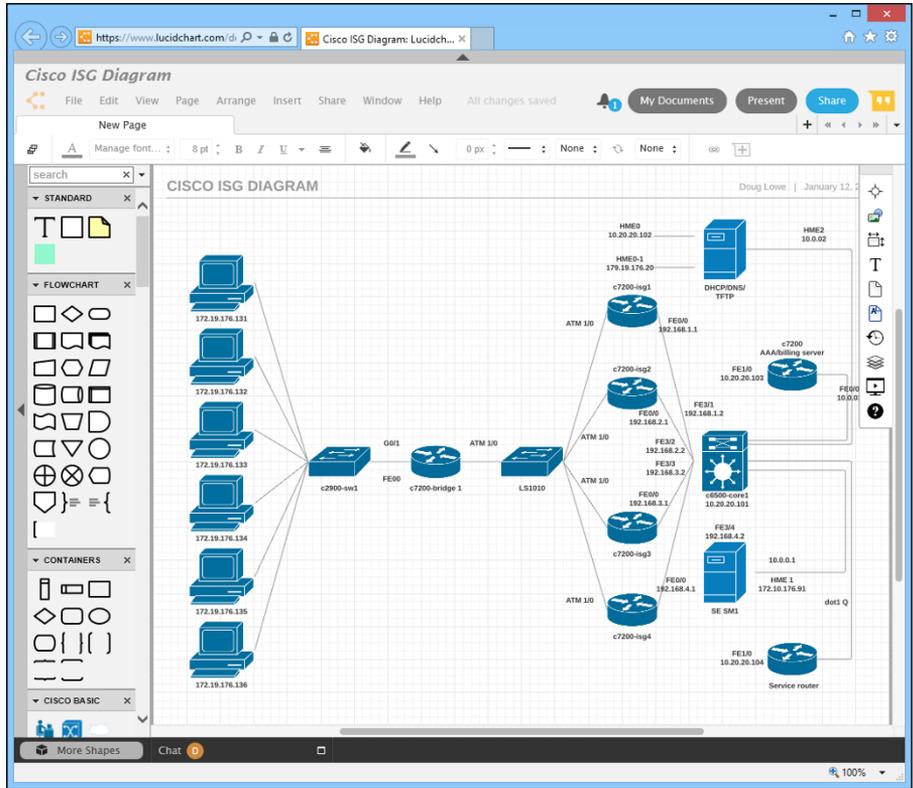


FIGURE 4-4: Using Lucidchart to draw a network diagram.

IN THIS CHAPTER

Getting a handle (or two) on the binary system

Digging into IP addresses

Finding out how subnetting works

Understanding private and public IP addresses

Looking at network address translation

Finding out how DHCP works

Understanding how DNS works

Chapter 5

Dealing with TCP/IP

Transfer Control Protocol/Internet Protocol — TCP/IP — is the basic protocol by which computers on a network talk to each other. Without TCP/IP, networks wouldn't work. In this chapter, I introduce you to the most important concepts of TCP/IP.



WARNING

This chapter is far and away the most technical chapter in this book. It helps you examine the binary system, the details of how IP addresses are constructed, how subnetting works, and how two of the most important TCP/IP services — Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) — work. You don't need to understand every detail in this chapter to set up a simple TCP/IP network. However, the more you understand the information in this chapter, the more TCP/IP will start to make sense. Be brave.

Understanding Binary

Before you can understand the details of how TCP/IP — in particular, IP — addressing works, you need to understand how the binary numbering system

works because binary is the basis of IP addressing. If you already understand binary, please skip right over this section to the next main section, “Introducing IP Addresses.” I don’t want to bore you with stuff that’s too basic.

Counting by ones

The *binary* counting system uses only two numerals: 0 and 1. In the decimal system to which most people are accustomed, you use ten numerals: 0 through 9. In an ordinary decimal number, such as 3,482, the rightmost digit represents ones; the next digit to the left, tens; the next, hundreds; the next, thousands; and so on. These digits represent powers of ten: first 10^0 (which is 1); next, 10^1 (10); then 10^2 (100); then 10^3 (1,000); and so on.

In binary, you have only two numerals rather than ten, which is why binary numbers look somewhat monotonous, as in 110011, 101111, and 100001.

The positions in a binary number (called *bits* rather than *digits*) represent powers of two rather than powers of ten — working from right to left, each bit represents the decimal values 1, 2, 4, 8, 16, 32, and so on. To figure the decimal value of a binary number, you multiply each bit by its corresponding power of two and then add the results. The decimal value of binary 10111, for example, is calculated as follows:

$$\begin{array}{r} 1 \times 2^0 = 1 \times 1 = 1 \\ + 1 \times 2^1 = 1 \times 2 = 2 \\ + 1 \times 2^2 = 1 \times 4 = 4 \\ + 0 \times 2^3 = 0 \times 8 = 0 \\ + 1 \times 2^4 = 1 \times 16 = \underline{16} \\ \hline 23 \end{array}$$

Fortunately, a computer is good at converting a number between binary and decimal — so good, in fact, that you’re unlikely ever to need to do any conversions yourself. The point of knowing binary isn’t to be able to look at a number, such as 1110110110110, and say instantly, “Ah! Decimal 7,606!” (If you could do that, you would probably be interviewed on the *Today* show, and they would even make a movie about you.)

Instead, the point is to have a basic understanding of how computers store information and — most important — to understand how the hexadecimal counting system works (which I describe in the following section).

Here are some of the more interesting characteristics of binary and how the system is similar to and differs from the decimal system:



TIP

- » The number of bits allotted for a binary number determines how large that number can be. If you allot eight bits, the largest value that number can store is 11111111, which happens to be 255 in decimal.
- » To quickly determine how many different values you can store in a binary number of a given length, use the number of bits as an exponent of two. An eight-bit binary number, for example, can hold 2^8 values. Because 2^8 is 256, an 8-bit number can have any of 256 different values, which is why a byte, which is eight bits, can have 256 different values.
- » This powers-of-two concept is why computers don't use nice, even, round numbers in measuring such values as memory or disk space. A value of 1K, for example, isn't an even 1,000 bytes — it's 1,024 bytes because 1,024 is 2^{10} . Similarly, 1MB isn't an even 1,000,000 bytes but rather is 1,048,576 bytes, which happens to be 2^{20} .

Doing the logic thing

One of the great things about binary is that it's very efficient at handling special operations called *logical operations*. Four basic logical operations exist, although additional operations are derived from the basic four operations. Three of the operations — AND, OR, and XOR — compare two binary digits (bits). The fourth (NOT) works on just a single bit.

The following list summarizes the basic logical operations:

- » **AND:** An AND operation compares two binary values. If both values are 1, the result of the AND operation is 1. If one value is 0 or both of the values are 0, the result is 0.
- » **OR:** An OR operation compares two binary values. If at least one of the values is 1, the result of the OR operation is 1. If both values are 0, the result is 0.
- » **XOR:** An XOR operation compares two binary values. If exactly one of them is 1, the result is 1. If both values are 0 or if both values are 1, the result is 0.
- » **NOT:** The NOT operation doesn't compare two values. Instead, it simply changes the value of a single binary value. If the original value is 1, NOT returns 0. If the original value is 0, NOT returns 1.



TIP

Logical operations are applied to binary numbers that have more than one binary digit by applying the operation one bit at a time. The easiest way to do this manually is to

1. Line one of the two binary numbers on top of the other.
2. Write the result of the operation beneath each binary digit.

The following example shows how you calculate 10010100 AND 11001101:

```
10010100
AND 11001101
10000100
```

As you can see, the result is 10000100.

Introducing IP Addresses

An *IP address* is a number that uniquely identifies every host on an IP network. IP addresses operate at the Network layer of the TCP/IP protocol stack, so they're independent of lower-level addresses, such as MAC addresses. (MAC stands for *Media Access Control*.)

IP addresses are 32-bit binary numbers, which means that theoretically, a maximum of something in the neighborhood of 4 billion unique host addresses can exist throughout the Internet. You'd think that'd be enough, but TCP/IP places certain restrictions on how IP addresses are allocated. These restrictions severely limit the total number of usable IP addresses, and about half of the total available IP addresses have already been assigned. However, new techniques for working with IP addresses have helped to alleviate this problem, and a new standard for 128-bit IP addresses (known as *IPv6*) is on the verge of winning acceptance.

Networks and hosts

The primary purpose of Internet Protocol (IP) is to enable communications between networks. As a result, a 32-bit IP address consists of two parts:

- » **The network ID (or network address):** Identifies the network on which a host computer can be found
- » **The host ID (or host address):** Identifies a specific device on the network indicated by the network ID

Most of the complexity of working with IP addresses has to do with figuring out which part of the complete 32-bit IP address is the network ID and which part is the host ID. The original IP specification uses the *address classes* system to determine which part of the IP address is the network ID and which part is the host ID.

A newer system — classless IP addresses — is rapidly taking over the address classes system. You come to grips with both systems later in this chapter.

The dotted-decimal dance

IP addresses are usually represented in a format known as *dotted-decimal notation*. In dotted-decimal notation, each group of eight bits — an *octet* — is represented by its decimal equivalent. For example, consider the following binary IP address:

```
11000000101010001000100000011100
```

The dotted-decimal equivalent to this address is

```
192.168.136.28
```

Here, 192 represents the first eight bits (11000000); 168, the second set of eight bits (10101000); 136, the third set of eight bits (10001000); and 28, the last set of eight bits (00011100). This is the format in which you usually see IP addresses represented.

Classifying IP Addresses

When the original designers of the IP protocol created the IP addressing scheme, they could have assigned an arbitrary number of IP address bits for the network ID. The remaining bits would then be used for the host ID. For example, the designers may have decided that half of the address (16 bits) would be used for the network and the remaining 16 bits would be used for the host ID. The result of that scheme would be that the Internet could have a total of 65,536 networks, and each of those networks could have 65,536 hosts.

In the early days of the Internet, this scheme probably seemed like several orders of magnitude more than would ever be needed. However, the IP designers realized from the start that few networks would actually have tens of thousands of hosts. Suppose that a network of 1,000 computers joins the Internet and is assigned one of these hypothetical network IDs. Because that network uses only 1,000 of its 65,536 host addresses, more than 64,000 IP addresses would be wasted.

As a solution to this problem, the idea of IP address classes was introduced. The IP protocol defines five different address classes: A, B, C, D, and E. Each of the first three classes, A–C, uses a different size for the network ID and host ID portion of the address. Class D is for a special type of address called a *multicast address*. Class E is an experimental address class that isn't used.

The first four bits of the IP address are used to determine into which class a particular address fits:

- » If the first bit is 0, the address is a Class A address.
- » If the first bit is 1 and the second bit is 0, the address is a Class B address.
- » If the first two bits are both 1 and the third bit is 0, the address is a Class C address.
- » If the first three bits are all 1 and the fourth bit is 0, the address is a Class D address.
- » If the first four bits are all 1, the address is a Class E address.

Because Class D and E addresses are reserved for special purposes, I focus the rest of this discussion on Class A, B, and C addresses. Table 5-1 summarizes the details of each address class.

TABLE 5-1 IP Address Classes

Class	Address Range	Starting Bits	Length of Network ID	Number of Networks	Number of Hosts
A	1–126 .x.y.z	0	8	126	16,777,214
B	128–191 .x.y.z	10	16	16,384	65,534
C	192–223 .x.y.z	110	24	2,097,152	254

Class A addresses

Class A addresses are designed for very large networks. In a Class A address, the first octet of the address is the network ID, and the remaining three octets are the host ID. Because only eight bits are allocated to the network ID and the first of these bits is used to indicate that the address is a Class A address, only 126 Class A networks can exist in the entire Internet. However, each Class A network can accommodate more than 16 million hosts.



TECHNICAL
STUFF

Only about 40 Class A addresses are assigned to companies or organizations. The rest are either reserved for use by the Internet Assigned Numbers Authority (IANA) or are assigned to organizations that manage IP assignments for geographic regions, such as Europe, Asia, and Latin America.

Just for fun, Table 5-2 lists some of the better-known Class A networks. You probably recognize many of them. In case you're interested, you can find a complete list of all the Class A address assignments at www.iana.org/assignments/ipv4-address-space.

TABLE 5-2 Some Well-Known Class A Networks

Network	Description	Network	Description
3	General Electric Company	20	Computer Sciences Corporation
4	Level 3 Communications	22	Defense Information Systems Agency
6	Army Information Systems Center	25	UK Ministry of Defense
8	Level-3 Communications	26	Defense Information Systems Agency
9	IBM	28	Decision Sciences Institute (North)
11	DoD Intel Information Systems	29-30	Defense Information Systems Agency
12	AT&T Bell Laboratories	33	DLA Systems Automation Center
15	Hewlett-Packard Company	44	Amateur Radio Digital Communications
16	Digital Equipment Corporation	48	Prudential Securities, Inc.
17	Apple Computer Inc.	53	Daimler
18	MIT	55	DoD Network Information Center
19	Ford Motor Company	56	U.S. Postal Service

Class B addresses

In a Class B address, the first two octets of the IP address are used as the network ID, and the second two octets are used as the host ID. Thus, a Class B address comes close to my hypothetical scheme of splitting the address down the middle, using half for the network ID and half for the host ID. It isn't identical to this scheme, however, because the first two bits of the first octet are required to be 10, to indicate that the address is a Class B address. Thus, a total of 16,384 Class B networks can exist. All Class B addresses fall within the range 128.x.y.z to 191.x.y.z. Each Class B address can accommodate more than 65,000 hosts.



The problem with Class B networks is that even though they're much smaller than Class A networks, they still allocate far too many host IDs. Very few networks have tens of thousands of hosts. Thus, the careless assignment of Class B addresses can lead to a large percentage of the available host addresses being wasted on organizations that don't need them.

Class C addresses

In a Class C address, the first three octets are used for the network ID, and the fourth octet is used for the host ID. With only eight bits for the host ID, each Class C network can accommodate only 254 hosts. However, with 24 network ID bits, Class C addresses allow for more than 2 million networks.

WHAT ABOUT IPV6?

Most of the current Internet is based on version 4 of the Internet Protocol, also known as IPv4. IPv4 has served the Internet well for more than 20 years. However, the growth of the Internet has put a lot of pressure on IPv4's limited 32-bit address space. This chapter describes how IPv4 has evolved to make the best possible use of 32-bit addresses, but eventually all the addresses will be assigned; the IPv4 address space will be filled to capacity. When that happens, the Internet will have to migrate to the next version of IP, known as IPv6.

IPv6 is also called *IP next generation*, or *IPng*, in honor of the favorite television show of most Internet gurus, *Star Trek: The Next Generation*.

IPv6 offers several advantages over IPv4, but the most important is that it uses 128 bits for Internet addresses rather than 32 bits. The number of host addresses possible with 128 bits is a number so large that it would make Carl Sagan proud. It doesn't just double or triple the number of available addresses. Just for the fun of it, here's the number of unique Internet addresses provided by IPv6:

340,282,366,920,938,463,463,374,607,431,768,211,456

This number is so large that it defies understanding. If the IANA had been around at the creation of the universe and started handing out IPv6 addresses at a rate of one per millisecond, it would now, 15 billion years later, have not yet allocated even 1 percent of the available addresses.

Unfortunately, the transition from IPv4 to IPv6 has been a slow one. Thus, the Internet will continue to be driven by IPv4 for at least a few more years.



The problem with Class C networks is that they're too small. Although few organizations need the tens of thousands of host addresses provided by a Class B address, many organizations need more than a few hundred. The large discrepancy between Class B networks and Class C networks led to the development of subnetting, which I describe in the next section.

Subnetting

Subnetting is a technique that lets network administrators use the 32 bits available in an IP address more efficiently by creating networks that aren't limited to the scales provided by Class A, B, and C IP addresses. With subnetting, you can create networks with more realistic host limits.

Subnetting provides a more flexible way to designate which portion of an IP address represents the network ID and which portion represents the host ID. With standard IP address classes, only three possible network ID sizes exist: 8 bits for Class A, 16 bits for Class B, and 24 bits for Class C. Subnetting lets you select an arbitrary number of bits to use for the network ID.

Two reasons compel me to use subnetting. The first is to allocate the limited IP address space more efficiently. If the Internet were limited to Class A, B, or C addresses, every network would be allocated 254, 65,000, or 16 million IP addresses for host devices. Although many networks with more than 254 devices exist, few (if any) exist with 65,000, let alone 16 million. Unfortunately, any network with more than 254 devices would need a Class B allocation and probably waste tens of thousands of IP addresses.

The second reason for subnetting is that even if a single organization has thousands of network devices, operating all those devices with the same network ID would slow the network to a crawl. The way TCP/IP works dictates that all the computers with the same network ID must be on the same physical network. The physical network comprises a single *broadcast domain*, which means that a single network medium must carry all the traffic for the network. For performance reasons, networks are usually segmented into broadcast domains that are smaller than even Class C addresses provide.

Subnets

A *subnet* is a network that falls within another (Class A, B, or C) network. Subnets are created by using one or more of the Class A, B, or C host bits to extend the network ID. Thus, rather than the standard 8-, 15-, or 24-bit network ID, subnets can have network IDs of any length.

Figure 5-1 shows an example of a network before and after subnetting has been applied. In the unsubnetted network, the network has been assigned the Class B address 144.28.0.0. All the devices on this network must share the same broadcast domain.

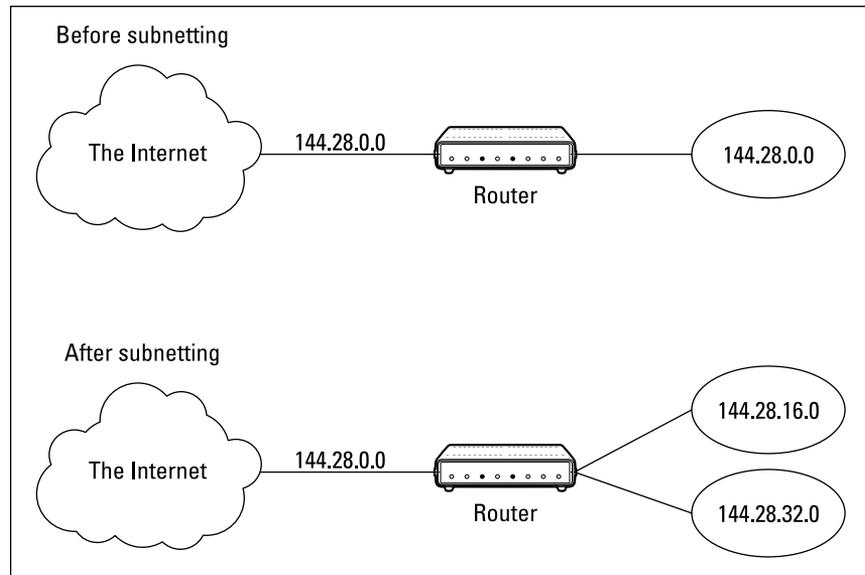


FIGURE 5-1:
A network before
and after
subnetting.

In the second network, the first four bits of the host ID are used to divide the network into two small networks, identified as subnets 16 and 32. To the outside world (that is, on the other side of the router), these two networks still appear to be a single network identified as 144.28.0.0. For example, the outside world considers the device at 144.28.16.22 to belong to the 144.28.0.0 network. As a result, a packet sent to this device is delivered to the router at 144.28.0.0. The router then considers the subnet portion of the host ID to decide whether to route the packet to subnet 16 or subnet 32.

Subnet masks

For subnetting to work, the router must be told which portion of the host ID to use for the subnet's network ID. This little sleight of hand is accomplished by using another 32-bit number, known as a *subnet mask*. Those IP address bits that represent the network ID are represented by a 1 in the mask, and those bits that represent the host ID appear as a 0 in the mask. As a result, a subnet mask always has a consecutive string of ones on the left, followed by a string of zeros.

For example, the subnet mask for the subnet, as shown in Figure 5-1, in which the network ID consists of the 16-bit network ID plus an additional 4-bit subnet ID, would look like this:

```
11111111 11111111 11110000 00000000
```

In other words, the first 20 bits are ones; the remaining 12 bits are zeros. Thus, the complete network ID is 20 bits in length, and the actual host ID portion of the subnetted address is 12 bits in length.

To determine the network ID of an IP address, the router must have both the IP address and the subnet mask. The router then performs a bitwise operation called a *logical AND* on the IP address to extract the network ID. To perform a logical AND, each bit in the IP address is compared to the corresponding bit in the subnet mask. If both bits are 1, the resulting bit in the network ID is set to 1. If either of the bits is 0, the resulting bit is set to 0.

For example, here's how the network address is extracted from an IP address using the 20-bit subnet mask from the previous example:

```
          144 .   28 .   16 .   17
IP address: 10010000 00011100 00100000 00001001
Subnet mask: 11111111 11111111 11110000 00000000
Network ID: 10010000 00011100 00100000 00000000
          144 .   28 .   16 .    0
```

Thus, the network ID for this subnet is 144.28.16.0.

The subnet mask itself is usually represented in dotted-decimal notation. As a result, the 20-bit subnet mask used in the previous example would be represented as 255.255.240.0:

```
Subnet mask: 11111111 11111111 11110000 00000000
            255 .   255 .   240 .    0
```



TIP

Don't confuse a subnet mask with an IP address. A subnet mask doesn't represent any device or network on the Internet. It's just a way of indicating which portion of an IP address should be used to determine the network ID. (You can spot a subnet mask right away because the first octet is always 255, and 255 isn't a valid first octet for any class of IP address.)

The great subnet roundup

You should know about a few additional restrictions that are placed on subnet masks — in particular:

- » The minimum number of network ID bits is eight. As a result, the first octet of a subnet mask is always 255.
- » The maximum number of network ID bits is 30. You have to leave at least two bits for the host ID portion of the address, to allow for at least two hosts. If you used all 32 bits for the network ID, that would leave no bits for the host ID. Obviously, that doesn't work. Leaving just one bit for the host ID doesn't work, either. That's because a host ID of all ones is reserved for a broadcast address — and all zeros refers to the network itself. Thus, if you used 31 bits for the network ID and left only one for the host ID, host ID 1 would be used for the broadcast address and host ID 0 would be the network itself, leaving no room for actual hosts. That's why the maximum network ID size is 30 bits.
- » Because the network ID is always composed of consecutive bits set to 1, only nine values are possible for each octet of a subnet mask (including counting 0). For your reference, these values are listed in Table 5-3.

TABLE 5-3 The Eight Subnet Octet Values

Binary Octet	Decimal	Binary Octet	Decimal
00000000	0	11111000	248
10000000	128	11111100	252
11000000	192	11111110	254
11100000	224	11111111	255
11110000	240		

Private and public addresses

Any host with a direct connection to the Internet must have a globally unique IP address. However, not all hosts are connected directly to the Internet. Some are on networks that aren't connected to the Internet. Some hosts are hidden behind firewalls, so their Internet connection is indirect.

Several blocks of IP addresses are set aside just for this purpose — for use on private networks that aren't connected to the Internet or to use on networks hidden behind

a firewall. Three such ranges of addresses exist, as summarized in Table 5-4. Whenever you create a private TCP/IP network, use IP addresses from one of these ranges.

TABLE 5-4 Private Address Spaces

Subnet Mask	Address Range
255.0.0.0	10.0.0.1–10.255.255.254
255.255.240.0	172.16.1.1–172.31.255.254
255.255.0.0	192.168.0.1–192.168.255.254

Understanding Network Address Translation

Many firewalls use a technique called *network address translation* (NAT) to hide the actual IP address of a host from the outside world. When that's the case, the NAT device must use a globally unique IP to represent the host to the Internet; behind the firewall, however, the host can use any IP address it wants. As packets cross the firewall, the NAT device translates the private IP address to the public IP address, and vice versa.

One of the benefits of NAT is that it helps slow down the rate at which the IP address space is assigned because a NAT device can use a single public IP address for more than one host. It does this by keeping track of outgoing packets so that it can match up incoming packets with the correct host. To understand how this process works, consider this sequence of steps:

1. A host whose private address is 192.168.1.100 sends a request to 216.58.192.4, which happens to be `www.google.com`. The NAT device changes the source IP address of the packet to 208.23.110.22, the IP address of the firewall. That way, Google will send its reply back to the firewall router. The NAT records that 192.168.1.100 sent a request to 216.58.192.4.
2. Now another host, at address 192.168.1.107, sends a request to 23.54.240.121, which happens to be `www.microsoft.com`. The NAT device changes the source of this request to 208.23.110.22 so that Microsoft will reply to the firewall router. The NAT records that 192.168.1.107 sent a request to 23.54.240.121.

3. A few seconds later, the firewall receives a reply from 216.58.192.4. The destination address in the reply is 208.23.110.22, the address of the firewall. To determine to whom to forward the reply, the firewall checks its records to see who's waiting for a reply from 216.58.192.4. It discovers that 192.168.1.100 is waiting for that reply, so it changes the destination address to 192.168.1.100 and sends the packet on.

Actually, the process is a little more complicated than that because it's very likely that two or more users may have pending requests from the same public IP. In that case, the NAT device uses other techniques to figure out to which user each incoming packet should be delivered.

Configuring Your Network for DHCP

Every host on a TCP/IP network must have a unique IP address. Each host must be properly configured so that it knows its IP address. When a new host comes online, it must be assigned an IP address within the correct range of addresses for the subnet — one that's not already in use. Although you can manually assign IP addresses to each computer on your network, that task quickly becomes overwhelming if the network has more than a few computers.

That's where Dynamic Host Configuration Protocol (DHCP) comes into play. DHCP automatically configures the IP address for every host on a network, thus ensuring that each host has a valid, unique IP address. DHCP even automatically reconfigures IP addresses as hosts come and go. As you can imagine, DHCP can save a network administrator many hours of tedious configuration work.

In this section, you discover the ins and outs of DHCP: what it is, how it works, and how to set it up.

Understanding DHCP

DHCP allows individual computers on a TCP/IP network to obtain their configuration information — in particular, their IP addresses — from a server. The DHCP server keeps track of which IP addresses have already been assigned so that when a computer requests an IP address, the DHCP servers offer it an IP address that isn't already in use.

The alternative to DHCP is to assign each computer on your network a static IP address, which can be good or problematic:



WARNING

- » Static IP addresses are okay for networks with a handful of computers.
- » For networks with more than a few computers, using static IP addresses is a huge mistake. Eventually, some poor, harried administrator (guess who?) will make the mistake of assigning two computers the same IP address. Then you have to manually check each computer's IP address to find the conflict. DHCP is a must for any but the smallest networks.

Although the primary job of DHCP is to assign IP addresses, DHCP provides more configuration information than just the IP address to its clients. The additional configuration information is referred to as *DHCP options*. The following list describes some common DHCP options that can be configured by the server:

- » Router address, also known as the default gateway address
- » Expiration time for the configuration information
- » Domain name
- » DNS server address
- » Windows Internet Name Service (WINS) server address

DHCP servers

A DHCP server can be a server computer located on the TCP/IP network. Fortunately, all modern server operating systems have a built-in DHCP server capability. To set up DHCP on a network server, all you have to do is enable the server's DHCP function and configure its settings. In the section "Managing a Windows Server 2016 DHCP Server," later in this chapter, I show you how to configure a DHCP server for Windows 2016.

A server computer running DHCP doesn't have to be devoted entirely to DHCP unless the network is very large. For most networks, a file server can share duty as a DHCP server, especially if you provide long leases for your IP addresses. (I explain the idea of leases later in this chapter.)

Most multifunction routers also have built-in DHCP servers. So if you don't want to burden one of your network servers with the DHCP function, you can enable the router's built-in DHCP server. An advantage of allowing the router to be your network's DHCP server is that you rarely need to power down a router. By contrast, you occasionally need to restart or power down a file server to perform system maintenance, to apply upgrades, or to do some needed troubleshooting.



TIP

Most networks require only one DHCP server. Setting up two or more servers on the same network requires that you carefully coordinate the IP address ranges (known as *scopes*) for which each server is responsible. If you accidentally set up two DHCP servers for the same scope, you may end up with duplicate address assignments if the servers attempt to assign the same IP address to two different hosts. To prevent this situation from happening, set up just one DHCP server unless your network is so large that one server can't handle the workload.

Understanding scopes

A *scope* is simply a range of IP addresses that a DHCP server is configured to distribute. In the simplest case, in which a single DHCP server oversees IP configuration for an entire subnet, the scope corresponds to the subnet. However, if you set up two DHCP servers for a subnet, you can configure each one with a scope that allocates only one part of the complete subnet range. In addition, a single DHCP server can serve more than one scope.

You must create a scope before you can enable a DHCP server. When you create a scope, you can provide it these properties:

- » A **scope name**, which helps you identify the scope and its purpose.
- » A **scope description**, which lets you provide additional details about the scope and its purpose.
- » A **starting IP address** for the scope.
- » An **ending IP address** for the scope.
- » A **subnet mask** for the scope. You can specify the subnet mask with dotted decimal notation or with Classless Inter Domain Routing (CIDR) notation.
- » **One or more ranges of excluded addresses.** These addresses aren't assigned to clients. (For more information, see the section "Feeling excluded?" later in this chapter.)
- » **One or more reserved addresses.** These addresses are always assigned to particular host devices. (For more information, see the section "Reservations suggested," later in this chapter.)
- » The **lease duration**, which indicates how long the host is allowed to use the IP address. The client attempts to renew the lease when half of the lease duration has elapsed. For example, if you specify a lease duration of eight days, the client attempts to renew the lease after four days have passed. The host then has plenty of time to renew the lease before the address is reassigned to some other host.



TIP

- » The **router address** for the subnet.
This value is also known as the *default gateway address*.
- » The **domain name and the IP address** of the network's DNS servers and WINS servers.

Feeling excluded?

Everyone feels excluded once in a while. With a wife and three daughters, I know how that feels. Sometimes, however, being excluded is a good thing. In the case of DHCP scopes, exclusions can help you prevent IP address conflicts and can enable you to divide the DHCP workload for a single subnet among two or more DHCP servers.

An *exclusion* is a range of addresses not included in a scope but falling within the range of the scope's starting and ending addresses. In effect, an exclusion range lets you punch a hole in a scope: The IP addresses that fall within the hole aren't assigned.

Here are a couple of reasons to exclude IP addresses from a scope:

- » **The computer that runs the DHCP service itself must usually have a static IP address assignment.** As a result, the address of the DHCP server should be listed as an exclusion.
- » **You may want to assign static IP addresses to your other servers.** In that case, each server IP address should be listed as an exclusion.



TIP

Reservations are often better solutions to this problem, as I describe in the next section.

Reservations suggested

In some cases, you may want to assign a specific IP address to a particular host. One way to do this is to configure the host with a static IP address so that the host doesn't use DHCP to obtain its IP configuration. However, two major disadvantages to that approach exist:

- » **TCP/IP configuration supplies more than just the IP address.** If you use static configuration, you must manually specify the subnet mask, default gateway address, DNS server address, and other configuration information required by the host. If this information changes, you have to change it not only at the DHCP server, but also at each host that you configured statically.

» You must remember to exclude the static IP address from the DHCP server's scope. Otherwise, the DHCP server doesn't know about the static address and may assign it to another host. Then comes the problem: You have two hosts with the same address on your network.



TIP

A better way to assign a fixed IP address to a particular host is to create a DHCP reservation. A *reservation* simply indicates that whenever a particular host requests an IP address from the DHCP server, the server should provide it the address that you specify in the reservation. The host doesn't receive the IP address until the host requests it from the DHCP server, but whenever the host does request IP configuration, it always receives the same address.

To create a reservation, you associate the IP address that you want assigned to the host with the host's MAC address. Accordingly, you need to get the MAC address from the host before you create the reservation:

» Usually, you can get the MAC address by running the command `ipconfig /all` from a command prompt.

» If TCP/IP has not yet been configured on the computer, you can get the MAC address by choosing the System Information command:

Choose Start ⇨ All Programs ⇨ Accessories ⇨ System Tools ⇨ System Information.



TIP

If you set up more than one DHCP server, be sure to specify the same reservations on each server. If you forget to repeat a reservation on one of the servers, that server may assign the address to another host.

How long to lease?

One of the most important decisions that you make when you configure a DHCP server is the length of time to specify for the lease duration. The default value is eight days, which is appropriate in many cases. However, you may encounter situations in which a longer or shorter interval may be appropriate:

» The more stable your network, the longer the lease duration can safely exist. If you only periodically add new computers to your network (or replace existing computers), you can safely increase the lease duration past eight days.

» The more volatile the network, the shorter the lease duration should be. For example, you may have a wireless network in a university library, used by students who bring their laptop computers into the library to work for a few hours at a time. For this network, a duration as short as one hour may be appropriate.



WARNING

Don't configure your network to allow leases of infinite duration. Although some administrators feel that this duration cuts down the workload for the DHCP server on stable networks, no network is permanently stable. Whenever you find a DHCP server that's configured with infinite leases, look at the active leases. I guarantee that you'll find IP leases assigned to computers that no longer exist.

Managing a Windows Server 2016 DHCP Server

The exact steps to follow when you configure and manage a DHCP server depend on the network operating system or router you're using. The following paragraphs describe how to work with a DHCP server in Windows Server 2016. The procedures for other operating systems are similar.

If you haven't already installed the DHCP server on the server, install it using the Server Manager (click Server Manager in the task bar, and then use Add Roles and Features to add the DHCP role). Once the DHCP server role is installed, you can manage it by opening the DHCP management console, as shown in Figure 5-2. To open this console, open System Manager and choose Tools ⇄ DHCP.

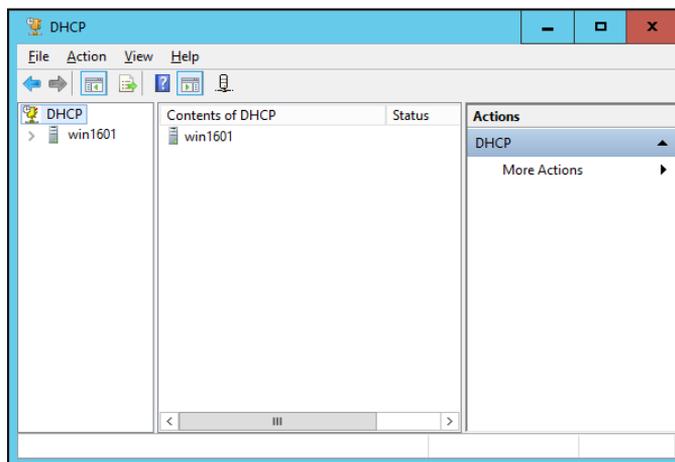


FIGURE 5-2:
The DHCP management console.

To get started with a DHCP server, you must create at least one scope. You can create a scope by using the New Scope Wizard, which you start by selecting the

server you want to create the scope on and then clicking New Scope. The wizard asks for the essential information required to define the scope, including the scope's name, its starting and ending IP addresses, and the subnet mask. You can also specify any IP addresses you want to exclude from the scope, the lease duration (the default is eight days), the IP address of your gateway router, the domain name for your network, and the IP addresses for the DNS servers you want the client computers to use. Figure 5-3 shows the New Scope Wizard in action.

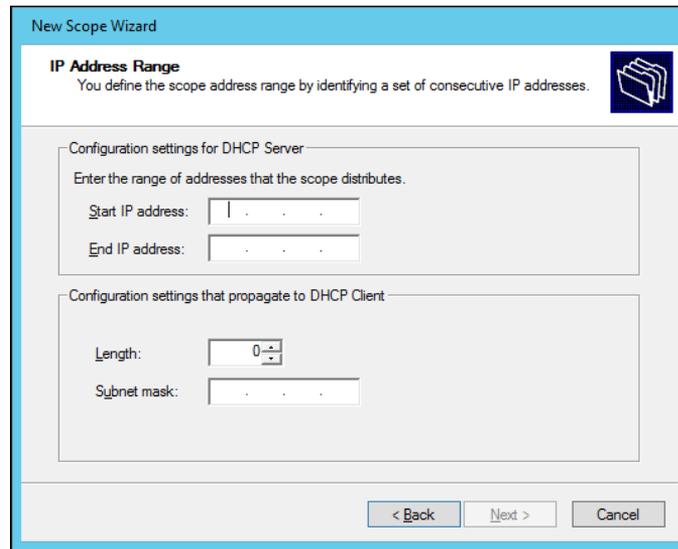


FIGURE 5-3:
The New Scope
Wizard.

Configuring a Windows DHCP Client

Configuring a Windows client for DHCP is easy. The DHCP client is included automatically when you install the TCP/IP protocol, so all you have to do is configure TCP/IP to use DHCP. To do this, open the Network Properties dialog box by choosing Network or Network Connections in the Control Panel (depending on which version of Windows the client is running). Then select the TCP/IP protocol and click the Properties button. This action opens the TCP/IP Properties dialog box, as shown in Figure 5-4. To configure the computer to use DHCP, select the Obtain an IP Address Automatically and Obtain DNS Server Address Automatically options. Click OK, and you're done.

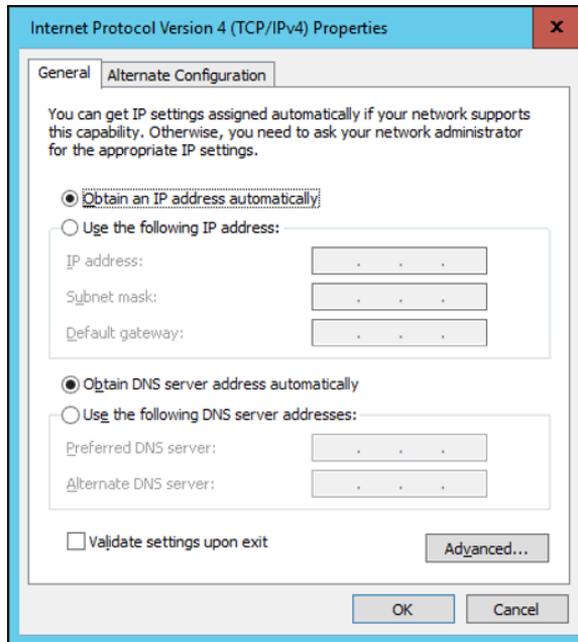


FIGURE 5-4:
Configuring a
Windows client to
use DHCP.

Using DNS

DNS (Domain Name System) is the TCP/IP facility that lets you use names rather than numbers to refer to host computers. Without DNS, you'd buy books from 54.239.25.208 rather than from `www.amazon.com`, you'd sell your used furniture at 23.4.21.125 rather than on `www.ebay.com`, and you'd search the web at 172.217.0.68 rather than at `www.google.com`.

Understanding how DNS works and how to set up a DNS server is crucial to setting up and administering a TCP/IP network. The rest of this chapter introduces you to the basics of DNS, including how the DNS naming system works and how to set up a DNS server.

Domains and domain names

To provide a unique DNS name for every host computer on the Internet, DNS uses a time-tested technique: divide and conquer. DNS uses a hierarchical naming system that's similar to the way folders are organized hierarchically on a Windows computer. Instead of folders, however, DNS organizes its names into *domains*. Each domain includes all the names that appear directly beneath it in the DNS hierarchy.

For example, Figure 5-5 shows a small portion of the DNS domain tree. At the top of the tree is the *root domain*, which is the anchor point for all domains. Directly beneath the root domain are four *top-level domains*, named `edu`, `com`, `org`, and `gov`.

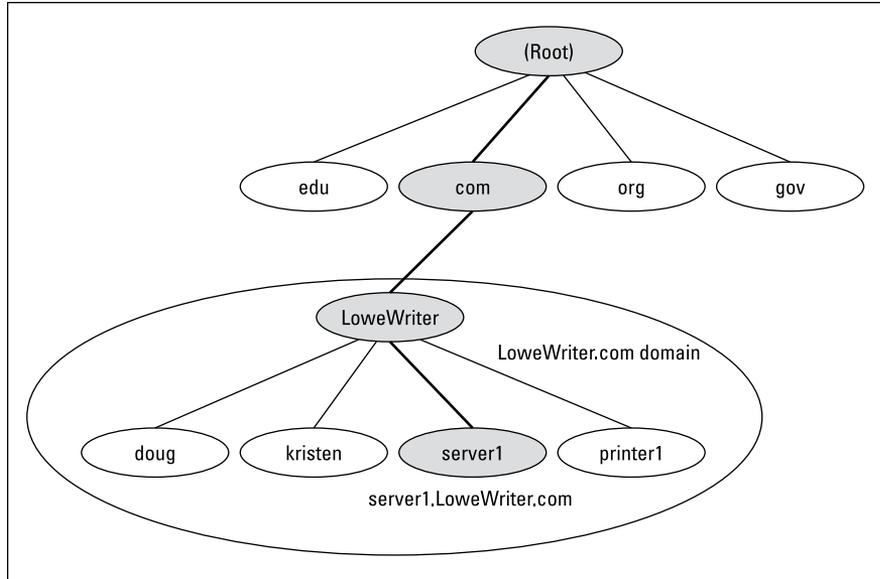


FIGURE 5-5:
DNS names.

In reality, many more top-level domains than this exist in the Internet's root domain. In fact, at the time I wrote this, there were about a thousand top-level domains in use.

Beneath the `com` domain in Figure 5-5 is another domain named `LoweWriter`, which happens to be my own personal domain. (Pretty clever, eh?) To completely identify this domain, you have to combine it with the name of its *parent domain* (in this case, `com`) to create the complete domain name: `LoweWriter.com`. Notice that the parts of the domain name are separated from each other by periods, which are pronounced “dot.” As a result, when you read this domain name, you should pronounce it “LoweWriter dot com.”

Beneath the `LoweWriter` node are four host nodes, named `doug`, `kristen`, `server1`, and `printer1`. These nodes correspond to three computers and a printer on my home network. You can combine the host name with the domain name to get the complete DNS name for each of my network's hosts. For example, the complete DNS name for my server is `server1.LoweWriter.com`. Likewise, my printer is `printer1.LoweWriter.com`.

Here are a few additional details that you need to remember about DNS names:



TIP

- » DNS names aren't case-sensitive. As a result, `Lowewriter` and `Lowewriter` are treated as the same name, as are `LOWEWRIter`, `LOWewriter`, and `Lowewriter`. When you use a domain name, you can use capitalization to make the name easier to read, but DNS ignores the difference between capital and lowercase letters.
- » The name of each DNS node can be up to 63 characters long (not including the dot) and can include letters, numbers, and hyphens. No other special characters are allowed.
- » A *subdomain* is a domain that's beneath an existing domain. For example, the `com` domain is a subdomain of the root domain. Likewise, `Lowewriter.com` is a subdomain of the `com` domain.
- » DNS is a hierarchical naming system that's similar to the hierarchical folder system used by Windows. However, one crucial difference exists between DNS and the Windows naming convention. When you construct a complete DNS name, you start at the bottom of the tree and work your way up to the root. Thus, `doug` is the lowest node in the name `doug.Lowewriter.com`. By contrast, Windows paths are the opposite: They start at the root and work their way down. For example, in the path `\Windows\System32\dns`, `dns` is the lowest node.
- » The DNS tree can be up to 127 levels deep. However, in practice, the DNS tree is pretty shallow. Most DNS names have just three levels (not counting the root), and although you sometimes see names with four or five levels, you rarely see more levels than that.
- » Although the DNS tree is shallow, it's very broad. In other words, each of the top-level domains has a huge number of second-level domains immediately beneath it. For example, at the time I wrote this book, the `com` domain had more than two million second-level domains beneath it.

Fully qualified domain names

If a domain name ends with a trailing dot, that trailing dot represents the root domain, and the domain name is said to be a *fully qualified domain name* (FQDN). A fully qualified domain name — also called an *absolute name* — is unambiguous because it identifies itself all the way back to the root domain. In contrast, if a domain name doesn't end with a trailing dot, the name may be interpreted in the context of some other domain. Thus, DNS names that don't end with a trailing dot are *relative names*.

This concept is similar to the way relative and absolute paths work in Windows. For example, if a path begins with a backslash, such as `\Windows\System32\dns`, the path is absolute. However, a path that doesn't begin with a backslash, such as `System32\dns`, uses the current folder as its starting point. If the current folder happens to be `\Windows`, `\Windows\System32\dns` and `System32\dns` refer to the same location.

In many cases, relative and fully qualified domain names are interchangeable because the software that interprets them always interprets relative names in the context of the root domain. That's why, for example, you can type `www.wiley.com` — without the trailing dot — rather than `www.wiley.com.` to go to the Wiley home page in a web browser. Some applications, such as DNS servers, may interpret relative names in the context of a domain other than the root.

Working with the Windows DNS Server

The procedure for installing and managing a DNS server depends on the network operating system you're using. This section is specific to working with a DNS server in Windows 2016. Working with a DNS server in a Linux or Unix environment is similar but without the help of a graphical user interface.

You can install the DNS server on Windows Server 2016 from the Server Manager (choose Server Manager on the taskbar). After you install the DNS server, you can manage it from the DNS management console. Here, you can perform common administrative tasks, such as adding additional zones, changing zone settings, or adding new records an existing zone. The DNS management console hides the details of the resource records from you, thus allowing you to work with a friendly graphical user interface instead.

To add a new host (which is defined by a DNS record called an A record) to a zone, right-click the zone in the DNS management console and choose the Add New Host command. This action opens the New Host dialog box, as shown in Figure 5-6.

Here, you specify the following information:

- » **Name:** The host name for the new host.
- » **IP Address:** The host's IP address.

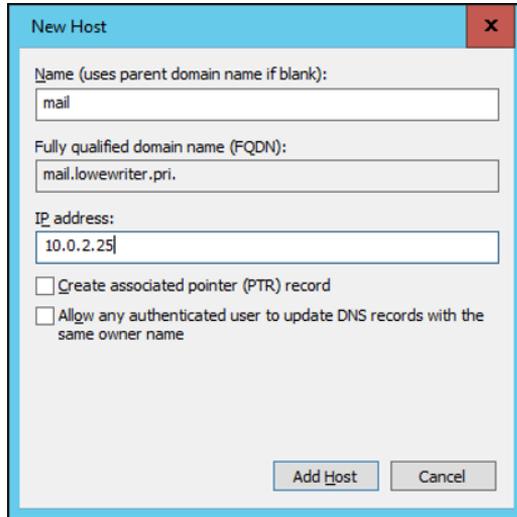


FIGURE 5-6:
The New Host
dialog box.

- » **Create Associated Pointer (PTR) Record:** Automatically creates a PTR record in the reverse lookup zone file. Select this option if you want to allow reverse lookups for the host. (A *reverse lookup* determines the domain name for a given IP address. It's called that because the normal type of DNS lookup determines the IP address for a given domain name.)
- » **Allow Any Authenticated User to Update:** Select this option if you want to allow other users to update this record or other records with the same host name. You should usually leave this option deselected.
- » **Time to Live:** The TTL value for this record, which indicates how long (in seconds) the data should be cached.

You can add other records, such as MX records, in the same way.

Configuring a Windows DNS Client

Client computers don't need much configuration to work properly with DNS. The client must have the address of at least one DNS server. Usually, this address is supplied by DHCP, so if the client is configured to obtain its IP address from a DHCP server, it also obtains the DNS server address from DHCP.

To configure a client computer to obtain the DNS server location from DHCP, open the Network Properties dialog box by choosing Network or Network Connections in the Control Panel (depending on which version of Windows the client is running). Then select the TCP/IP protocol and click the Properties button. This action summons the TCP/IP Properties dialog box, which is shown back in Figure 5-4. To configure the computer to use DHCP, select the Obtain an IP Address Automatically and the Obtain DNS Server Address Automatically options. Click OK, and you're done.

IN THIS CHAPTER

Getting a whiff of Ethernet

Checking out the different types of network cable

Installing twisted-pair cable

Working with hubs and switches

Examining routers

Chapter 6

Oh, What a Tangled Web We Weave: Cables, Switches, and Routers

Cable is the plumbing of your network. In fact, working with network cable is a lot like working with pipe: You have to use the right pipe (cable), the right valves and connectors (switches and routers), and the right fixtures (network interface cards).

And network cables have an advantage over pipes: You don't get wet when they leak.

This chapter tells you far more about network cables than you probably need to know. I introduce you to *Ethernet*, the most common system of network cabling for small networks. Then you find out how to work with the cables used to wire an Ethernet network. You also find out how to install a network interface card (NIC), which enables you to connect the cables to your computer.

What Is Ethernet?

Ethernet is a standardized way of connecting computers to create a network.

You can think of Ethernet as a kind of municipal building code for networks: It specifies what kind of cables to use, how to connect the cables, how long the cables can be, how computers transmit data to one another by using the cables, and more.



TECHNICAL
STUFF

Although Ethernet is now the overwhelming choice for networking, that wasn't always the case. In ye olde days, Ethernet had two significant competitors:

- » **Token Ring:** This IBM standard for networking is still in some organizations (especially where IBM mainframe or midrange systems are in use).
- » **ARCnet:** This standard is still commonly used for industrial network applications, such as building automation and factory robot control.

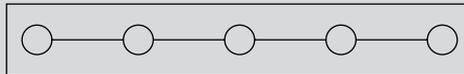
But these older networks are now pretty much obsolete, so you don't need to worry about them. Ethernet is now the only real choice for new networks — small or large.

OBLIGATORY INFORMATION ABOUT NETWORK TOPOLOGY



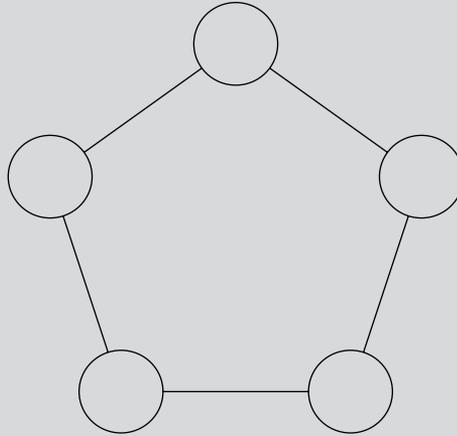
TECHNICAL
STUFF

A networking book wouldn't be complete without the usual textbook description of the three basic network topologies. One type of network topology is a *bus*, in which network nodes (that is, computers) are strung together in a line, like this:



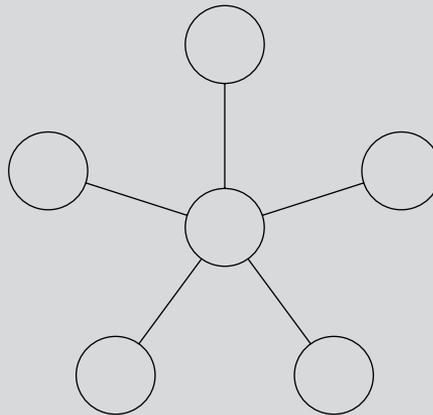
A *bus* is the simplest type of topology, but it has some drawbacks. If the cable breaks somewhere in the middle, the whole network breaks.

A second type of topology is the ring:



A ring is very much like a bus except with no end to the line: The last node on the line is connected to the first node, forming an endless loop.

A third type of topology is a star:



In a star network, all the nodes are connected to a central hub. In effect, each node has an independent connection to the network, so a break in one cable doesn't affect the others.

Ethernet networks are based on a bus design. However, fancy cabling tricks make an Ethernet network appear to be wired like a star when twisted-pair cable is used.

Here are a few tidbits you're likely to run into at parties where the conversation is about Ethernet standards:

- » Ethernet is a set of standards for the infrastructure on which a network is built. All the operating systems that I discuss in this book can operate on an Ethernet network. If you build your network on a solid Ethernet base, you can change network operating systems later.
- » Ethernet is often referred to by network gurus as 802.3 (pronounced "eight-oh-two-dot-three"), which is the official designation used by the *IEEE* (pronounced "eye-triple-ee," not "aiieee!"), a group of electrical engineers who wear bow ties and have nothing better to do than argue all day long about things like inductance and cross-talk — and it's a good thing they do. If not for them, you couldn't mix and match Ethernet components made by different companies.
- » The original vintage Ethernet transmits data at a rate of 10 million bits per second, or 10 Mbps. (*Mbps* is usually pronounced "megabits per second.") Because 8 bits are in a byte, that translates into roughly 1.2 million bytes per second. In practice, Ethernet can't move information that fast because data must be transmitted in packages of no more than 1,500 bytes, called *packets*. So 150KB of information has to be split into 100 packets.



Ethernet's transmission speed has nothing to do with how fast electrical signals move on the cable. The electrical signals travel at about 70 percent of the speed of light, or as Captain Kirk would say, "Warp factor point-seven-oh."

- » A faster version of Ethernet, called *100 Mbps Ethernet* or *Fast Ethernet*, moves data ten times as fast as normal Ethernet.
- » The most common version of Ethernet today is *gigabit Ethernet*, which moves data at 1,000 Mbps.
- » Most networking components that you can buy these days support 10, 100 Mbps and 1,000 Mbps. These components are called *10/100/1000 Mbps components*.
- » Some networking components support 10 gigabit Ethernet, which moves data at 10,000 Mbps (or 10 Gbps). Ten Gbps Ethernet is usually used for high-speed connections between servers and network switches.

All about Cable

Although you can use wireless technology to create networks without cables, most networks still use cables to physically connect each computer to the network. Over the years, various types of cables have been used with Ethernet networks. Almost

all networks are now built with twisted-pair cable. In this type of cable, pairs of wires are twisted around each other to reduce electrical interference. (You almost need a PhD in physics to understand why twisting the wires helps to reduce interference, so don't feel bad if this concept doesn't make sense.)

You may encounter other types of cable in an existing network; for example, on older networks, you may encounter two types of *coaxial* cable (also known as *coax*, pronounced "COE-ax"). The first type, which resembles television cable, is RG-58 cable. The second type is a thick, yellow cable that used to be the only type of cable used for Ethernet. You may also encounter fiber optic cables that span long distances at high speeds or thick twisted-pair bundles that carry multiple sets of twisted-pair cable between wiring closets in a large building. Most networks, however, use simple twisted-pair cable.

Twisted-pair cable is sometimes called *UTP*. (The *U* stands for *unshielded*, but "twisted-pair" is the standard name.) Figure 6-1 shows a twisted-pair cable.

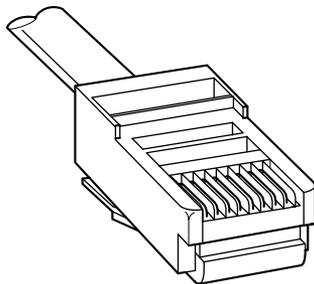


FIGURE 6-1:
Twisted-pair
cable.

When you use UTP cable to construct an Ethernet network, you connect the computers in a star arrangement, as Figure 6-2 illustrates. In the center of this star is a *switch*. Depending on the model, Ethernet switches enable you to connect 4 to 48 computers (or more) by using twisted-pair cable.



TIP

In the UTP star arrangement, if one cable goes bad, only the computer attached to that cable is affected. The rest of the network continues to chug along.

Cable categories

Twisted-pair cable comes in various grades: categories. These categories are specified by the ANSI/EIA Standard 568. (*ANSI* stands for American National Standards Institute; *EIA* stands for Electronic Industries Association). The standards indicate the data capacity — bandwidth — of the cable. Table 6-1 lists the various categories of twisted-pair cable.

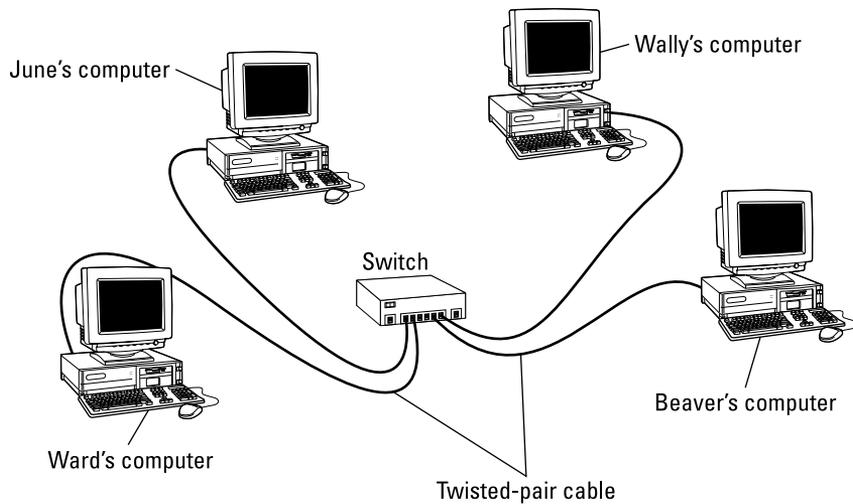


FIGURE 6-2:
A network cabled with twisted-pair cable.

TABLE 6-1 Twisted-Pair Cable Categories

Category	Maximum Data Rate	Intended Use
1	1 Mbps	Voice only
2	4 Mbps	4 Mbps Token Ring
3	16 Mbps	10BaseT Ethernet
4	20 Mbps	16 Mbps Token Ring
5	100 Mbps (2-pair)	100BaseT Ethernet
	1000 Mbps (4-pair)	1000BaseTX
5e	1000 Mbps (2-pair)	1000BaseT
6	1000 Mbps (2-pair)	1000BaseT and faster broadband applications
6a	10000 Mbps (2-pair)	Provides for 10 Gbps Ethernet

Although higher-category cables are more expensive, the real cost of installing Ethernet cabling is the labor required to pull the cables through the walls. You should never install anything less than Category (Cat) 5e cable. And if at all possible, invest in Cat 6 cable to allow for upgrades to your network.



TIP

To sound like the cool kids, say “Cat 6” rather than “Category 6.”

What's with the pairs?

Most twisted-pair cables have four pairs of wires, for a total of eight wires. Standard Ethernet uses only two of the pairs, so the other two pairs are unused. You may be tempted to save money by purchasing cable with just two pairs of wires, but that's a bad idea. If a network cable develops a problem, you can sometimes fix it by switching over to one of the extra pairs. If you use two-pair cable, though, you don't have any spare pairs to use.



WARNING

Don't use the extra pairs for some other purpose, such as a voice line or a second data line. The electrical "noise" in the extra wires can interfere with your network.

To shield or not to shield

Unshielded twisted-pair cable (UTP) is designed for normal office environments. When you use UTP cable, you must be careful not to route cable close to fluorescent light fixtures, air conditioners, or electric motors (such as automatic door motors or elevator motors). UTP is the least expensive type of cable.

In environments that have a lot of electrical interference (such as factories), you may want to use shielded twisted-pair cable (STP). Because STP can be as much as three times more expensive than regular UTP, you don't want to use STP unless you have to. With a little care, UTP can withstand the amount of electrical interference found in a normal office environment. But for harsh environments, where cable will be placed outdoors, buried in the ground, or near industrial equipment, you should use STP cable instead.

Most STP cable is shielded by a layer of aluminum foil. For buildings with unusually high amounts of electrical interference, the more expensive braided-copper shielding offers even more protection.

When to use plenum cable

The outer sheath of shielded and unshielded twisted-pair cable comes in two kinds:

- » **PVC:** The most common and least expensive type.
- » **Plenum:** A special type of fire-retardant cable designed for use in the plenum space (definition coming right up) of a building. Plenum cable has a special Teflon coating that not only resists heat, but also gives off fewer toxic fumes if it does burn. Unfortunately, plenum cable costs more than twice as much as ordinary PVC cable.



WARNING

Most local building codes require plenum cable when the wiring is installed in the building's *plenum space* (a compartment that's part of the building's air-distribution system, usually the space above a suspended ceiling or under a raised floor).



TIP

The area above a suspended ceiling is *not* a plenum space if both the delivery and return lines of the air-conditioning and heating systems are ducted. Plenum cable is required only if the air-conditioning and heating systems aren't ducted. When in doubt, have the local inspector look at your facility before you install cable.

Sometimes solid, sometimes stranded

The actual copper wire that makes up the cable comes in two varieties: solid and stranded. Your network will have some of each:

» **Stranded:** Each conductor is made from a bunch of very small wires that are twisted together. Stranded cable is more flexible than solid cable, so it doesn't break as easily. However, stranded cable is more expensive than solid cable and isn't very good at transmitting signals over long distances. Stranded cable is best used for *patch cables* (such as patch panels to hubs and switches).

Strictly speaking, the cable that connects your computer to the wall jack is a station cable — not a patch cable — but it's an appropriate use for stranded cable. (Although not technically correct, most people refer to the cable that connects a computer to a wall jack as a "patch cable.")

» **Solid:** Each conductor is a single, solid strand of wire. Solid cable is less expensive than stranded cable and carries signals farther, but it isn't very flexible. If you bend it too many times, it breaks. Typically, you find solid cable in use as permanent wiring within the walls and ceilings of a building.



TIP

Installation guidelines

The hardest part of installing network cable is the physical task of pulling the cable through ceilings, walls, and floors. This job is just tricky enough that I recommend you don't attempt it yourself, except for small offices. For large jobs, hire a professional cable installer. You may even want to hire a professional for small jobs if the ceiling and wall spaces are difficult to access.

Keep these pointers in mind if you install cable yourself:

» You can purchase twisted-pair cable in prefabricated lengths, such as 10, 15, or 20 feet. Longer lengths, such as 50 feet or 100 feet, are also available.



WARNING

- » Alternatively, you can purchase cable in bulk rolls, cut them to length, and attach the connectors yourself.
- » Always use a bit more cable than you need, especially if you're running cable through walls. For example, when you run a cable up a wall, leave a few feet of slack in the ceiling above the wall. That way, you have plenty of cable if you need to make a repair.
- » When running cable, avoid sources of interference, such as fluorescent lights, big motors, X-ray machines, nuclear reactors, cyclotrons, or other gadgets you may have hidden in behind closed doors in your office.

Fluorescent lights are the most common sources of interference for cables behind ceiling panels. Give light fixtures a wide berth. Three feet should do it.

- » The maximum allowable cable length between the hub and the computer is 100 meters (about 328 feet).
- » If you must run cable across the floor where people walk, cover the cable so no one trips over it. Cable protectors are available at most hardware stores.
- » When running cables through walls, label each cable at both ends. Most electrical supply stores carry pads of cable labels that are perfect for the job. These pads contain 50 sheets or so of precut labels with letters and numbers. They look much more professional than wrapping a loop of masking tape around the cable and writing on the tape with a marker.

Alternatively, you can just write directly on the label with a permanent marker.



TIP

- » If you're installing cable in new construction, label each end of the cable at least three times, leaving about a foot of space between the labels. The drywallers or painters will probably spray mud or paint all over your cables, making the labels difficult to find.
- » When several cables come together, tie them with plastic cable ties. Avoid masking tape if you can; the tape doesn't last, but the sticky glue stuff does. It's a mess a year later. Cable ties are available at electrical supply stores.



TIP

- » Cable ties have all sorts of useful purposes. Once, on a backpacking trip, I used a pair of cable ties to attach an unsuspecting buddy's hat to a high tree limb. He wasn't impressed with my innovative use of the cable ties, but my other hiking companions were.
- » When you run cable above suspended ceiling panels, use cable ties, hooks, or clamps to secure the cable to the ceiling or to the metal frame that supports the ceiling tiles. Don't just lay the cable on top of the panels.

The tools you need

Of course, to do a job right, you must have the right tools:



TIP

- » Start with a basic set of computer tools, which you can get for about \$15 from any computer store and most office-supply stores. These kits include socket wrenches and screwdrivers to open your computers and insert adapter cards.

The computer tool kit probably contains everything you need if

- All your computers are in the same room.
- You're running the cables along the floor.
- You're using prefabricated cables.



TIP

If you don't have a computer tool kit, make sure that you have several flat-head and Phillips screwdrivers of various sizes.

- » If you're using bulk cable and plan on attaching your own connectors, you also need the following tools in addition to the basic computer tool kit:

- *Wire cutters*: You need big ones for coax; smaller ones work for twisted-pair cable. For yellow cable, you need the Jaws of Life.
- *A crimp tool*: You need the crimp tool to attach the connectors to the cable. Don't use a cheap \$25 crimp tool. A good crimp tool costs \$100 and will save you many headaches in the long run.

When you crimp, you mustn't scrimp.

- *Wire stripper*: You need this tool only if the crimp tool doesn't include a wire stripper.
- *A cable tester*, which lets you determine whether the cable will work.

- » If you plan on running cables through walls, you need these additional tools:

- *A hammer*
- *A keyhole saw*: This one is useful if you plan on cutting holes through walls to route your cable.
- *A flashlight*
- *A ladder*
- *Someone to hold the ladder*
- *Fish tape*: Possibly. A *fish tape* is a coiled-up length of stiff metal tape. To use it, you feed the tape into one wall opening and fish it toward the other opening, where a partner is ready to grab it when the tape arrives. Next, your partner attaches the cable to the fish tape and yells something like



REMEMBER

“Let ‘er rip!” or “Bombs away!” Then you reel in the fish tape and the cable along with it. (You can find fish tape in the electrical section of most well-stocked hardware stores.)



If you plan on routing cable through a concrete subfloor, you need to rent a jackhammer and a backhoe and then hire someone to hold a yellow flag while you work. Better yet, find some other route for the cable.

Pinouts for twisted-pair cables

Each pair of wires in a twisted-pair cable is one of four colors: orange, green, blue, or brown. The two wires that make up each pair are complementary: one is white with a colored stripe; the other is colored with a white stripe. For example, the orange pair has an orange wire with a white stripe (the *orange wire*) and a white wire with an orange stripe (the *white/orange wire*). Likewise, the blue pair has a blue wire with a white stripe (the *blue wire*) and a white wire with a blue stripe (the *white/blue wire*).

When you attach a twisted-pair cable to a modular connector or jack, you must match up the right wires to the right pins. It’s harder than it sounds; you can use any of several different standards to wire the connectors. To confuse matters further, you can use one of the two popular standard ways of hooking up the wires: EIA/TIA 568A or EIA/TIA 568B, also known as AT&T 258A. Both of these wiring schemes are shown in Table 6-2.

TABLE 6-2 Pin Connections for Twisted-Pair Cable

Pin	Function	EIA/TIA 568A	EIA/TIA568B AT&T 258A
1	Transmit +	White/Green	White/orange wire
2	Transmit -	Green	Orange wire
3	Receive +	White/Orange	White/green wire
4	Unused	Blue	Blue wire
5	Unused	White/Blue	White/blue wire
6	Receive -	Orange	Green wire
7	Unused	White/Brown	White/brown wire
8	Unused	Brown	Brown wire



WARNING

It doesn't matter which of these wiring schemes you use, but pick one and stick with it. If you use one wiring standard on one end of a cable and the other standard on the other end, the cable doesn't work.

Ethernet only uses two of the four pairs, connected to Pins 1, 2, 3, and 6. One pair transmits data; the other receives data. The only difference between the two wiring standards is which pair transmits and which receives. In the EIA/TIA 568A standard, the green pair is used for transmit, and the orange pair is used for receive. In the EIA/TIA 568B and AT&T 258A standards, the orange pair is used for transmit and the green pair for receive.

If you want, you can get away with connecting only Pins 1, 2, 3, and 6. However, I suggest that you connect all four pairs, as indicated in Table 6-2.

RJ-45 connectors

RJ-45 connectors for twisted-pair cables aren't too difficult to attach if you have the right crimping tool. The only trick is making sure that you attach each wire to the correct pin and then press the tool hard enough to ensure a good connection.

Here's the procedure for attaching an RJ-45 connector:

- 1. Cut the end of the cable to the desired length.**
Make sure that you make a square cut — not a diagonal cut.
- 2. Insert the cable into the stripper portion of the crimp tool so that the end of the cable is against the stop.**
Squeeze the handles and slowly pull out the cable, keeping it square. This strips off the correct length of outer insulation without puncturing the insulation on the inner wires.
- 3. Arrange the wires so that they lie flat and line up according to Table 6-2.**
You have to play with the wires a little bit to get them to lay out in the right sequence.
- 4. Slide the wires into the pinholes on the connector.**
Double-check to make sure all the wires are slipped into the correct pinholes.
- 5. Insert the plug and wire into the crimping portion of the tool and then squeeze the handles to crimp the plug.**
Squeeze it tight!
- 6. Remove the plug from the tool and double-check the connection.**
You're done!

Here are a few other points to remember when dealing with RJ-45 connectors and twisted-pair cable:



TIP

- » The pins on the RJ-45 connectors aren't numbered.

You can tell which is Pin 1 by holding the connector so that the metal conductors are facing up, as shown in Figure 6-3. Pin 1 is on the left.

- » Some people wire the cable differently — using the green-and-white pair for Pins 1 and 2, and the orange-and-white pair for Pins 3 and 6. Doing it this way doesn't affect the operation of the network (the network is color blind) *as long as the connectors on both ends of the cable are wired the same way!*
- » When you attach the connectors, don't untwist more than half an inch of cable. And, don't try to stretch the cable runs beyond the 100-meter maximum. When in doubt, have the cable professionally installed.

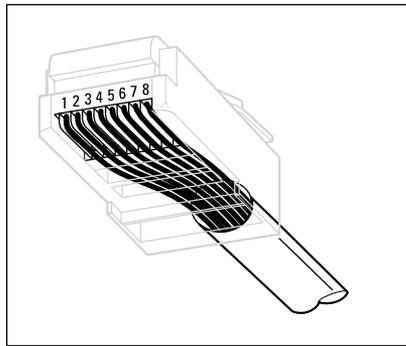


FIGURE 6-3:
Attaching an RJ-45 connector to twisted-pair cable.

Crossover cables

A *crossover cable* can directly connect two devices without a switch. You can use a crossover cable to connect two computers directly to each other, but crossover cables are more often used to daisy-chain hubs and switches to each other.

If you want to create your own crossover cable, you must reverse the wires on one end of the cable, as shown in Table 6-3. This table shows how you should wire both ends of the cable to create a crossover cable. Connect one of the ends according to the Connector A column and the other according to the Connector B column.

Note that you don't need to use a crossover cable if one of the switches or hubs that you want to connect has a crossover port, usually labeled Uplink or Daisy-chain. If the hub or switch has an Uplink port, you can daisy-chain it by using a

normal network cable. For more information about daisy-chaining hubs and switches, see the section “Daisy-Chaining Switches,” later in this chapter.



TIP

If you study Table 6-3 long enough and then compare it with Table 6-2, you may notice that a crossover cable is a cable that’s wired according to the 568A standard on one end and the 568B standard on the other end.

TABLE 6-3 **Creating a Crossover Cable**

Pin	Connector A	Connector B
1	White/green	White/orange
2	Green	Orange
3	White/orange	White/green
4	Blue	Blue
5	White/blue	White/blue
6	Orange	Green
7	White/brown	White/brown
8	Brown	Brown

Wall jacks and patch panels

If you want, you can run a single length of cable from a network hub or switch in a wiring closet through a hole in the wall, up the wall to the space above the ceiling, through the ceiling space to the wall in an office, down the wall, through a hole, and all the way to a desktop computer. That’s not a good idea. For example, every time someone moves the computer or even cleans behind it, the cable will get moved a little bit. Eventually, the connection will fail, and the RJ-45 plug will have to be replaced. Then the cables in the wiring closet will quickly become a tangled mess.

The alternative is to put a wall jack in the wall at the user’s end of the cable and connect the other end of the cable to a patch panel. Then the cable itself is completely contained within the walls and ceiling spaces. To connect a computer to the network, you plug one end of a patch cable (properly called a *station cable*) into the wall jack and plug the other end into the computer’s network interface. In the wiring closet, you use a patch cable to connect the wall jack to the network hubs or switches. Figure 6-4 shows how this arrangement works.

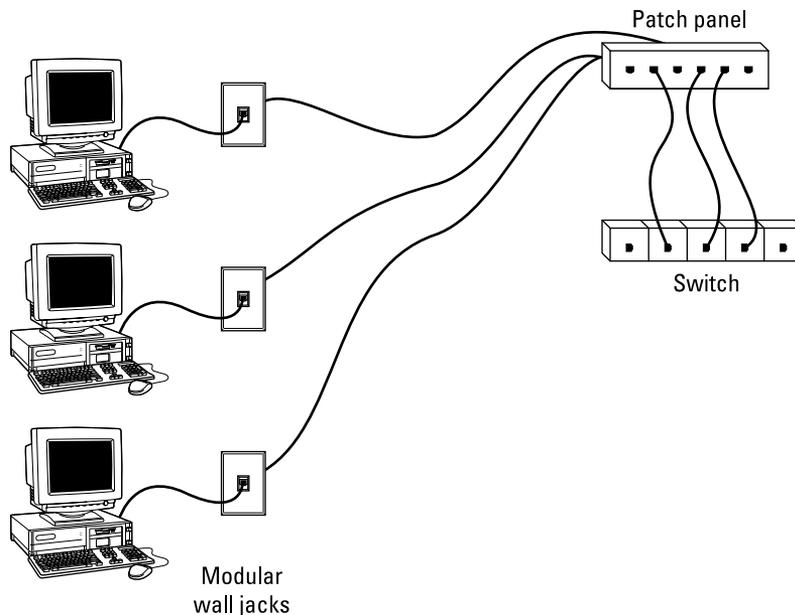


FIGURE 6-4: Using wall jacks and patch panels.

Connecting a twisted-pair cable to a wall jack or a patch panel is similar to connecting it to an RJ-45 plug. However, you don't usually need any special tools. Instead, the back of the jack has a set of slots that you lay each wire across. You then snap a removable cap over the top of the slots and press it down. This action forces the wires into the slots, where little metal blades pierce the insulation and establish the electrical contact.



TIP

When you connect the wire to a jack or a patch panel, be sure to untwist as little of the wire as possible. If you untwist too much of the wire, the signals that pass through the wire may become unreliable.

Working with Switches

When you use twisted-pair cable to wire a network, you don't plug the computers into each other. Instead, each computer plugs into a separate device called a *switch*.

You need to know only a few details when working with switches. Here they are:

- » Installing a switch is usually very simple. Just plug in the power cord and then plug in patch cables to connect the network.



TIP

- » Each port on the switch has an RJ-45 jack and a single LED indicator, labeled *Link*, that lights up when a connection is made on the port.

If you plug one end of a cable into the port and the other end into a computer or other network device, the Link light should come on. If it doesn't, something is wrong with the cable, the hub or switch port, or the device on the other end of the cable.



TIP

- » Each port may have an LED indicator that flashes to indicate network activity.

If you stare at a switch for a while, you can find out who uses the network most by noting which activity indicators flash the most.

- » The ports may also have a collision indicator that flashes whenever a packet collision occurs on the port.

It's perfectly acceptable for the collision indicator to flash now and then, but if it flashes a lot, you may have a problem with the network:

- Usually, the flashing means that the network is overloaded and should be segmented with a switch to improve performance.
- In some cases, the flashing may be caused by a faulty network node that clogs the network with bad packets.



WARNING

Daisy-Chaining Switches

If a single switch doesn't have enough ports for your entire network, you can connect switches by *daisy-chaining* them, as shown in Figure 6-5.



TIP

You can often increase the overall performance of your network by using two (or more) connections between switches. For example, you may use two patch cables to create two connections between a pair of switches.



WARNING

Don't chain more than two switches together. If you do, the network may not transmit your data reliably. However, you can get around this rule by using *stackable switches* (switches with a special cable connector that connects two or more switches so that they function as a single switch). Stackable switches are musts for large networks.



TIP

If your building is prewired and has a network jack near each desk, you can use a small switch to connect two or more computers to the network by using a single jack. Just use one cable to plug the daisy-chain port of the hub into the wall jack and then plug each computer into one of the hub's ports.

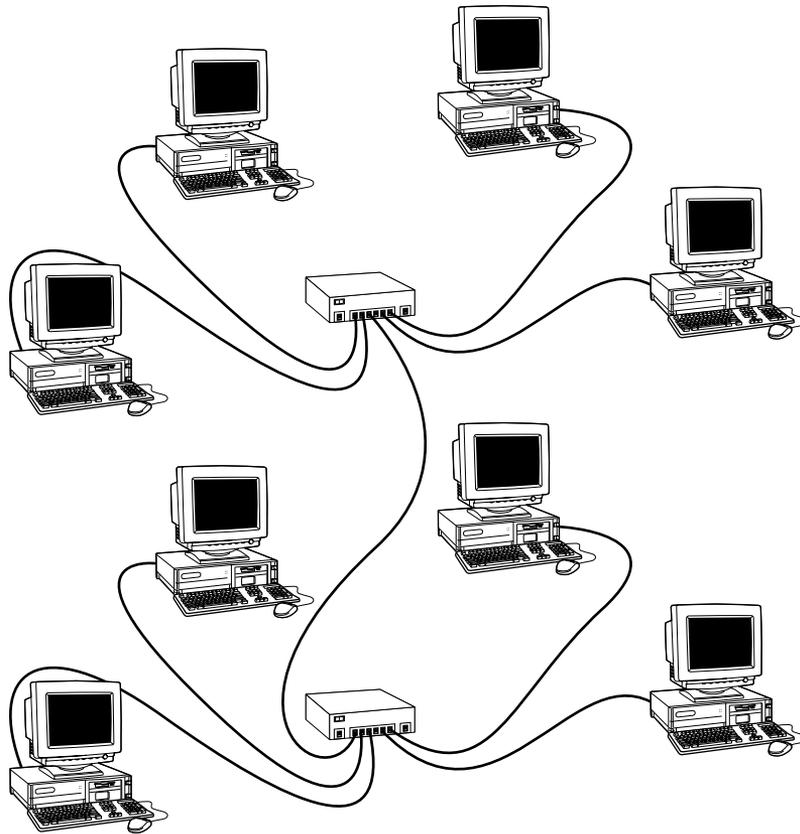


FIGURE 6-5:
You can
daisy-chain
switches
together.

Using a Router

A *router* is a device that is capable of passing data between two networks. The most common reason for using a router is to connect a LAN to the Internet. However, routers can perform many other functions as well. For example, a router can filter data based on its content, allowing some types of data to pass through while blocking other types.



TIP

You can configure a network with several routers that can work cooperatively. For example, some routers can monitor the network to determine the most efficient path for sending a message to its ultimate destination. If a part of the network is extremely busy, a router can automatically route messages along a less-busy route. In this respect, the router is kind of like a traffic reporter flying in a helicopter. The router knows that the 101 is bumper to bumper all the way through Sunnyvale, so it sends the message on the 280 instead.

Here's some additional information about routers:

- » Routers used to be expensive and used only on large networks. However, the price of small routers has dropped substantially in recent years, so they're now becoming common even on small networks.
- » The functional distinctions between bridges and routers — and switches and hubs, for that matter — get blurrier all the time. *Multifunction routers* (which combine the functions of routers, bridges, hubs, and switches) are often used to handle some chores that used to require separate devices.
- » A pair of routers can be used to create a secure connection between two locations that are geographically distant from each other.
- » One of the main reasons for using routers is to connect a LAN to the Internet. Figure 6-6 shows a router used for this purpose.

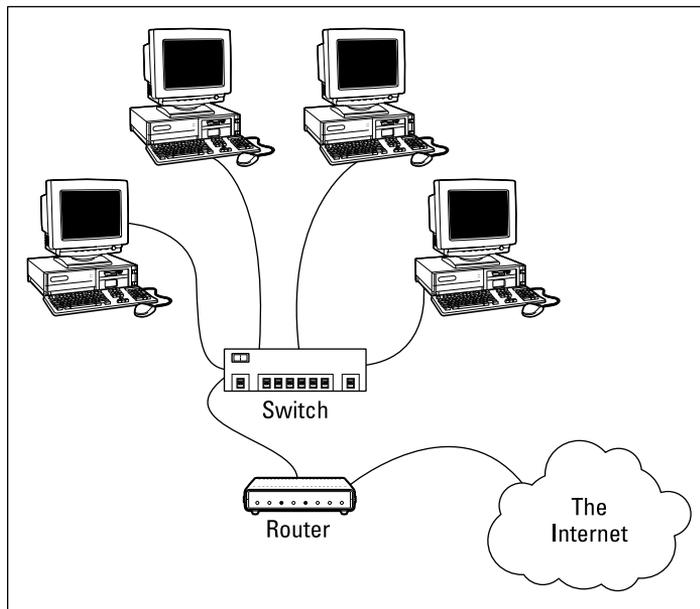


FIGURE 6-6: Using a router to connect a LAN to the Internet.

Chapter 7

Configuring Windows Clients

Before your network setup is complete, you must configure the network's client computers. In particular, you have to configure each client computer's network interface card (NIC) so that it works properly, and you have to install the right protocols so that the clients can communicate with other computers on the network.

Fortunately, the task of configuring client computers for the network is child's play in Windows. For starters, Windows automatically recognizes your network interface card when you start up your computer. All that remains is to make sure that Windows properly installed the network protocols and client software.

With each version of Windows, Microsoft has simplified the process of configuring client network support. In this chapter, I describe the steps for configuring networking for Windows 10. The procedures for previous versions of Windows are similar.

Configuring Network Connections

Windows usually detects the presence of a network adapter automatically; typically, you don't have to install device drivers manually for the adapter. When

Windows detects a network adapter, Windows automatically creates a network connection and configures it to support basic networking protocols. You may need to change the configuration of a network connection manually, however.

The following steps show you how to configure your network adapter on a Windows 10 system:

1. **Click the Start icon (or press the Start button on the keyboard), and then tap or click Settings.**

The Settings page appears, as shown in Figure 7-1.

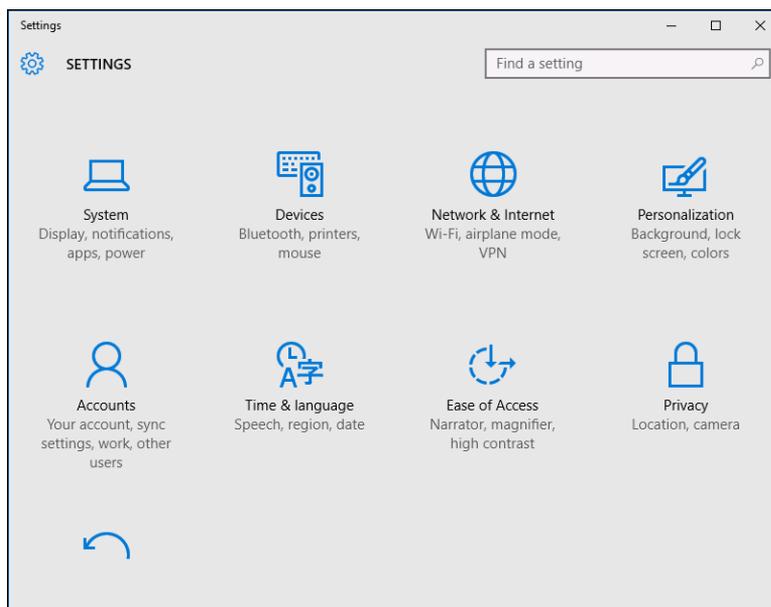


FIGURE 7-1:
The Settings page.

2. **Click Network & Internet.**

The Network & Internet page appears, as shown in Figure 7-2.

3. **Click Ethernet.**

The Ethernet settings page appears, as shown in Figure 7-3.

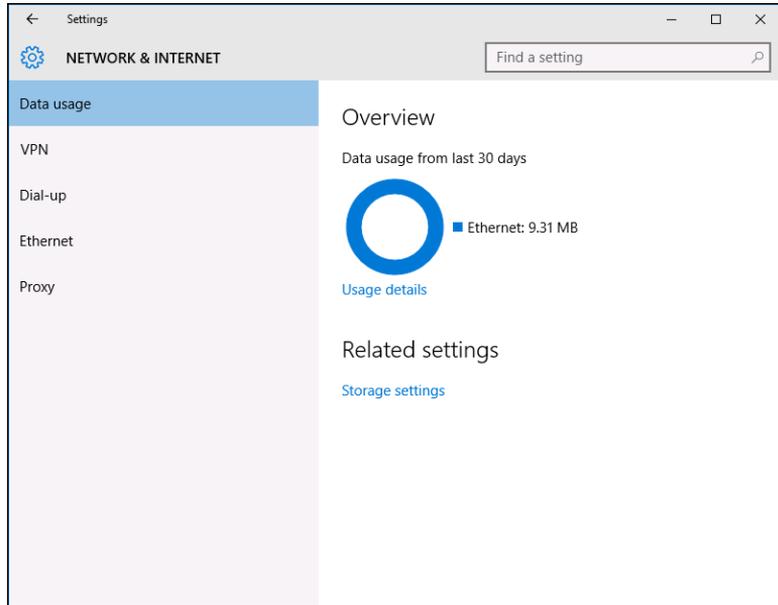


FIGURE 7-2:
The Network & Internet page.

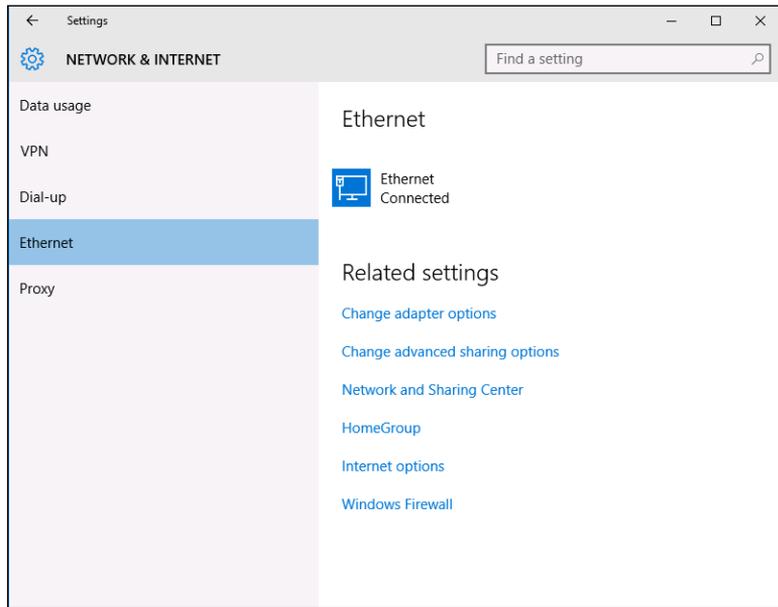


FIGURE 7-3:
The Ethernet settings page.

4. Click Change Adapter Options.

The Network Connections page appears, as shown in Figure 7-4. This page lists each of your network adapters. In this case, only a single wired Ethernet adapter is shown. If the device has more than one adapter, additional adapters will appear on this page.

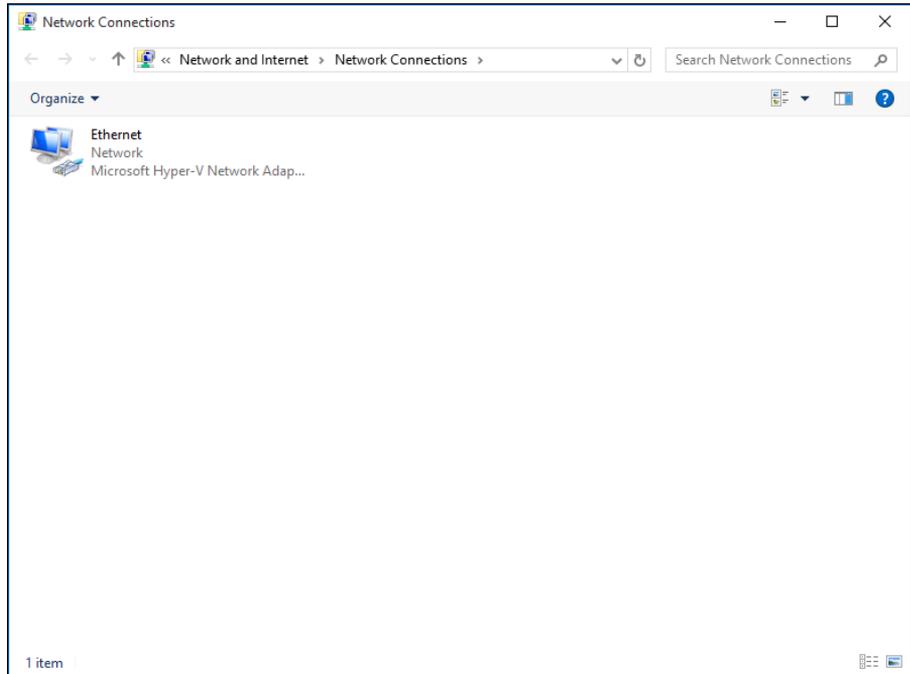


FIGURE 7-4:
The Network
Connections
page.

5. Right-click the connection that you want to configure and then choose Properties from the contextual menu that appears.

This action opens the Ethernet Properties dialog box, as shown in Figure 7-5.

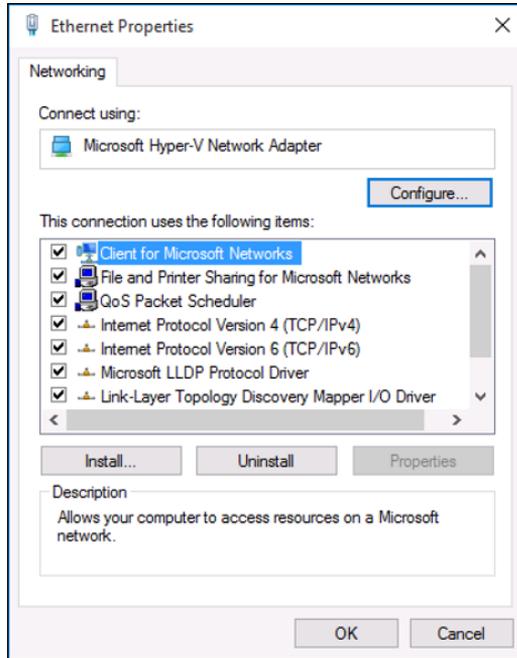


FIGURE 7-5:
The Ethernet
Properties
dialog box.

6. To configure the network adapter card settings, click **Configure**.

The Properties dialog box for your network adapter appears, as shown in Figure 7-6. This dialog box has seven tabs that let you configure the adapter:

- *General*: Shows basic information about the adapter, such as the device type and status.
- *Advanced*: Lets you set a variety of device-specific parameters that affect the operation of the adapter.
- *About*: Displays information about the device's patent protection.
- *Driver*: Displays information about the device driver that's bound to the NIC and lets you update the driver to a newer version, roll back the driver to a previously working version, or uninstall the driver.
- *Details*: With this tab, you can inspect various properties of the adapter such as the date and version of the device driver. To view the setting of a particular property, select the property name from the drop-down list.
- *Events*: Lists recent events that have been logged for the device.
- *Power Management*: Lets you configure power management options for the device.

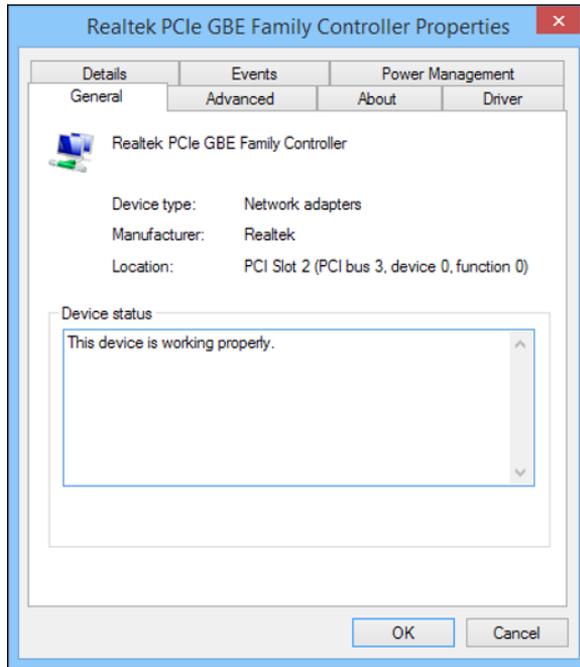


FIGURE 7-6:
The Properties dialog box for a network adapter.



TIP

When you click OK to dismiss the dialog box, the network connection's Properties dialog box closes and you're returned to the Network Connections page (refer to Figure 7-4). Right-click the network adapter and choose Properties again to continue the procedure.

7. Review the list of connection items listed in the Properties dialog box.

The most important items you commonly see are:

- *Client for Microsoft Networks*: This item is required if you want to access a Microsoft Windows network. It should always be present.
- *File and Printer Sharing for Microsoft Networks*: This item allows your computer to share its files or printers with other computers on the network.



TIP

This option is usually used with peer-to-peer networks, but you can use it even if your network has dedicated servers. If you don't plan to share files or printers on the client computer, however, you should disable this item.

- *Internet Protocol Version 4 (TCP/IPv4)*: This item enables the client computer to communicate by using the version 4 standard TCP/IP protocol.
- *Internet Protocol Version 6 (TCP/IPv6)*: This item enables version 6 of the standard TCP/IP protocol. Typically, both IP4 and IP6 are enabled, even though most networks rely primarily on IP4.

8. If a protocol that you need isn't listed, click the Install button to add the needed protocol.

A dialog box appears, asking whether you want to add a network client, protocol, or service. Click Protocol and then click Add. A list of available protocols appears. Select the one you want to add; then click OK.

9. To remove a network item that you don't need (such as File and Printer Sharing for Microsoft Networks), select the item, and click the Uninstall button.

For security reasons, you should make it a point to remove any clients, protocols, or services that you don't need.

10. To configure TCP/IP settings, click Internet Protocol (TCP/IP); click Properties to display the TCP/IP Properties dialog box (shown in Figure 7-7); adjust the settings; and then click OK.

The TCP/IP Properties dialog box lets you choose among these options:

- *Obtain an IP Address Automatically*: Choose this option if your network has a DHCP server that assigns IP addresses automatically. Choosing this option dramatically simplifies administering TCP/IP on your network. (See Chapter 5 for more information about DHCP.)
- *Use the Following IP Address*: If your computer must have a specific IP address, choose this option and then type the computer's IP address, subnet mask, and default gateway address. (For more information about these settings, see Chapter 5.)
- *Obtain DNS Server Address Automatically*: The DHCP server can also provide the address of the Domain Name System (DNS) server that the computer should use. Choose this option if your network has a DHCP server. (See Chapter 5 for more information about DNS.)
- *Use the Following DNS Server Addresses*: Choose this option if a DNS server isn't available. Then type the IP addresses of the primary and secondary DNS servers.

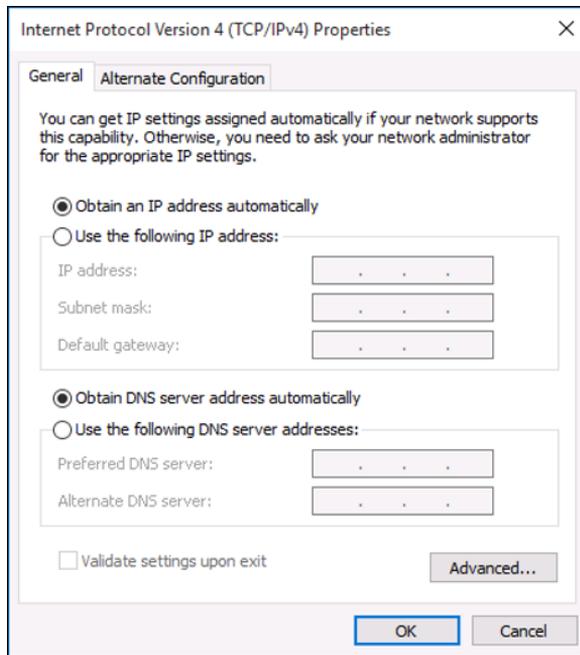


FIGURE 7-7:
Configuring
TCP/IP.

Joining a Domain

When Windows first installs, it isn't joined to a domain network. Instead, it's available as part of a workgroup, which is an unmanaged network suitable only for the smallest of networks with just a few computers and without dedicated servers. To use a computer in a domain network, you must join the computer to the domain. Here are the steps for Windows 10:

- 1. Click the Start icon (or press the Start button on the keyboard), and then tap or click Settings.**

The Settings page appears (refer to Figure 7-1).

- 2. Click System.**

The System settings page appears.

- 3. Click About.**

The PC settings page appears, as shown in Figure 7-8.

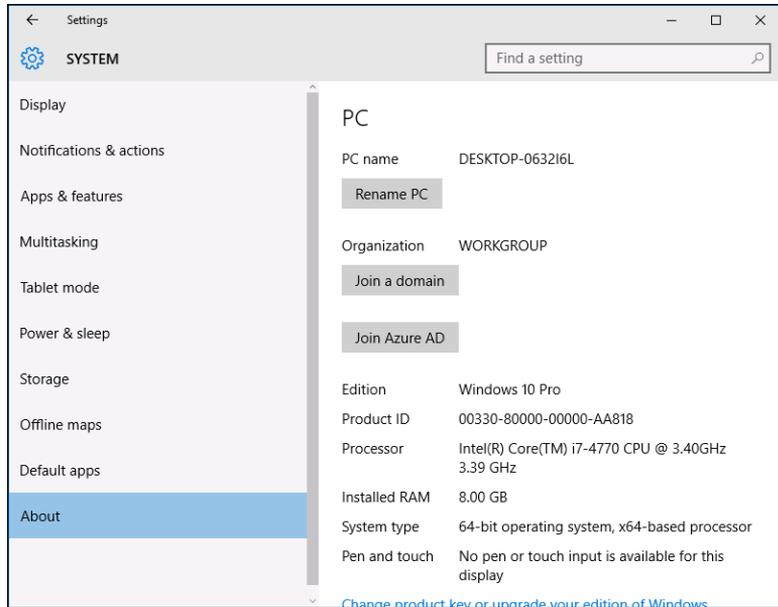


FIGURE 7-8:
The PC settings page.

4. (Optional) To change the name of the computer, click Rename PC.

You're prompted to enter a new name and then reboot the computer.

Before you join a domain, you should ensure that the computer's name won't be the same as the name of a computer that's already a member of the domain. If it is, you should first change the name.



TIP

If you do change the computer's name, repeat the procedure from Step 1 after the reboot.

5. Click Join a Domain.

The Join a Domain dialog box appears, as shown in Figure 7-9.

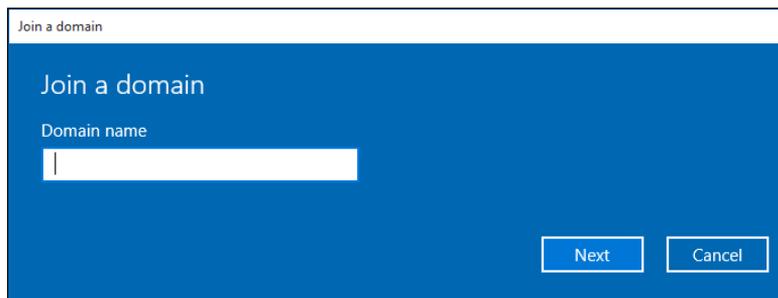


FIGURE 7-9:
Joining a domain.

6. Enter the domain name and click Next.

You're prompted for the user name and password of a user who has administration privileges on the domain, as shown in Figure 7-10.

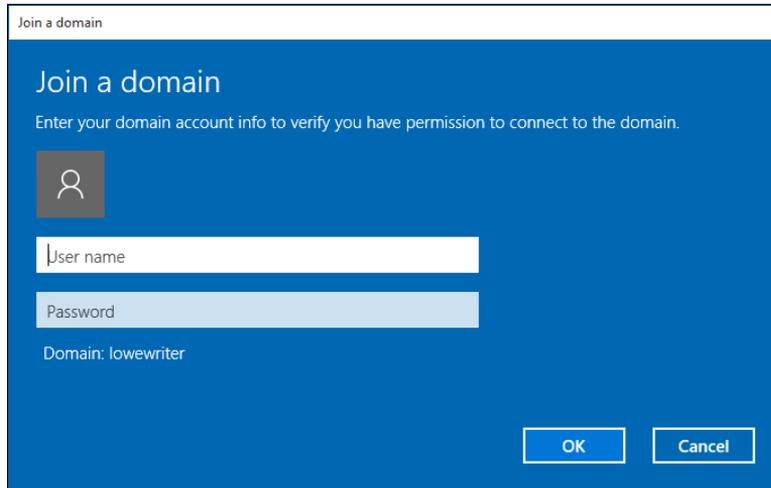


FIGURE 7-10: You must provide domain administrator credentials to join a domain.

7. Click OK.

8. Enter the username and password for an Administrator account when prompted.

You're asked to provide this information only if a computer account hasn't already been created for the client computer.

9. When informed that you need to restart the computer, click Restart Now.

The computer is restarted and added to the domain.

IN THIS CHAPTER

Looking at DSL and cable

Examining T1 and T3 connections

Using a router

Securing your connection with a firewall

Using the firewall that comes with Windows

Chapter 8

Connecting Your Network to the Internet

So you decided to connect your network to the Internet. All you have to do is call the cable company and have them send someone out, right? Wrong. Unfortunately, connecting to the Internet involves more than just calling the cable company. For starters, you have to make sure that cable is the right way to connect. Then you have to select and configure the software you use to access the Internet. Finally, you have to lie awake at night worrying whether hackers are breaking into your network via its Internet connection.

Not to worry. The advice in this chapter helps you decide how to connect to the Internet and, once the decision is made, how to do it safely.

Connecting to the Internet

Connecting to the Internet isn't free. For starters, you have to purchase the computer equipment necessary to make the connection. Then you have to obtain a connection from an Internet service provider (ISP). The ISP charges you a monthly fee that depends on the speed and capacity of the connection.

The following sections describe the most commonly used methods of connecting network users to the Internet.

Connecting with cable or DSL

For small and home offices, the two most popular methods of connecting to the Internet are cable and digital subscriber line (DSL). Cable and DSL connections are often called *broadband connections* for technical reasons you don't really want to know.

Cable Internet access works over the same cable that brings 40 billion TV channels into your home, whereas DSL is a digital phone service that works over a standard phone line. Both offer three major advantages over old-fashioned dialup connections:

» **Cable and DSL are much faster than dialup connections.**

A cable connection can be anywhere from 10 to 200 times faster than a dialup connection, depending on the service you get. And the speed of a DSL line is comparable with cable. (Although DSL is a dedicated connection, cable connections are shared among several subscribers. The speed of a cable connection may slow down when several subscribers use the connection simultaneously.)

» **With cable and DSL, you're always connected to the Internet.**

You don't have to connect and disconnect each time you want to go online like you would if you use a modem. No more waiting for the modem to dial your service provider and listening to the annoying modem shriek while it attempts to establish a connection.

» **Cable and DSL don't tie up a phone line while you're online.**

With cable, your Internet connection works over TV cables, not over phone cables. With DSL, the phone company installs a separate phone line for the DSL service, so your regular phone line isn't affected.

Unfortunately, there's no such thing as a free lunch, and the high-speed, always-on connections offered by cable and DSL don't come without a price. For starters, you can expect to pay a higher monthly access fee for cable or DSL. In most areas of the United States, cable runs about \$50 per month for residential users; business users can expect to pay more, especially if more than one user will be connected to the Internet via the cable.

The cost for DSL service depends on the access speed you choose. In some areas, residential users can get a relatively slow DSL connection for as little as \$30 per month. For higher access speeds or for business users, DSL can cost substantially more.

Too, cable and DSL access aren't available everywhere. But if you live in an area where cable or DSL isn't available, you can still get high-speed Internet access by using a satellite hookup or a cellular network.

Connecting with high-speed private lines

If your network is large and high-speed Internet access is a high priority, contact your local phone company (or companies) about installing a dedicated high-speed digital line. These lines can cost you plenty (on the order of hundreds of dollars per month), so they're best suited for large networks in which 20 or more users are accessing the Internet simultaneously.

The following paragraphs describe three basic options for high-speed private lines:

» **T1 and T3 lines:** T1 and T3 lines run over standard copper phone lines. A T1 line has a connection speed of up to 1.544 Mbps. A T3 line is a bit faster: It transmits data at 44.184 Mbps. Of course, T3 lines are more expensive than T1 lines.

If you don't have enough users to justify the expense of an entire T1 or T3 line, you can lease just a portion of the line. With a fractional T1 line, you can get connections with speeds of 128 Kbps to 768 Kbps; with a fractional T3 line, you can choose speeds ranging from 4.6 Mbps to 32 Mbps.

You may be wondering whether T1 or T3 lines are really any faster than cable or DSL connections. After all, T1 runs at 1.544 Mbps and T3 runs at 44.184 Mbps, and cable and DSL claim to run at comparable speeds. But there are many differences that justify the substantial extra cost of a T1 or T3 line. In particular, a T1 or T3 line is a *dedicated* line — not shared by any other users. T1 and T3 are higher-quality connections, so you actually get the 1.544 or 44.184 connection speeds. In contrast, both cable and DSL connections usually run at substantially



TIP

less than their advertised maximum speeds because of poor-quality connections and because the connections are often shared with other users.

Also, both cable and DSL connections download data much faster than they upload data. So, while you may be able to download data at 100 Mbps over a cable connection, you'll be lucky to upload data at much more than 7 or 8 Mbps.

- » **Business-class cable:** Cable TV providers (such as Comcast) offer business-class service on their cable network. The price and speed depends on your area. For example, where I live, I can get 100Mbps service for about \$400/month.

Like residential cable, the upload speed for business-class cable is usually much slower than the download speed. For example, a typical plan that allows 100Mbps for downloads can support only 10Mbps for uploads. Thus, if you need to upload large amounts of data, you'll notice the performance drop.

Another drawback of business-class cable service is that it is, well, cable service. Your Internet connection is service by the same people who service cable TV in your community. Although business-class customers get priority service over residential customers, business-class service usually does not include response-time guarantees the way T1/T3 or fiber service does. So if your connection goes down, you might find yourself down for hours or even a few days instead of minutes or, at worse, a few hours.

- » **Fiber optic:** The fastest, most reliable, and most expensive form of Internet connection is fiber optic. Fiber optic cable uses strands of glass to transmit data over light signals at very high speeds. Because the light signals traveling within the fiber cables are not subject to electromagnetic interference, fiber connections are extremely reliable; about the only thing that can interrupt a fiber connection is if someone physically cuts the wire.

Fiber is also very expensive. A 50 Mbps fiber connection can cost well over \$1,000 per month. However, the connection is extremely reliable, and response time to service interruptions is measured in minutes instead of hours.

- » **Wireless providers:** In areas where wired service (such as cable or fiber) is not available, you may be able to find wireless service, which provides Internet access using cellular or other wireless technology.

Sharing an Internet connection

After you choose a method to connect to the Internet, you can turn your attention to setting up the connection so that more than one user on your network can share

it. The best way to do that is by using a separate device called a *router*. You can pick up an inexpensive router for a small network for less than \$75. Routers suitable for larger networks will, naturally, cost a bit more.

Because all communications between your network and the Internet must go through the router, the router is a natural place to provide the security measures necessary to keep your network safe from the many perils of the Internet. As a result, a router used for Internet connections often doubles as a firewall, as described in the section “Using a firewall,” later in this chapter.

Securing Your Connection with a Firewall

If your network is connected to the Internet, a whole host of security issues bubbles to the surface. You probably connected your network to the Internet so that your network’s users could get out to the Internet. Unfortunately, however, your Internet connection is a two-way street. It not only enables your network’s users to step outside the bounds of your network to access the Internet, but it also enables others to step in and access your network.

And step in they will. The world is filled with hackers who are looking for networks like yours to break into. They may do it just for the fun of it, or they may do it to steal your customers’ credit card numbers or to coerce your mail server into sending thousands of spam messages on behalf of the bad guys. Whatever their motive, rest assured that your network will be broken into if you leave it unprotected.

Using a firewall

A *firewall* is a security-conscious router that sits between the Internet and your network with a single-minded task: preventing *them* from getting to *us*. The firewall acts as a security guard between the Internet and your local area network (LAN). All network traffic into and out of the LAN must pass through the firewall, which prevents unauthorized access to the network.



WARNING

Some type of firewall is a must-have if your network has a connection to the Internet, whether that connection is broadband (cable modem or DSL), T1, or some other high-speed connection. Without it, sooner or later a hacker will discover your unprotected network and tell his friends about it, and within a few hours, your network will be toast.

You can set up a firewall in two basic ways:

» **Firewall appliance:** The easiest way, and usually the best choice. A firewall appliance is basically a self-contained router with built-in firewall features.

Most firewall appliances include web-based interfaces that enable you to connect to the firewall from any computer on your network by using a browser. You can then customize the firewall settings to suit your needs.

» **Server computer:** Can be set up to function as a firewall computer.

The server can run just about any network operating system, but most dedicated firewall systems run Linux.

Whether you use a firewall appliance or a firewall computer, the firewall must be located between your network and the Internet, as shown in Figure 8-1. Here, one end of the firewall is connected to a network hub, which is, in turn, connected to the other computers on the network. The other end of the firewall is connected to the Internet. As a result, all traffic from the LAN to the Internet (and vice versa) must travel through the firewall.

The term *perimeter* is sometimes used to describe the location of a firewall on your network. In short, a firewall is like a perimeter fence that completely surrounds your property and forces all visitors to enter through the front gate.

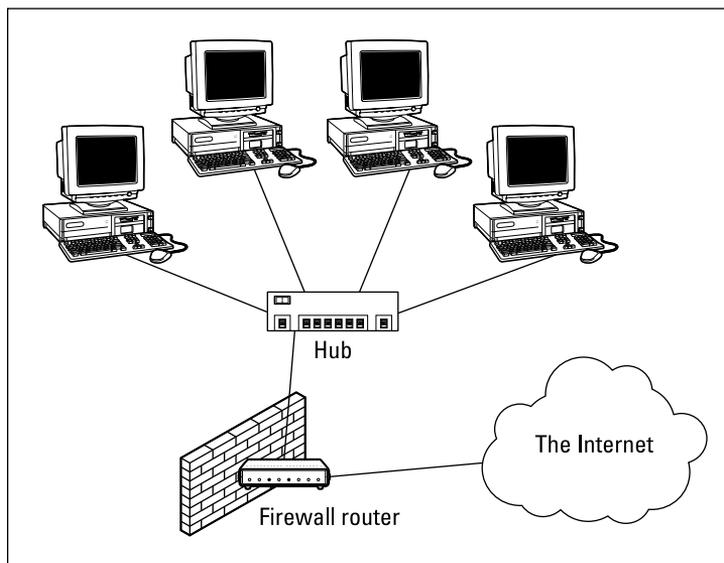


FIGURE 8-1: A firewall router creates a secure link between a network and the Internet.



WARNING

In large networks, figuring out exactly where the perimeter is located can be a little difficult. If your network has two or more Internet connections, make sure that every one of those connections connects to a firewall — and not directly to the network. You can do this by providing a separate firewall for each Internet connection or by using a firewall with more than one Internet port.



TIP

Some firewall routers can also enforce virus protection for your network. For more information about virus protection, see Chapter 20.

The built-in Windows firewall

Windows includes a built-in firewall that provides basic packet-filtering firewall protection. In most cases, you're better off using a dedicated firewall router because these devices provide better security features than the built-in Windows firewall does. Still, the built-in firewall is suitable for home networks or very small office networks.

Here are the steps that activate the built-in firewall in Windows:

1. **Open the Control Panel.**
2. **Click the System and Security link.**

The System and Security page appears.

3. **Click the Windows Firewall link.**

The Windows Firewall page appears.

4. **Click the Turn Windows Firewall On or Off link.**

The page shown in Figure 8-2 appears.

5. **Select the Turn on Windows Firewall option.**

Note that you can independently turn the firewall on or off for public network — that is, for your connection to the Internet — and for your home or work network — that is, if you have a network that connects other computers in your home or office. I recommend you either turn the firewall on for both or turn it off for both. Turn the firewall off if you're using a separate firewall built into the router that connects your computer or home or work network to the Internet. Turn the firewall on if you don't have a separate firewall.

I also recommend leaving the Notify Me When Windows Firewall Blocks a New Program option enabled. That way, you'll be notified when the firewall blocks a suspicious program.

6. **Click OK.**

The firewall is enabled.

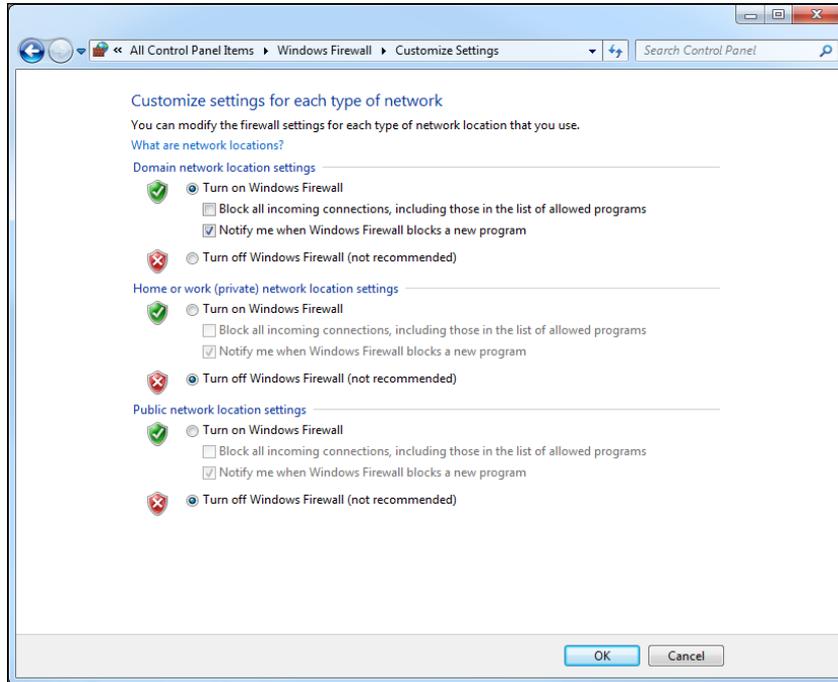


FIGURE 8-2:
Activating the
Windows firewall.

Note that the Windows firewall includes additional options you can configure. However, I recommend against fiddling with those options unless you've taken an upper-division college course in computer security.



WARNING

Do *not* enable the Windows firewall if you're using a separate firewall router to protect your network. Because the other computers on the network are connected directly to the router and not to your computer, the Windows firewall doesn't protect the rest of the network. Additionally, as an unwanted side effect, the rest of the network will lose the capability of accessing your computer.

IN THIS CHAPTER

Understanding wireless network standards

Reviewing basic radio terms

Considering infrastructure and ad-hoc networks

Working with a wireless access point

Configuring Windows for wireless networking

Securing a wireless network

Chapter 9

Setting Up a Wireless Network

Since the beginning of Ethernet networking, cable has been getting smaller and easier to work with. The original Ethernet cable was about as thick as your thumb, weighed a ton, and was difficult to bend around tight corners. Then came coaxial cable, which was lighter and easier to work with. Coaxial cable was supplanted by unshielded twisted-pair (UTP) cable, which is the cable used for most networks today.

Although cable through the years has become smaller, cheaper, and easier to work with, it is still *cable*. So you have to drill holes in walls, pull cable through ceilings, and get insulation in your hair to wire your entire home or office.

The alternative to networking with cables is, of course, networking *without* cables. . . also known as *wireless networking*. Wireless networks use radio waves to send and receive network signals. As a result, a computer can connect to a wireless network at any location in your home or office.

Wireless networks are especially useful for notebook computers. After all, the main benefit of a notebook computer is that you can carry it around with you wherever you go. At work, you can use your notebook computer at your desk, in the conference room, in the break room, or even out in the parking lot. At home, you can use it in the bedroom, kitchen, den, or game room, or out by the pool. With wireless networking, your notebook computer can be connected to the network no matter where you take it.

Wireless networks have also become extremely useful for other types of mobile devices, such as smartphones and tablet computers. Sure, these devices can connect via a cell network, but that can get real pricey real quick. With a wireless network, though, you can connect your smartphone or tablet without having to pay your cellphone company for the connection time.

This chapter introduces you to the ins and outs of setting up a wireless network. I tell you what you need to know about wireless networking standards, how to plan your wireless network, and how to install and configure wireless network components. And if you end up with a hybrid network of wired and wireless, I show you how to create that, too.

Diving into Wireless Networking

As I mention earlier, a wireless network is just a network that uses radio signals rather than direct cable connections to exchange information. Simple as that. A computer with a wireless network connection is like a cellphone. Just as you don't have to be connected (tethered) to a phone line to use a cellphone, you don't have to be connected to a network cable to use a wireless networked computer.

Here are the key concepts and terms you need to understand to set up and use a basic wireless network:

- » **WLAN:** A wireless network is often referred to as a wireless local area network (WLAN). Some people prefer to switch the acronym around to local area wireless network, or LAWN.
- » **Wi-Fi:** The term *Wi-Fi* is often used to describe wireless networks although it technically refers to just one form of wireless network: the 802.11 standard. (See the section "Eight-Oh-Two-Dot-Eleventy Something?: Understanding Wireless Standards," later in this chapter for more information.)
- » **SSID:** A wireless network has a name, known as a SSID. *SSID* stands for *service set identifier*. (Wouldn't that make a great *Jeopardy!* question? I'll take obscure

four-letter acronyms for \$400, please!) All the computers that belong to a single wireless network must have the same SSID.

- » **Channels:** Wireless networks can transmit over any of several channels. For computers to talk to one another, though, they must be configured to transmit on the same channel.
- » **Ad-hoc:** The simplest type of wireless network consists of two or more computers with wireless network adapters. This type of network is an *ad-hoc mode network*.
- » **Infrastructure mode:** A more complex type of network is an infrastructure mode network. All this really means is that a group of wireless computers can be connected not only to one another, but also to an existing cabled network via a device called a *wireless access point (WAP)*. (I tell you more about ad-hoc and infrastructure networks later in this chapter.)

A Little High School Electronics

I was a real nerd in high school: I took three years of electronics. The electronics class at my school was right next door to the auto shop. All the cool kids took auto shop, of course, and only nerds like me took electronics. We hung in there, though, and learned all about capacitors and diodes while the cool kids were learning how to raise their cars and install 2-gigawatt stereo systems.

It turns out that a little of that high school electronics information proves useful when it comes to wireless networking — not much, but a little. You'll understand wireless networking much better if you know the meanings of some basic radio terms.

Waves and frequencies

For starters, radio consists of electromagnetic waves sent through the atmosphere. You can't see or hear them, but radio receivers can pick them up and convert them to sounds, images, or — in the case of wireless networks — data. Radio waves are actually cyclical waves of electronic energy that repeat at a particular rate: the *frequency*.

Figure 9-1 shows two frequencies of radio waves. The first is one cycle per second; the second is two cycles per second. (Real radio doesn't operate at that low a frequency, but I figured that one and two cycles per second would be easier to draw than 680,000 and 2.4 million cycles per second.)

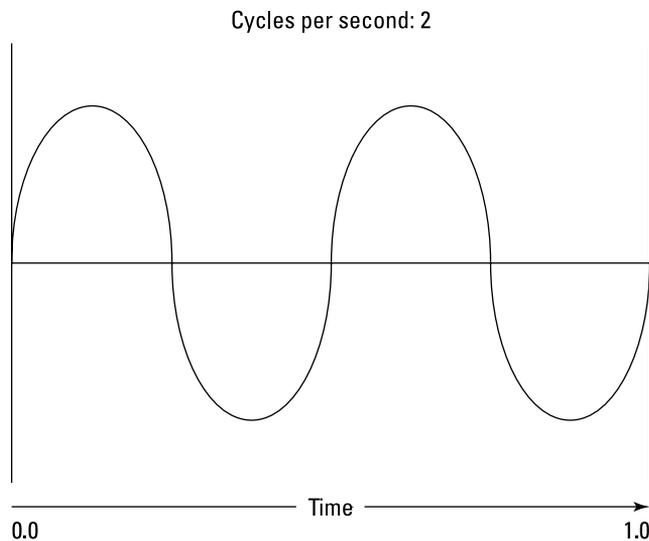
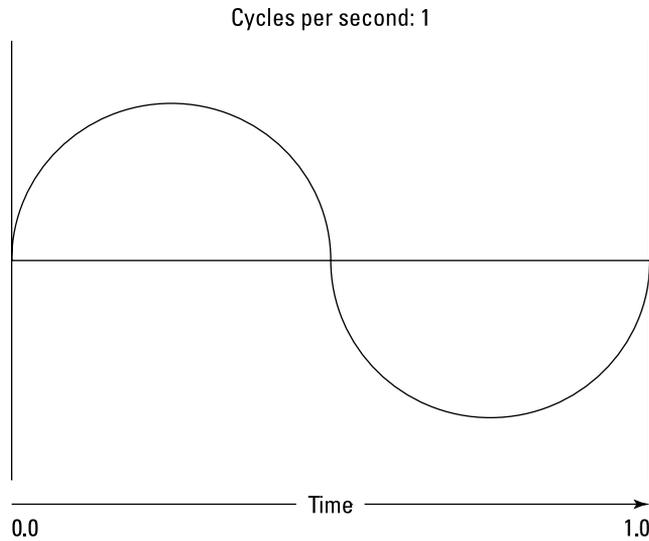


FIGURE 9-1:
Radio waves
frequently have
frequency.



TIP

The measure of a frequency is *cycles per second*, which indicates how many complete cycles the wave makes in 1 second (duh). In honor of Heinrich Hertz — who was the first person to successfully send and receive radio waves (it happened in the 1880s) — cycles per second is usually referred to as *Hertz*, abbreviated *Hz*. Thus, 1 Hz is one cycle per second.



TECHNICAL
STUFF

Incidentally, when the prefix *K* (for *kilo*, or 1,000), *M* (for *mega*, 1 million), or *G* (for *giga*, 1 billion) is added to the front of Hz, the *H* is still capitalized. Thus, 2.4 MHz is correct (not 2.4 Mhz).

So the beauty of radio frequencies is that transmitters can be tuned to broadcast radio waves at a precise frequency. Likewise, receivers can be tuned to receive radio waves at a precise frequency, ignoring waves at other frequencies. That's why you can tune the radio in your car to listen to dozens of radio stations: Each station broadcasts at its own frequency.

Wavelength and antennas

A term related to frequency is *wavelength*. Radio waves travel at the speed of light, and *wavelength* refers to how far the radio signal travels with each cycle. Or, put another way, *wavelength* refers to the physical distance between the crest of each wave. Because the speed of light is roughly 182,282 miles per second, for example, the wavelength of a 1 Hz radio wave is about 182,282 miles. The wavelength of a 2 Hz signal is about half that: a mere 91,141 miles.

As you can see, the wavelength decreases as the frequency increases. The wavelength of a typical AM radio station broadcasting at 580 KHz is about 522 yards. For a TV station broadcasting at 100 MHz, it's about 3 yards. For a wireless network broadcasting at 2.4 GHz, the wavelength is just shorter than 5 inches.

And the shorter the wavelength, the smaller the antenna needs to be to adequately receive the signal. As a result, higher-frequency transmissions need smaller antennas. You may have noticed that AM radio stations usually have huge antennas mounted on top of tall towers, but cellphone transmitters are much smaller, and their towers aren't nearly as tall because cellphones operate on a higher frequency than AM radio stations do. So who decides what type of radio gets to use specific frequencies? That's where spectrums and the FCC come in.

Spectrums and the FCC

Spectrum refers to a continuous range of frequencies on which radio can operate. In the United States, the Federal Communications Commission (FCC) regulates not only how much of Janet Jackson can be shown at the Super Bowl, but also how various portions of the radio spectrum can be used. Essentially, the FCC has divided the radio spectrum into dozens of small ranges — *bands* — and restricted certain uses to certain bands. AM radio, for example, operates in the band from 535 KHz to 1,700 KHz.

Table 9-1 lists some of the most popular bands. Note that some of these bands are wide — UHF television begins at 470 MHz and ends at 806 MHz — but other bands are restricted to a specific frequency. The difference between the lowest and highest frequency within a band is the *bandwidth*.

TABLE 9-1 Popular Bands of the Radio Spectrum

Band	Use
535 KHz–1,700 KHz	AM radio
5.9 MHz–26.1 MHz	Shortwave radio
26.96 MHz–27.41 MHz	Citizens Band (CB) radio
54 MHz–88 MHz	Television (VHF channels 2–6)
88 MHz–108 MHz	FM radio
174 MHz–220 MHz	Television (VHF channels 7–13)
470 MHz–806 MHz	Television (UHF channels)
806 MHz–890 MHz	Cellular networks
900 MHz	Cordless phones and wireless networks (802.11ah)
1850 MHz–1990 MHz	PCS cellular
2.4 GHz–2.4835 GHz	Cordless phones and wireless networks (802.11b and 802.11n)
4 GHz–5 GHz	Large-dish satellite TV
5 GHz	Wireless networks (802.11a)
11.7 GHz–12.7 GHz	Small-dish satellite TV

AND NOW, A WORD FROM THE IRONY DEPARTMENT

I was an English-literature major in college, so I like to use literary devices such as irony. I don't get to use it much in the computer books I write, so when I get the chance to use irony, I like to jump on it like a hog out of water.

So here's my juicy bit of irony for today: The very first Ethernet system was actually a wireless network. Ethernet traces its roots to a network called AlohaNet, developed at the University of Hawaii in 1970. This network transmitted its data by using small radios. If two computers tried to broadcast data at the same time, the computers detected the collision and tried again after a short, random delay. This technique was the inspiration for the basic technique of Ethernet, now called "carrier sense multiple access with collision detection" or CSMA/CD. The wireless AlohaNet was the network that inspired Robert Metcalfe to develop his cabled network, which he called Ethernet, as his doctoral thesis at Harvard in 1973.

For the next 20 years or so, Ethernet was pretty much a cable-only network. It wasn't until the mid-1990s that Ethernet finally returned to its wireless roots.

Two of the bands in the spectrum are allocated for use by wireless networks: 2.4 GHz and 5 GHz. Note that these bands aren't devoted exclusively to wireless networks. In particular, the 2.4 GHz band shares its space with cordless phones. As a result, cordless phones sometimes interfere with wireless networks. Note also that, as of 2016, some wireless networks can also operate in the 900 MHz spectrum.

Eight-Oh-Two-Dot-Eleventy Something?: Understanding Wireless Standards

The most popular standards for wireless networks are the IEEE 802.11 standards. These standards are essential wireless Ethernet standards and use many of the same networking techniques that the cabled Ethernet standards (in other words, 802.3) use. Most notably, 802.11 networks use the same CSMA/CD technique as cabled Ethernet to recover from network collisions.

The 802.11 standards address the bottom two layers of the IEEE seven-layer model: the Physical layer and the Media Access Control (MAC) layer. Note that TCP/IP protocols apply to higher layers of the model. As a result, TCP/IP runs just fine on 802.11 networks.

The original 802.11 standard was adopted in 1997. Two additions to the standard, 802.11a and 802.11b, were adopted in 1999. Then came 802.11g in 2003 and 802.11n in 2009.

802.11n ruled the roost for a few years, until the latest to gain widespread acceptance came out in 2014: 802.11ac. Still more variations are in the works, including 802.11ah, which will operate in the 900 MHz spectrum.

Table 9-2 summarizes the basic characteristics of the five most popular variants of 802.11 as of early 2016. Currently, most wireless networks are based on the 802.11n and 802.11ac standards.

TABLE 9-2 802.11 Variations

Standard	Speeds	Frequency	Typical Range (Indoors)
802.11a	Up to 54 Mbps	5 GHz	150 feet
802.11b	Up to 11 Mbps	2.4 GHz	300 feet
802.11g	Up to 54 Mbps	2.4 GHz	300 feet
802.11n	Up to 600 Mbps (but most devices are in the 100 Mbps range)	2.4 GHz	230 feet
802.11ac	Up to 1,300 Mbps	5 GHz	230 feet

Home on the Range

The maximum range of an 802.11ac wireless device indoors is about 230 feet. This can have an interesting effect when you get a bunch of wireless computers together such that some of them are in range of one another but others are not. Suppose that Wally, Ward, and the Beaver all have wireless notebooks. Wally's computer is 150 feet away from Ward's computer, and Ward's computer is 150 feet away from Beaver's in the opposite direction (see Figure 9-2). In this case, Ward can access both Wally's and Beaver's computers, but Wally can access only Ward's computer, and Beaver can access only Ward's computer. In other words, Wally and Beaver won't be able to access each other's computers because they are 300 feet away from each other, well beyond the 230-foot range limit. (This is starting to sound suspiciously like an algebra problem. Now suppose that Wally starts walking toward Ward at 2 miles per hour, and Beaver starts running toward Ward at 4 miles per hour. . . .)

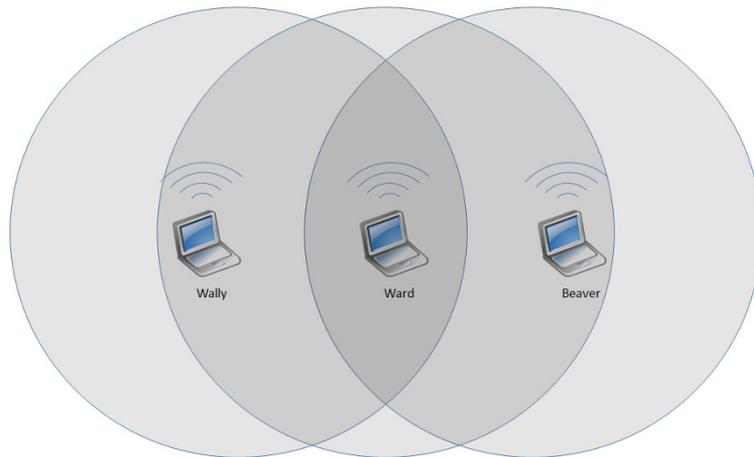


FIGURE 9-2: Ward, Wally, and the Beaver playing with their wireless network.

Note: Although the normal range for 802.11ac is 230 feet, the range may be less in actual practice. Obstacles — solid walls, bad weather, cordless phones, microwave ovens, backyard nuclear reactors, and so on — can all conspire to reduce the effective range of a wireless adapter. If you're having trouble connecting to the network, sometimes just adjusting the antenna helps.

Also, wireless networks tend to slow down when the distance increases. 802.11ac network devices claim to operate at 1,300 Mbps, but they usually achieve that speed only at close range. The farther out you get, the slower the actual speed becomes. At maximum distance, you may be able to connect, but your effective

connection speed will be slow. You should also realize that when you're at the edge of the wireless device's range, you're more likely to lose your connection suddenly due to bad weather.

Using Wireless Network Adapters

Each computer that will connect to your wireless network needs a wireless network adapter, which is similar to the network interface card (NIC) used for a standard Ethernet connection. Instead of having a cable connector on the back, however, a wireless network adapter has an antenna.

Just about all portable computers, such as notebooks and tablets, come with wireless networking built in, so you don't have to add a separate wireless network adapter to a portable computer. Desktop computers, though, are a different story. They often don't have built-in wireless networking. If yours doesn't, you'll have to purchase one of two types of wireless adapters:

- » **A wireless PCI card:** You install this wireless network adapter in an available slot inside your desktop computer. Yup, you need to take your computer apart, so use this type of card only if you have the expertise and the nerves to dig into your computer's guts. A typical 802.11 ac version costs about \$75.
- » **A wireless USB adapter:** This gizmo is a separate device that plugs into a USB port on your computer. A USB adapter should cost about the same as a PCI adapter. And you can install it without taking your computer apart.

Setting Wireless Access Points

Unlike cabled networks, wireless networks don't need a separate switch. If all you want to do is network a group of wireless computers and the computers are close to one another, you just purchase a wireless adapter for each computer and — *voilà!* — instant network.

But what if you already have an existing cabled network? Suppose that you work at an office with 15 computers all cabled up nicely, and you just want to add a couple of wireless notebook computers to the network. Or suppose that you have two computers in your den connected with network cable, but you want to link up a computer in your bedroom without pulling cable through the attic.

That's where a WAP comes in. A WAP actually performs two functions:

- » It acts as a central connection point for all your computers that have wireless network adapters. In effect, the WAP performs the same function that a network switch performs for a wired network.
- » It links your wireless network to your existing wired network so that your wired computer and your wireless computers get along like one big happy family. This sounds like the makings of a Dr. Seuss story. ("Now the wireless sneeches had hubs without wires. But the twisted-pair sneeches had cables to thires. . . .")



TIP

Wireless access points are sometimes just called access points (APs), which is basically a box with an antenna (or often a pair of antennae) and an RJ-45 Ethernet port. You plug the AP into a network cable and then plug the other end of the cable into a hub or switch, and your wireless network should be able to connect to your cabled network.

Figure 9-3 shows how an access point acts as a central connection point for wireless computers and also how it bridges your wireless network to your wired network.

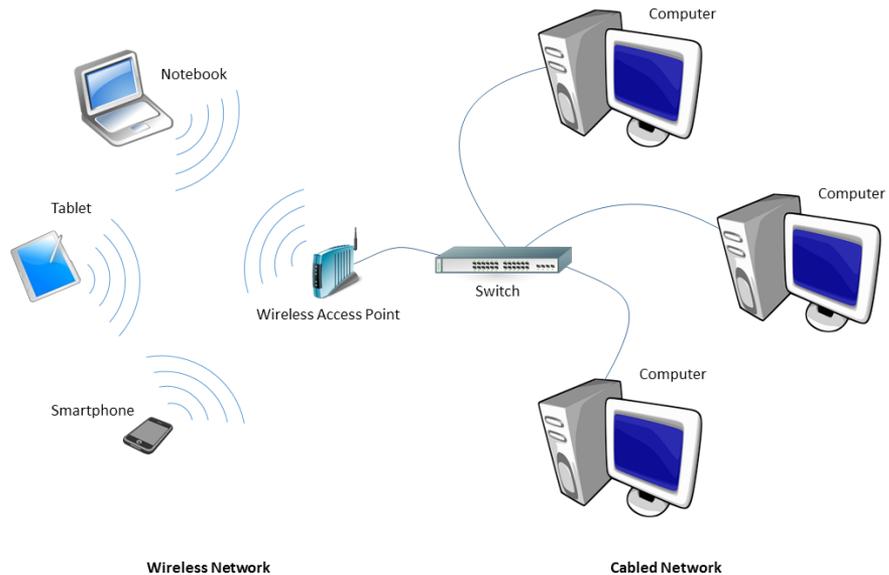


FIGURE 9-3: A wireless access point connects a wireless network to a cabled network.

Infrastructure mode

When you set up a wireless network with an AP, you're creating an *infrastructure mode network*: The AP provides a permanent infrastructure for the network. The

APs are installed at fixed physical locations, so the network has relatively stable boundaries. Whenever a mobile computer wanders into the range of one of the APs, it has come into the sphere of the network and can connect.

An access point and all the wireless computers that are connected to it are called a *Basic Service Set* (BSS). Each BSS is identified by a *Service Set Identifier* (SSID). When you configure an access point, you specify the SSID that you want to use. The default SSID on most access points is a generic name that often includes the name of the manufacturer of the device. However, you can easily change the SSID to something more meaningful.

Multifunction WAPs



TIP

Wireless access points often include other built-in features. Some access points double as Ethernet switches, sporting a bank of RJ-45 ports that you can plug other computers or devices into. In addition, some access ports include broadband cable or DSL firewall routers that enable you to connect to the Internet. Figure 9-4 shows a Linksys EA8500, an 802.11ac wireless router intended for home or small office use. This device includes the following features:

- » An 802.11ac wireless access point that can support multiple wireless devices
- » A router with firewall capabilities that can be connected directly to the Internet output from a cable or DSL router
- » A four-port gigabit Ethernet switch to connect cabled computers or other devices
- » Two USB ports that enable you to connect USB printers or disk drives to your network



FIGURE 9-4:
A typical wireless router.



A multifunction access point designed to serve as an Internet gateway for home networks is sometimes called a *wireless router* or a *residential gateway*.

Roaming Capabilities

You can use two or more wireless access points to create a large wireless network in which computer users can roam from area to area and still be connected to the wireless network. As the user moves out of the range of one access point, another access point automatically picks up the user and takes over without interrupting the user's network service.

To set up two or more access points for roaming, you must carefully place the access points so that all areas of the office or building that are being networked are in range of at least one of the access points. Then just make sure that all the computers and access points use the same SSID.

Two or more access points joined for roaming, along with all the wireless computers connected to any of the access points, form an *Extended Service Set (ESS)*. The access points in the ESS are usually connected to a wired network.

One limitation of roaming is that each access point in an ESS must be on the same TCP/IP subnet. That way, a computer that roams from one access point to another within the ESS retains the same IP address. If the access points had a different subnet, a roaming computer would have to change IP addresses when it moved from one access point to another.

Wireless bridging

Another use for wireless APs is to bridge separate subnets that can't easily be connected by cable. Suppose that you have two office buildings that are only about 50 feet apart. To run cable from one building to the other, you'd have to bury conduit — a potentially expensive job. Because the buildings are so close, though, you can probably connect them with a pair of wireless access points that function as a *wireless bridge* between the two networks. Connect one of the access points to the first network and the other access points to the second network. Then configure both APs to use the same SSID and channel.

Ad-hoc networks

A WAP isn't necessary to set up a wireless network. Any time two or more wireless devices come within range of each other, they can link up to form an ad-hoc

network. If you and a few of your friends all have notebook computers with wireless adapters, for example, you can meet anywhere and form an ad-hoc network.

All the computers within range of one another in an ad-hoc network are an *Independent Basic Service Set* (IBSS).

Configuring a Wireless Access Point

The physical setup for a wireless access point is pretty simple: You take it out of the box, put it on a shelf or on top of a bookcase near a network jack and a power outlet, plug in the power cable, and plug in the network cable.

The software configuration for an access is a little more involved but still not very complicated. It's usually done via a web interface. To get to the configuration page for the access, you need to know its IP address. Then you just type that address in the address bar of a browser on any computer on the network.

Multifunction access points usually provide DHCP and NAT services for the networks and double as the network's gateway router. As a result, they typically have a private IP address that's either at the beginning of one of the Internet's private IP address ranges, typically 192.168.0.1 or 10.0.0.1. Consult the documentation that came with the AP to find out more.



TIP

If you use a multifunction AP that serves as both your wireless AP and your Internet router, and you can't remember the IP address, run the `IPCONFIG` command at a command prompt on any computer on the network. The Default Gateway IP address should be the IP address of the access point.

Basic configuration options

Figure 9-5 shows the main configuration screen for a typical router. I called up this configuration page by entering 192.168.1.1 in the address bar of a web browser and then supplying the login password when prompted.

On the main setup page of this router, you configure information such as the hostname and IP address of the router and whether the router's DHCP server should be enabled. Options found on additional tabs allow you to configure wireless settings, such as the network name (SSID), the type of security to enforce, and a variety of other settings.

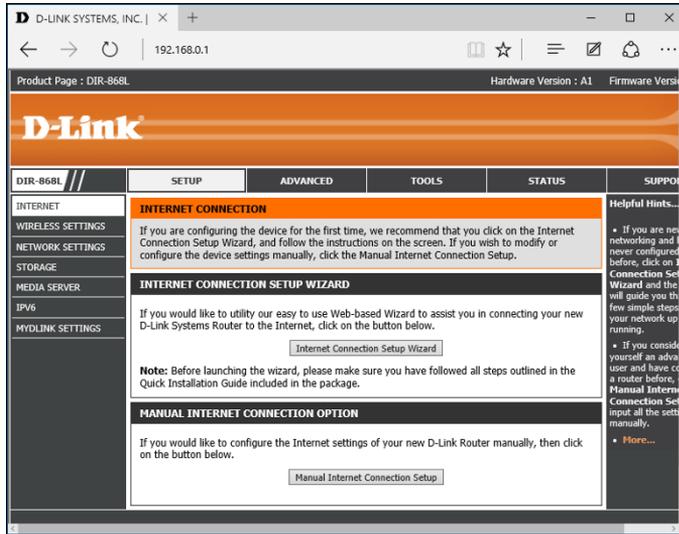


FIGURE 9-5:
The main configuration page for a typical wireless router.

DHCP configuration

You can configure most multifunction access points to operate as a DHCP server. For small networks, the access point is commonly the DHCP server for the entire network. In that case, you need to configure the access point's DHCP server.

Figure 9-6 shows the DHCP configuration page for a typical wireless router. To enable DHCP, you select the Enable option and then specify the other configuration options to use for the DHCP server.

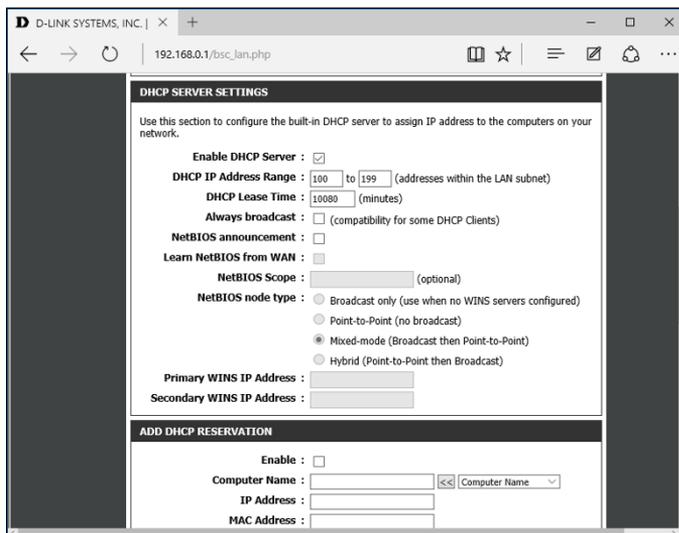


FIGURE 9-6:
Configuring DHCP for a typical wireless router.

Larger networks with more demanding DHCP requirements are likely to have a separate DHCP server running on another computer. In that case, you can defer to the existing server by disabling the DHCP server in the access point.

For more information on configuring a DHCP server, refer to Chapter 5.

Connecting to a Wireless Network

Connecting to a wireless network on a Windows computer is straightforward. Windows automatically detects any wireless networks that are in range and displays them in a list when you tap the Wireless icon at the bottom of the screen, as shown in Figure 9-7.

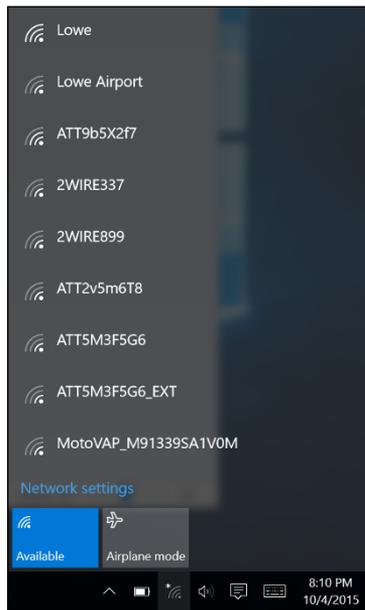


FIGURE 9-7:
Choosing a wireless network.

To connect to a network, just tap it, and then enter the security key when prompted. If the key is correct, you'll be connected.

At the time you connect, you can choose to connect to the network automatically whenever it's in range. If you select this option, you won't have to select the network manually or enter the security key; you'll just be connected automatically.

Windows remembers every network you connect to, which is a plus for networks you frequently use but a drawback for networks you'll likely never use again. To tell Windows to forget a network, follow these steps:

1. Click Start, and then click Settings.

The Settings window appears.

2. Click Network & Internet.

This brings up the Network & Internet page, which lists the known networks.

3. Scroll to the bottom of the list of known networks, and then click Manage Wi-Fi Settings.

This brings up the Manage Wi-Fi Settings page, which includes a section titled Manage Known Networks.

4. In the Manage Known Networks section, click the network you want to forget.

The network is selected, as shown in Figure 9-8.

5. Click Forget.

The network will be forgotten. To log into this network again, you'll have to enter the security key.

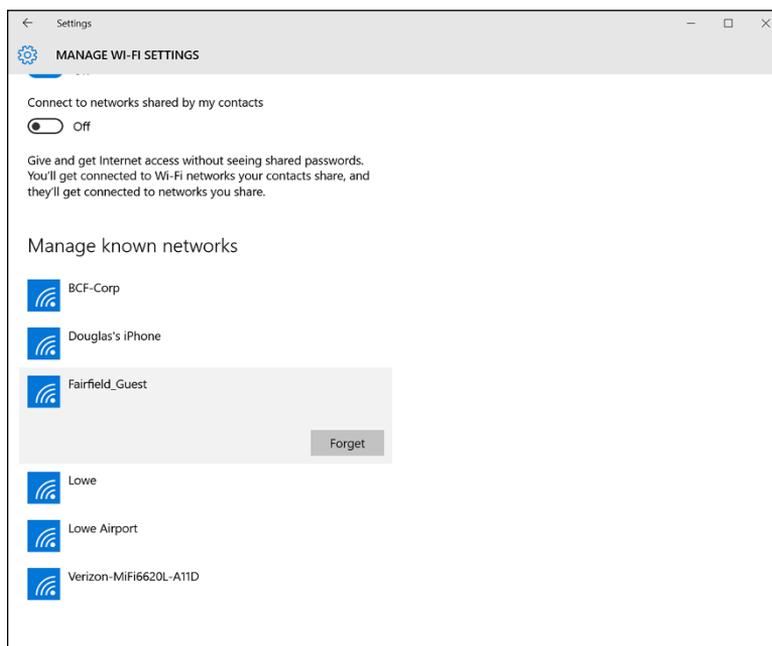


FIGURE 9-8:
Forgetting a wireless network in Windows 10.

Paying Attention to Wireless Network Security

Before you unleash a wireless access point on your network, you should first consider the inherent security risks that come along with wireless networking. Unless you take a few basic precautions, adding a wireless access point will expose the innards of your network to anyone within range.

Refer to Chapter 19 for details about network security in general. The information in this section will help you prevent unwanted visitors from gaining access to your network via your wireless access point, but the security techniques outlined in Chapter 19 are a must regardless of whether you provide wireless access.

The following paragraphs describe the types of security threats that wireless networks are most likely to encounter. You should take each of these kinds of threats into consideration when you plan your network's security:

» **Intruders:** With a wired network, an intruder usually must gain access to your facility to physically connect to your network. That's not so with a wireless network. Anyone with a wireless device can gain access to your network if she can get within range of your network's radio signals. At home, your neighbors can probably see your wireless network. And in an office, kids sitting on a bench outside your building can probably see your wireless network.

» **Freeloaders:** *Freeloaders* are intruders who want to piggyback on your wireless network to get free access to the Internet. If they manage to gain access to your wireless network, they probably won't do anything malicious: They'll just fire up their web browsers and surf the web. These are folks who are too cheap to spend \$40 per month on their own broadband connection at home, so they'd rather drive into your parking lot and steal yours.

Even though freeloaders may be relatively benign, they can be a potential source of trouble. They suck up your bandwidth. They may use your network to download illegal pornography, or they may try to hijack your email server to send spam. And they may start out innocently looking for free Internet access, but their curiosity may grow once they get in, leading them to snoop around your network.

» **Eavesdroppers:** *Eavesdroppers* just like to listen to your network traffic. They don't actually try to gain access via your wireless network — at least, not at first. They just listen. They spy on the packets that you're sending over the wireless network, hoping to find useful information such as passwords or credit card numbers.

» **Spoilers:** A *spoiler* is a hacker who gets his kicks from jamming networks so that they become unusable. A spoiler usually accomplishes this act by flooding the network with meaningless traffic so that legitimate traffic gets lost in the flow. Spoilers may also try to place viruses or worm programs on your network via an unsecured wireless connection.



ROGUE ACCESS POINTS

One of the biggest problems that business networks face is the problem of *rogue access points*, which are access points that suddenly appears on your network out of nowhere. What usually happens is that an employee wants to connect his iPad or smartphone to your company network, but you won't give him the password. So the user stops at Walmart on the way home from work one day, buys a cheap wireless router, and plugs it into your network without asking permission.

Now, in spite of all the elaborate security precautions you've taken to fence in your network, this well-meaning user has opened the barn door. It's *very* unlikely that the user will enable the security features of the wireless access point; in fact, he probably isn't even aware that wireless access devices *have* security features.

Unless you take some kind of action to find it, a rogue access point can operate undetected on your network for months or even years. You may not discover it until you report to work one day and find that your network has been trashed by an intruder who found his way into your network via an unprotected wireless access point that you didn't even know existed.

Here are some steps you can take to reduce the risk of rogue access points appearing on your system:

- **Establish a policy prohibiting users from installing wireless access points on their own.** Then make sure that you inform all network users of the policy, and let them know why installing an access point on their own can be such a major problem.
- **If possible, establish a program that quickly and inexpensively grants wireless access to users who want it.** Rogue access points show up in the first place because users who want access can't get it. If you make it easier for users to get legitimate wireless access, you're less likely to find wireless access points hidden behind file cabinets or in flower pots.
- **Once in a while, take a walk through the premises, looking for rogue access points.** Take a look at every network outlet in the building and see what's connected to it.
- **Turn off all your wireless access points and then walk around the premises with a wireless-equipped mobile device such as a smartphone and look for wireless networks that pop up.** Just because you detect a wireless network, of course, doesn't mean you've found a rogue access point — you may have stumbled onto a wireless network in a nearby office or home. But knowing what wireless networks are available from within your office will help you determine whether any rogue access points exist.

Hopefully I've convinced you that wireless networks do, indeed, pose many security risks. Here are some steps you can take to help secure your wireless network:

- » **Create a secure wireless password.** The first thing you should do when you set up a wireless network is change the default password required to access the network. Most manufacturers of wireless routers secure the SSID with a standard password that is well known. Make sure you change it to something that only you know, and then only share that password with those you want to grant access to.
- » **Change the administrative password.** Most access points have a web-based setup page that you can access from any web browser to configure the access point's settings. The setup page is protected by a username and password, but the username and password are initially set to default values that are easy to guess. Anyone who gains access to your network can then log in to the administrative page and take control of your network.
- » **Hide the SSID.** A simple step you can take to secure your wireless network is to disable the automatic broadcast of your network's SSID. That way, only those who know of your network's existence will be able to access it. (Securing the SSID isn't a complete security solution, so you shouldn't rely on it as your only security mechanism.)
- » **Disable guest mode.** Many access points have a guest-mode feature that enables client computers to specify a blank SSID or to specify "any" as the SSID. If you want to ensure that only clients that know the SSID can join the network, you must disable this feature.
- » **Use MAC address filtering.** One of the most effective ways to secure a wireless network is to use a technique called *MAC address filtering*. MAC address filtering allows you to specify a list of MAC addresses for the devices that are allowed to access the network or are prohibited from accessing the network. If a computer with a different MAC address tries to join the network via the access point, the access point will deny access.



TIP

MAC address filtering is a great idea for wireless networks with a fixed number of clients. If you set up a wireless network at your office so that a few workers can connect their notebook computers, you can specify the MAC addresses of those computers in the MAC filtering table. Then other computers won't be able to access the network via the access point.

MAC address filtering isn't bulletproof, but it can go a long way toward keeping unwanted visitors off your network. Unfortunately, MAC address filtering is also pretty inconvenient. Whenever you want to grant access for a new device, you'll have to find out that device's MAC address and add it to the list of permitted devices.

» **Place your access points outside the firewall.** The most effective security technique for wireless networking is placing all your wireless access points *outside* your firewall. That way, all network traffic from wireless users will have to travel through the firewall to access the network.

As you can imagine, doing this can significantly limit network access for wireless users. To get around those limitations, you can enable a virtual private network (VPN) connection for your wireless users. The VPN will allow full network access to authorized wireless users.

Obviously, this solution requires a bit of work to set up and can be a little inconvenient for your users, but it's an excellent way to fully secure your wireless access points.

DON'T NEGLECT THE BASICS

The security techniques described in this chapter are specific to wireless networks. They should be used alongside the basic security techniques that are presented in Chapter 19. In other words, don't forget the basics, such as the following:

- Use strong passwords for your user accounts.
- Apply security patches to your servers.
- Change default server account information (especially the administrator password).
- Disable unnecessary services.
- Check your server logs regularly.
- Install virus protection.
- Back up!

IN THIS CHAPTER

Examining the basics of virtualization

Weighing the benefits of virtualization

Looking at what a hypervisor does

Considering how disks and networks are virtualized

Activating Hyper-V

Creating and using virtual machines

Chapter 10

Virtual Networking

Virtualization is one of the hottest trends in networking today. According to some industry pundits, virtualization is the best thing to happen to computers since the invention of the transistor. If you haven't already begun to virtualize your network, you're standing on the platform watching as the train is pulling out.

This chapter is a brief introduction to virtualization, with an emphasis on using it to leverage your network server hardware to provide more servers using less hardware. In addition to the general concepts of virtualization, you find out how to experiment with virtualization by activating Microsoft's Hyper-V, which is a free virtualization product included with Windows.

Understanding Virtualization

The basic idea behind virtualization is to use software to simulate the existence of hardware. This powerful idea enables you to run more than one independent computer system on a single physical computer system. Suppose that your organization requires a total of 12 servers to meet its needs. You could run each of these

12 servers on a separate computer, in which case you would have 12 computers in your server room, or you could use virtualization to run these 12 servers on just 2 computers. In effect, each of those computers would simulate six separate computer systems, each running one of your servers.

Each of the simulated computers is called a *virtual machine* (VM). For all intents and purposes, each virtual machine appears to be a complete, self-contained computer system with its own processor (or, more likely, processors), memory, disk drives, CD-ROM/DVD drives, keyboard, mouse, monitor, network interfaces, USB ports, and so on.

Like a real computer, each virtual machine requires an operating system to do productive work. In a typical network server environment, each virtual machine runs its own copy of Windows Server. The operating system has no idea that it's running on a virtual machine rather than on a real machine.

Here are a few terms you need to be familiar with if you expect to discuss virtualization intelligently:

- » **Host:** The actual physical computer on which one or more virtual machines run.
- » **Bare metal:** Another term for the host computer that runs one or more virtual machines.
- » **Guest:** Another term for a virtual machine running on a host.
- » **Guest operating system:** An operating system that runs within a virtual machine. By itself, a guest is just a machine; it requires an operating system to run. The guest operating system is what brings the guest to life.

As far as licensing is concerned, Microsoft treats each virtual machine as a separate computer. Thus, if you run six guests on a single host, and each guest runs Windows Server, you need six licenses of Windows Server.

- » **Hypervisor:** The virtualization operating system that creates and runs virtual machines. For more information about hypervisors, read the next section, "Understanding Hypervisors."
- » **Hardware abstraction layer (HAL):** A layer of software that acts as a go-between to separate actual hardware from the software that interacts with it. An operating system provides a HAL, because it uses device drivers to communicate with actual hardware devices so that software running in the operating system doesn't have to know the details of the specific device it's interacting with. A hypervisor also provides a HAL that enables the guest operating systems in virtual machines to interact with virtualized hardware.



WARNING

THE LONG TREK OF VIRTUALIZATION

Kids these days think they invented everything, including virtualization.

Little do they know.

Virtualization was developed for PC-based computers in the early 1990s, around the time Captain Picard was flying the Enterprise around in *Star Trek: The Next Generation*.

But the idea is much older than that.

The first virtualized server computers predate Captain Picard by about 20 years. In 1972, IBM released an operating system called simply VM, which had nearly all the basic features found in today's virtualization products.

VM allowed the administrators of IBM's System/370 mainframe computers to create multiple independent virtual machines, each of which was called (you guessed it) a virtual machine, or VM. This terminology is still in use today.

Each VM could run one of the various guest operating systems that were compatible with the System/370 and appeared to this guest operating system to be a complete, independent System/370 computer with its own processor cores, virtual memory, disk partitions, and input/output devices.

The core of the VM system itself was called the *hypervisor* — another term that persists to this day.

The VM product that IBM released in 1972 was actually based on an experimental product that IBM released on a limited basis in 1967.

So whenever someone tells you about this new technology called *virtualization*, you can tell them that it was invented when *Star Trek* was TV. When they ask, "You mean the one with Picard?" you can say, "No, the one with Kirk."

Understanding Hypervisors

At the core of virtualization is a *hypervisor*, a layer of software that manages the creation and execution of virtual machines. A hypervisor provides several core functions:

- » It provides a HAL, which virtualizes all the hardware resources of the host computer on which it runs. This includes processor cores, RAM, and I/O

devices such as disk drives, keyboards, mice, monitors, USB devices, and so on.

- » It creates pools of these abstracted hardware resources that can be allocated to virtual machines.
- » It creates virtual machines. A virtual machine is a complete implementation of an idealized computer system that has the hardware resources of the host available to it. The hardware for each virtual machine is drawn from the pools of available hardware resources managed by the hypervisor.
- » It manages the execution of its virtual machines, allocating host hardware resources as needed to each virtual machine and starting and stopping virtual machines when requested by users.
- » It ensures that each virtual machine is completely isolated from all other virtual machines, so that if a problem develops in one virtual machine, none of the other virtual machines is affected.
- » It manages communication among the virtual machines over virtual networks, enabling the virtual machines to connect with each other and with a physical network that reaches beyond the host.

There are two basic types of hypervisors you should know about:

- » **Type-1:** A type-1 hypervisor runs directly on the host computer, with no intervening operating system. This is the most efficient type of hypervisor because it has direct access to the hardware resources of the host system.

The two best-known examples of type-1 hypervisors are VMware's ESXi and Microsoft's Hyper-V. ESXi is part of a suite of popular virtualization products from VMware, and Hyper-V is the built-in virtualization platform that is included with recent versions of Windows Server.

- » **Type-2:** A type-2 hypervisor runs as an application within an operating system that runs directly on the host computer. Type-2 hypervisors are less efficient than type-1 hypervisors because when you use a type-2 hypervisor, you add an additional layer of hardware abstraction — the first provided by the operating system that runs natively on the host, and the second by the hypervisor that runs as an application on the host operating system.



TIP

For production use, you should always use type-1 hypervisors because they're much more efficient than type-2 hypervisors. Type-1 hypervisors are considerably more expensive than type-2 hypervisors, however. As a result, many people use inexpensive or free type-2 hypervisors to experiment with virtualization before making a commitment to purchase an expensive type-1 hypervisor.

Understanding Virtual Disks

Computers aren't the only things that are virtualized in a virtual environment. In addition to creating virtual computers, virtualization also creates virtual disk storage. Disk virtualization lets you combine a variety of physical disk storage devices to create pools of disk storage that you can then parcel out to your virtual machines as needed.

Virtualization of disk storage is nothing new. In fact, there are actually several layers of virtualization involved in an actual storage environment. At the lowest level are the actual physical disk drives. Physical disk drives are usually bundled together in arrays of individual drives. This bundling is a type of virtualization in that it creates the image of a single large disk drive that isn't really there. For example, four 2TB disk drives might be combined in an array to create a single 8TB disk drive.

Note that disk arrays are usually used to provide data protection through redundancy. This is commonly called RAID, which stands for *redundant array of inexpensive disks*.

One common form of RAID, called RAID-10, lets you create mirrored pairs of disk drives so that data is always written to both of the drives in a mirror pair. So, if one of the drives in a mirror pair fails, the other drive can carry the load. With RAID-10, the usable capacity of the complete array is equal to one-half of the total capacity of the drives in the array. For example, a RAID-10 array consisting of four 2TB drives contains two pairs of mirrored 2TB disk drives, for a total usable capacity of 4TB.

Another common form of RAID is RAID-5, in which disk drives are combined and one of the drives in the group is used for redundancy. Then, if any one of the drives in the array fails, the remaining drives can be used to re-create the data that was on the drive that failed. The total capacity of a RAID-5 array is equal to the sum of the capacities of the individual drives, minus one of the drives. For example, an array of four 2TB drives in a RAID-5 configuration has a total usable capacity of 6TB.

In a typical virtual environment, the host computers can be connected to disk storage in several distinct ways:

» **Local disk storage:** In local disk storage, disk drives are mounted directly on the host computer and are connected to the host computer via its internal disk drive controllers. For example, a host computer might include four 1TB disk drives mounted within the same chassis as the computer itself. These

four drives might be used to form a RAID-10 array with a usable capacity of 2TB.

The main drawbacks of local disk storage is that it's limited to the physical capacity of the host computers and is available only to the host computer that it's installed in.

» **Storage Area Network (SAN):** In a SAN, disk drives are contained in a separate device that is connected to the host via a high-speed controller. The difference between a SAN and local storage is that the SAN is a separate device. Its high-speed connection to the host is often just as fast as the internal connection of local disk storage, but the SAN includes a separate storage controller that is responsible for managing the disk drives.

A typical SAN can hold a dozen or more disk drives and can allow high-speed connections to more than one host. A SAN can often be expanded by adding one or more expansion chassis, which can contain a dozen or more disk drives each. Thus, a single SAN can manage hundreds of terabytes of disk data.

» **Network attached storage (NAS):** This type of storage is similar to a SAN, but instead of connecting to the hosts via a high-speed controller, a NAS connects to the host computers via standard Ethernet connections and TCP/IP. NAS is the least expensive of all forms of disk storage, but it's also the slowest.

Regardless of the way the storage is attached to the host, the hypervisor consolidates its storage and creates virtual pools of disk storage typically called *data stores*. For example, a hypervisor that has access to three 2TB RAID-5 disk arrays might consolidate them to create a single 6TB data store.

From this data store, you can create *volumes*, which are essentially virtual disk drives that can be allocated to a particular virtual machine. Then, when an operating system is installed in a virtual machine, the operating system can mount the virtual machine's volumes to create drives that the operating system can access.

For example, let's consider a virtual machine that runs Windows Server. If you were to connect to the virtual machine, log in, and use Windows Explorer to look at the disk storage that's available to the machine, you might see a C: drive with a capacity of 100GB. That C: drive is actually a 100GB volume that is created by the hypervisor and attached to the virtual machine. The 100GB volume, in turn, is allocated from a data store, which might be 4TB in size. The data store is created from disk storage contained in a SAN attached to the host, which might be made up of a RAID-10 array consisting of four 2TB physical disk drives.

So, you can see that there are at least four layers of virtualization required to make the raw storage available on the physical disk drives available to the guest operating system:

- » Physical disk drives are aggregated using RAID-10 to create a unified disk image that has built-in redundancy. RAID-10 is, in effect, the first layer of virtualization. This layer is managed entirely by the SAN.
- » The storage available on the SAN is abstracted by the hypervisor to create data stores. This is, effectively, a second level of virtualization.
- » Portions of a data store are used to create volumes that are then presented to virtual machines. Volumes represent a third layer of virtualization.
- » The guest operating system sees the volumes as if they're physical devices, which can be mounted and then formatted to create usable disk storage accessible to the user. This is the fourth layer of virtualization.

Although it may seem overly complicated, these layers of virtualization give you a lot of flexibility when it comes to storage management. New disk arrays can be added to a SAN, or a new NAS can be added to the network, and then new data stores can be created from them without disrupting existing data stores. Volumes can be moved from one data store to another without disrupting the virtual machines they're attached to. In fact, you can increase the size of a volume on the fly, and the virtual machine will immediately see the increased storage capacity of its disk drives, without even requiring so much as a reboot.

Understanding Network Virtualization

When you create one or more virtual machines on a host system, you need to provide a way for those virtual machines to communicate not only with each other but also with the other physical computers on your network. To enable such connections, you must create a *virtual network* within your virtualization environment. The virtual network connects the virtual machines to each other and to the physical network.

To create a virtual network, you must create a *virtual switch*, which connects the virtual machines to each other and to a physical network via the host computer's network interfaces. Like a physical switch, a virtual switch has ports. When you create a virtual switch, you connect the virtual switch to one or more of the host computer's network interfaces. These interfaces are then connected with network cable to physical switches, which effectively connects the virtual switch to the physical network.

Then, when you create virtual machines, you connect each virtual machine to a port on the virtual switch. When all the virtual machines are connected to the switch, the VMs can communicate with each other via the switch. And they can communicate with devices on the physical network via the connections through the host computer's network interfaces.

Looking at the Benefits of Virtualization

You might suspect that virtualization is inefficient because a real computer is inherently faster than a simulated computer. Although it's true that real computers are faster than simulated computers, virtualization technology has become so advanced that the performance penalty for running on a virtualized machine rather than a real machine is only a few percent.

The small amount of overhead imposed by virtualization is usually more than made up for by the simple fact that even the most heavily used servers spend most of their time twiddling their digital thumbs, waiting for something to do. In fact, many servers spend nearly *all* their time doing nothing. As computers get faster and faster, they spend even more of their time with nothing to do.

Virtualization is a great way to put all this unused processing power to good use.

Besides this basic efficiency benefit, virtualization has several compelling benefits:

- » **Hardware cost:** You typically can save a lot of money by reducing hardware costs when you use virtualization. Suppose that you replace ten servers that cost \$4,000 each with one host server. Granted, you'll probably spend more than \$4,000 on that server, because it needs to be maxed out with memory, processor cores, network interfaces, and so on. So you'll probably end up spending \$10,000 or \$15,000 for the host server. Also, you'll end up spending something like \$5,000 for the hypervisor software. But that's still a lot less than the \$40,000 you would have spent on ten separate computers at \$4,000 each.
- » **Energy costs:** Many organizations have found that going virtual has reduced their overall electricity consumption for server computers by 80 percent. This savings is a direct result of using less computer hardware to do more work. One host computer running ten virtual servers uses approximately one-tenth the energy that would be used if each of the ten servers ran on separate hardware.

» **Reduced downtime:** Virtual environments typically have less downtime than nonvirtual environments. For example, suppose you need to upgrade the BIOS on one of your server computers. With physical servers, this type of upgrade will ordinarily require that you shut down the operating system that runs on the server, upgrade the BIOS, and then restart the server. During the upgrade, the server will be unavailable.

In a virtual environment, you don't need to shut down the servers to upgrade the BIOS on the host computer that runs the server. Instead, all you do is move the servers that run on the host that needs the upgrade to another host. When the servers are moved (an operation that can be done without shutting them down), you can shut down the host and upgrade its BIOS. Then, after you restart the host, you can move the servers back to the host — again, without shutting down the servers.

» **Recoverability:** One of the biggest benefits of virtualization isn't the cost savings, but the ability to recover quickly from hardware failures. Suppose that your organization has ten servers, each running on separate hardware. If any one of those servers goes down due to a hardware failure — say, a bad motherboard — that server will remain down until you can fix the computer. On the other hand, if those ten servers are running as virtual machines on two different hosts, and one of the hosts fails, the virtual machines that were running on the failed host can be brought up on the other host in a matter of minutes.

Granted, the servers will run less efficiently on a single host than they would have on two hosts, but the point is that they'll all be running after only a short downtime.

In fact, with the most advanced hypervisors available, the transfer from a failing host to another host can be done automatically and instantaneously, so downtime is all but eliminated.

» **Disaster recovery:** Besides the benefit of recoverability when hardware failures occur, an even bigger benefit of virtualization comes into play in a true disaster-recovery situation. Suppose that your organization's server infrastructure consists of 20 separate servers. In the case of a devastating disaster, such as a fire in the server room that destroys all hardware, how long will it take you to get all 20 of those servers back up and running on new hardware? Quite possibly, the recovery time will be measured in weeks.

By contrast, virtual machines are actually nothing more than files that can be backed up onto tape. As a result, in a disaster-recovery situation, all you have to do is rebuild a single host computer and reinstall the hypervisor software. Then you can restore the virtual-machine backups from tape, restart the virtual machines, and get back up and running in a matter of days instead of weeks.

Introducing Hyper-V

Virtualization is a complex subject, and mastering the ins and outs of working with a full-fledged virtualization system like VMware Infrastructure is a topic that's beyond the scope of this book. You can dip your toes into the shallow end of the virtualization pond, however, by downloading and experimenting with Microsoft's free virtualization product, called Hyper-V, which comes with all server versions of Windows since Windows Server 2008 and all desktop versions of Windows since Windows 8.

The version of Hyper-V that comes with desktop versions of Windows is called Client Hyper-V. The nice thing about starting with Client Hyper-V is that it's similar to the enterprise-grade version of Hyper-V that is included with Windows Server. Much of what you learn about Hyper-V on desktop Windows applies to the server version as well.

Understanding the Hyper-V hypervisor

Although Hyper-V is built into all modern versions of Windows, Hyper-V is *not* a type-2 hypervisor that runs as an application within Windows. Instead, Hyper-V is a true type-1 hypervisor that runs directly on the host computer hardware. This is true even for the Hyper-V versions that are included with desktop versions of Windows.

In Hyper-V, each virtual machine runs within an isolated space called a *partition*. Each partition has access to its own processor, RAM, disk, network, and other virtual resources.

There are two types of partitions in Hyper-V: a *parent partition* and one or more *child partitions*. The parent partition is a special partition that hosts the Windows operating system that Hyper-V is associated with. Child partitions host additional virtual machines that you create as needed.

When you activate the Hyper-V feature, the hypervisor is installed and the existing Windows operating system is moved into a virtual machine that runs in the parent partition. Then, whenever you start the host computer, the hypervisor is loaded, the parent partition is created, and Windows is started in a virtual machine within the parent partition.

Although it may appear that the hypervisor is running within Windows, actually the reverse is true: Windows is running within the hypervisor.

In addition to the Windows operating system, the parent partition runs software that enables the management of virtual machines on the hypervisor. This includes creating new virtual machines, starting and stopping virtual machines, changing the resources allocated to existing virtual machines (for example, adding more processors, RAM, or disk storage), and moving virtual machines from one host to another.

Understanding virtual disks

Every Hyper-V virtual machine must have at least one virtual disk associated with it. A *virtual disk* is nothing more than a disk file that resides in the file system of the host operating system. The file has one of two file extensions, depending on which of two data formats you choose for the virtual disk:

- » .vhd: An older format that has a maximum virtual disk size of 2TB
- » .vhdx: A newer format that can support virtual disks up to 64TB

For either of these virtual disk formats, Hyper-V lets you create two different types of virtual disks:

- » **Fixed-size disk:** A virtual disk whose disk space is preallocated to the full size of the drive when you create the disk. For example, if you create a 100GB fixed-size disk using the .vhdx format, a .vhdx file of 100GB will be allocated to the drive. Even if the drive contains only 10GB of data, it will still consume 100GB of space on the host system's disk drive.
- » **Dynamically expanding disk:** A virtual disk that has a maximum disk space, but that actually consumes only the amount of disk space that is required to hold the data on the disk. For example, if you create a dynamically expanding disk with a maximum of 100GB but only put 10GB of data on it, the .vhdx file for the disk will occupy just 10GB of the host system's disk drive.



TIP

Don't be confused by the names *fixed size* and *dynamically expanding*. Both types of disk can be expanded later if you run out of space. The main difference is whether the maximum amount of disk space allowed for the drive is allocated when the drive is first created or as needed when data is added to the drive. Allocating the space when the drive is created results in better performance for the drive, because Hyper-V doesn't have to grab more disk space every time data is added to the drive. Both types of drives can be expanded later if necessary.

Enabling Hyper-V

Hyper-V is not automatically enabled when you install Windows; you must first enable this feature before you can use Hyper-V.

To enable Hyper-V on a server version of Windows, call up the Server Manager and open the Add Roles and Features Wizard. Then enable the Hyper-V role. When you complete the wizard, Hyper-V will install the Type-1 hypervisor and move the existing Windows Server operating system into the parent partition. You can then start building virtual machines.

To enable Hyper-V on a desktop version of Windows, follow these steps:

1. Open the Control Panel.

2. Choose Programs and Features.

The Programs and Features window appears.

3. Click Turn Windows Features On or Off.

The Windows Features dialog box appears, as shown in Figure 10-1.

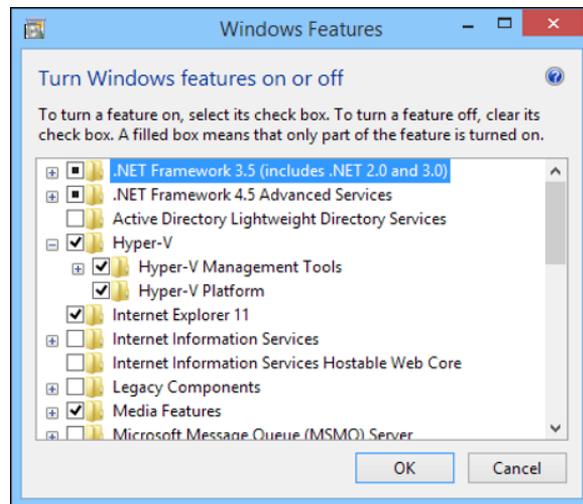


FIGURE 10-1: Enabling Hyper-V on a desktop version of Windows.

4. Select the Hyper-V feature and click OK.

The Client Hyper-V hypervisor is installed as an application on the existing desktop Windows operating system, and you can begin using Hyper-V.

5. When prompted, restart the computer.

The reboot is required to start the Hyper-V hypervisor. When your computer restarts, it's actually the Hyper-V hypervisor that starts, not Windows. The hypervisor then loads your desktop Windows into the parent partition.

Getting Familiar with Hyper-V

To manage Hyper-V, you use the Hyper-V Manager, shown in Figure 10-2. To start this program, click the Start button, type **Hyper-V**, and then choose Hyper-V Manager.

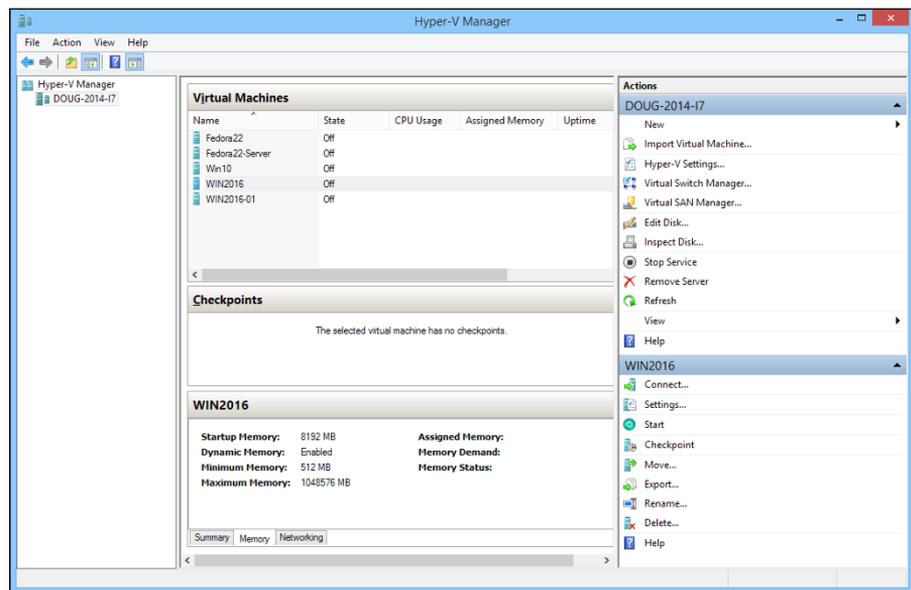


FIGURE 10-2: Hyper-V Manager.

The Hyper-V Manager window is divided into five panes:

- » **Navigation:** On the left side of the window is a navigation pane that lists the Hyper-V hosts, which Hyper-V calls *virtualization servers*. In Figure 10-2, just one host is listed: my Windows computer. In an enterprise environment where you have more than one host, each of the hosts will be listed in this pane.
- » **Virtual Machines:** This pane lists the virtual machines that are defined for the selected host. In Figure 10-2, you can see several of the Hyper-V virtual

machines that I created while I wrote this book — a couple of Linux machines, a Windows 10 machine, and a couple of Windows Server 2016 machines.

- » **Checkpoints:** In Hyper-V, a *checkpoint* is a recovery point for a virtual machine. You can create a checkpoint when you're going to make a modification to a virtual machine. Then, if something goes wrong, you can revert to the checkpoint.
- » **Virtual Machine Summary pane:** Below the Checkpoints pane is a pane that provides summary information for the virtual machine selected in the Virtual Machines pane. In Figure 10-2, you can see the summary information for one of the Windows Server 2016 machines. This pane has three tabs: Summary, Memory, and Networking. In the figure, the Memory tab is selected so you can see the memory that has been allocated to the machine.
- » **Actions:** The Actions tab contains buttons you can click to initiate actions for the selected host (DOUG-2014-I7) and the selected machine (WIN2016).

Creating a Virtual Switch

Before you start creating virtual machines in Hyper-V, you should create a virtual switch so that your virtual machines can communicate with each other and with the outside world. To do that, you use the Virtual Switch Manager. Here are the steps:

- 1. In Hyper-V Manager, click Virtual Switch Manager.**

This brings up the Virtual Switch Manager window, as shown in Figure 10-3.

- 2. Select the type of virtual switch you want to create.**

Hyper-V lets you create three types of switches:

- *External:* A virtual switch that binds to a physical network adapter, which allows virtual machines to communicate with each other, as well as with other computers on your physical network. This is usually the type of switch you should create.
- *Internal:* A virtual switch that does not bind with a physical network adapter. This type of switch lets the virtual machines on this computer communicate with each other and with the host computer, but not with other computers on your physical network.
- *Private:* A virtual switch that lets virtual machines communicate with each other but not with the host computer or with any computers on your physical network.

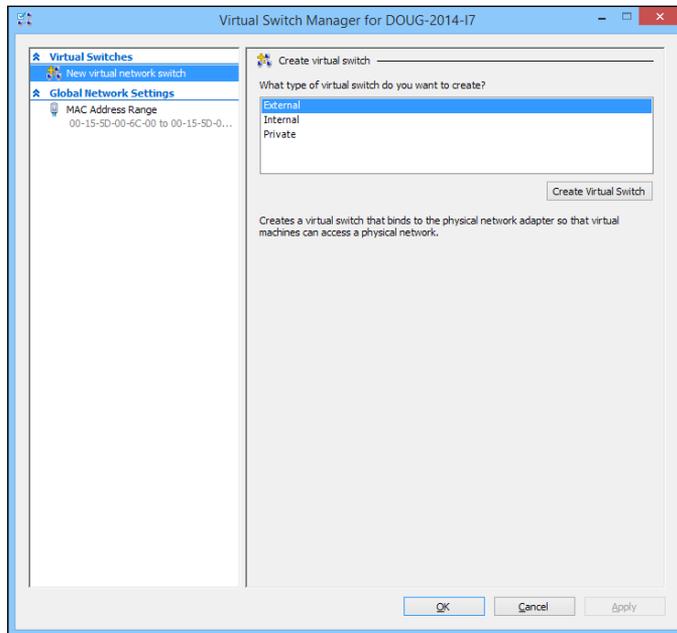


FIGURE 10-3:
The Virtual Switch Manager window.

3. Click Create Virtual Switch.

The settings for the new virtual switch appear, as shown in Figure 10-4.

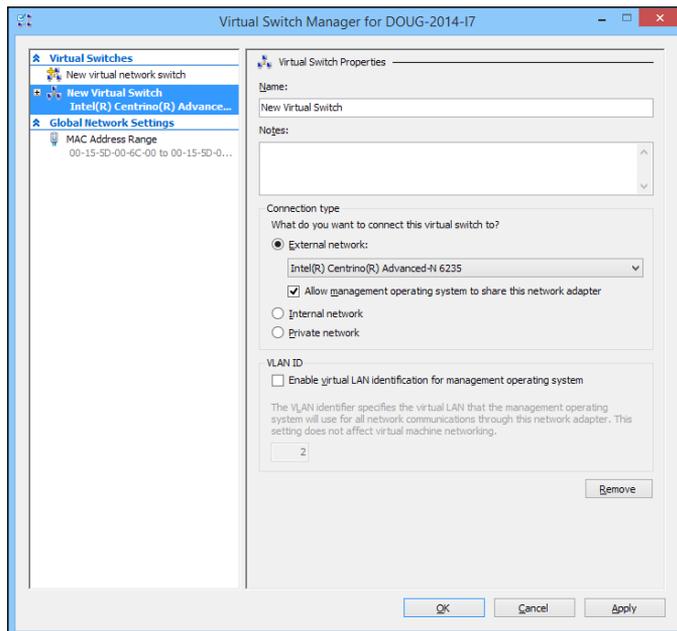


FIGURE 10-4:
Creating a new virtual switch.

- 4. Type a name for the new virtual switch in the Name field.**
Use any name you want.
- 5. Select the physical network adapter you want to bind the virtual switch to.**
If your computer has more than one network adapter, select the one you want to use. Binding the virtual switch to a physical network adapter allows the virtual machines to communicate not only with each other but also with other computers connected via the adapter you select.
- 6. If your network has multiple VLANs, click the Enable Virtual LAN Identification for Management Operating System check box and enter the VLAN ID for the VLAN you want this switch to connect to.**
If your network doesn't have multiple VLANs, you can skip this step.
- 7. Click OK.**
The virtual switch is created. Your Hyper-V environment now has a virtual network in place, so you can start creating virtual machines.

Creating a Virtual Disk

Before you create a virtual machine, it's best to first create a virtual disk for the machine to use. Note that you can create a virtual disk at the same time that you create a virtual machine. However, creating the virtual disk first gives you more flexibility. So, I recommend you create virtual disks and virtual machines separately. Here are the steps to create a virtual disk:

- 1. In Hyper-V Manager, click New and then choose Hard Disk.**
This brings up the New Virtual Hard Disk Wizard, as shown in Figure 10-5.
- 2. Click Next.**
You're asked which disk format to use, as shown in Figure 10-6. I recommend you always use the VHDX format, which can support drives larger than 2TB.
- 3. Select VHDX, and then click Next.**
When you click Next, the Choose Disk Type option page is displayed, as shown in Figure 10-7.
- 4. Select the disk type you want to use.**
The options are Fixed Size, Dynamically Expanding, and Differencing. Choose Fixed Size if you're concerned about the performance of the disk; otherwise, choose Dynamically Expanding.

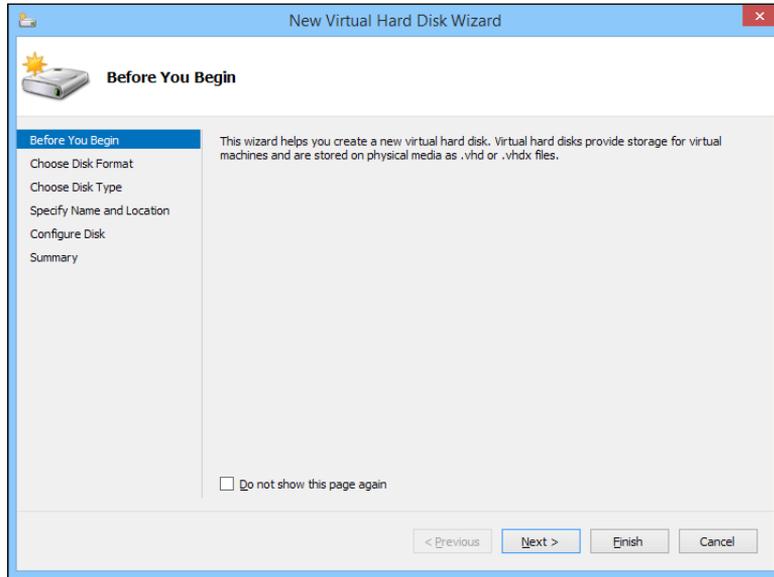


FIGURE 10-5:
The New Virtual
Hard Disk Wizard.

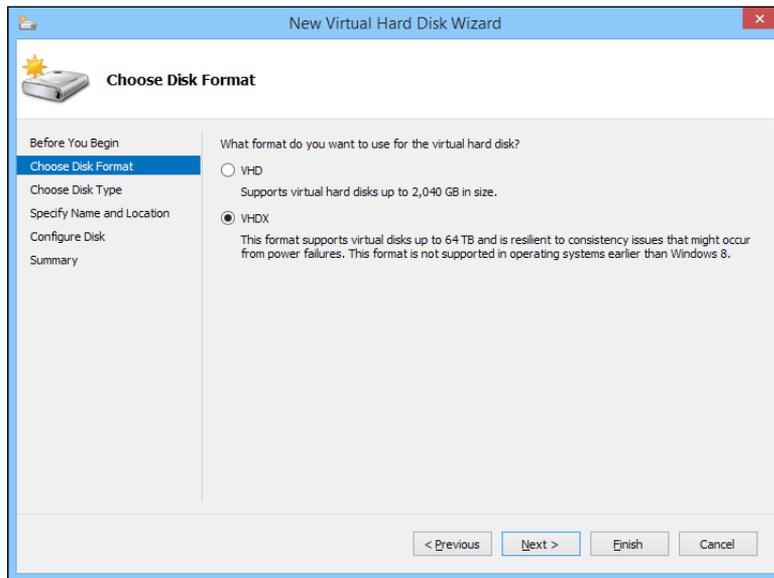


FIGURE 10-6:
Choose your
disk format.

5. Click Next.

The Specify Name and Location page, shown in Figure 10-8, appears.

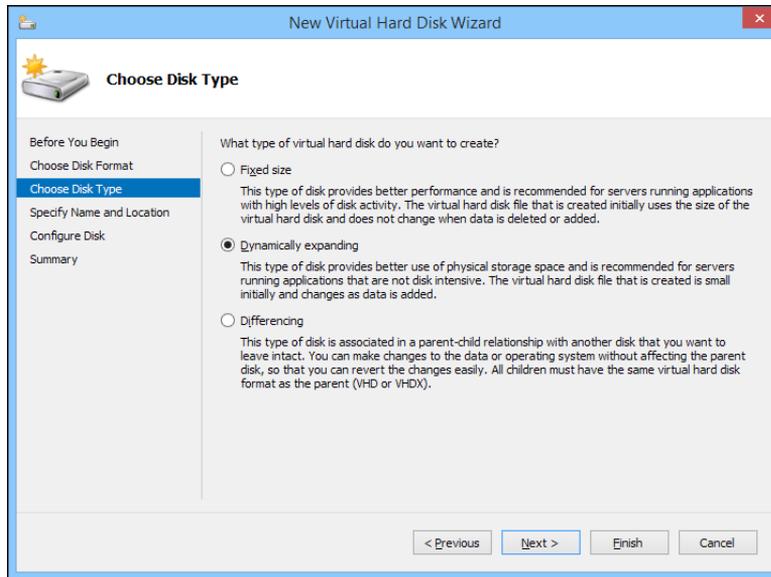


FIGURE 10-7: Choose your disk type.

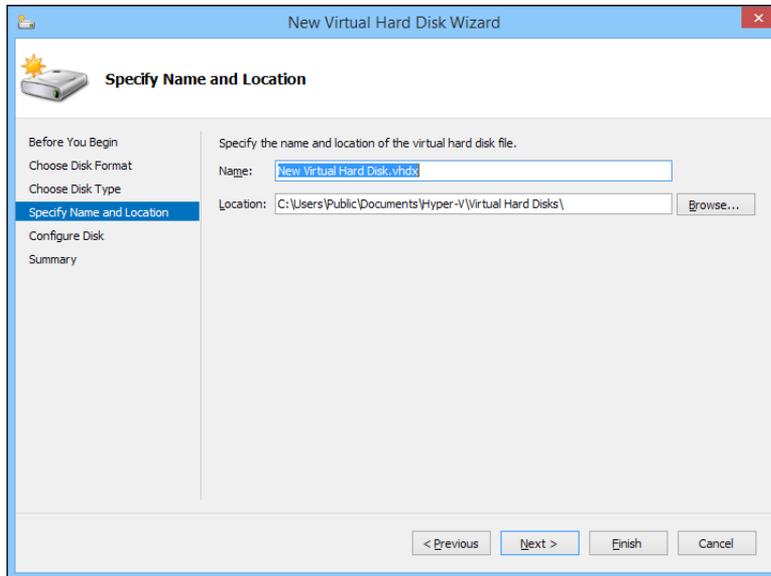


FIGURE 10-8: Specify the name and location of the disk.

6. Specify the name and location of the new disk.

Type any name you want for the virtual disk drive. Then click the Browse button to browse to the disk location where you want Hyper-V to create the .vhdx file.



TIP

Make sure you choose a location that has enough disk space to create the .vhdx file. If you're creating a dynamically expanding disk, you should ensure that the location has enough space to accommodate the drive as it grows.

7. Click Next.

The Configure Disk page appears, as shown in Figure 10-9.

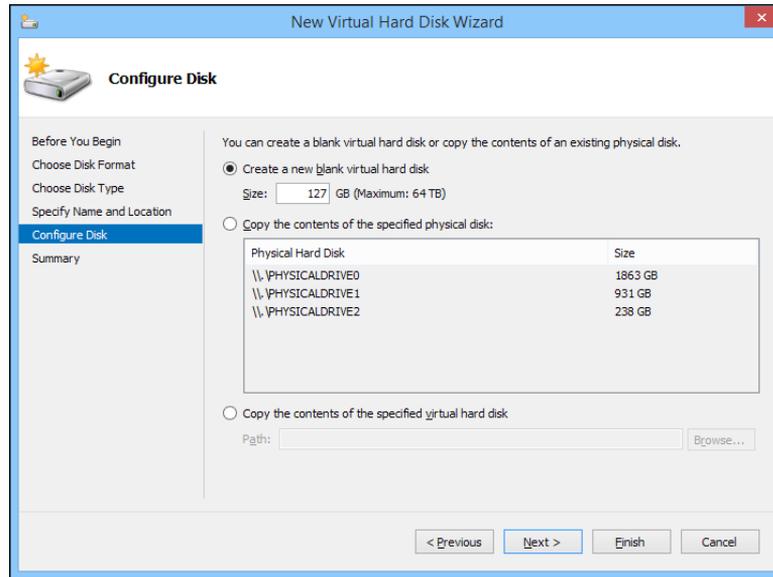


FIGURE 10-9: Specify the size of the disk.

8. Specify the maximum size for the disk drive.



TIP

This page also allows you to copy data either from an existing physical disk drive or from an existing virtual disk drive. Copying data from an existing physical drive is a quick way to convert a physical computer to a virtual computer; just copy the physical disk to a virtual disk, and then use the new virtual disk as the basis for a new virtual machine.

9. Click Next.

A confirmation screen appears, summarizing the options you've selected for your new disk.

10. Click Finish.

The new disk is created. Note that if you selected Fixed Disk as the disk type, creating the disk can take a while because the entire amount of disk storage you specified is allocated to the disk. Be patient.

You're done! You've now created a virtual disk that can be used as the basis for a new virtual machine.

Creating a Virtual Machine

After you've created a virtual disk, creating a virtual machine to use it is a straightforward affair. Follow these steps:

- 1. From the Hyper-V Manager, choose New and then choose Virtual Machine.**

This brings up the New Virtual Machine Wizard, as shown in Figure 10-10.

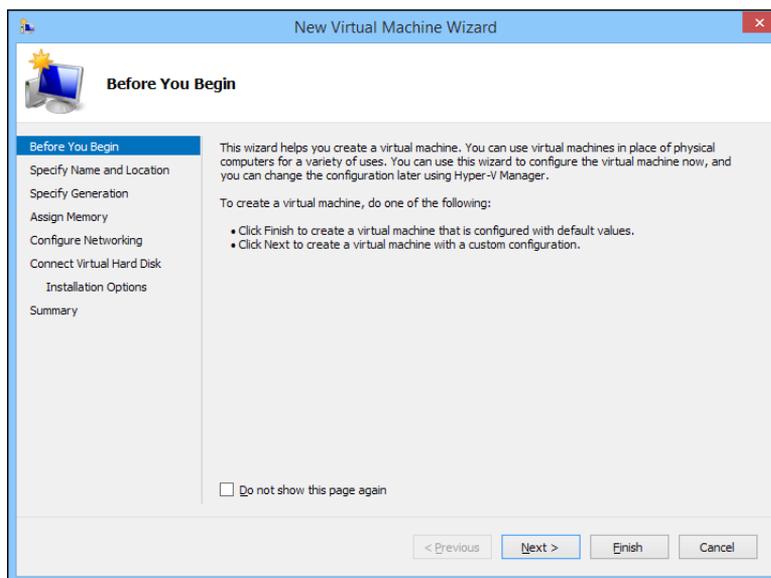


FIGURE 10-10:
Say hello to the
New Virtual
Machine Wizard.

- 2. Click Next.**

The Specify Name and Location page appears, as shown in Figure 10-11.

- 3. Enter the name you want to use for your virtual machine.**

You can choose any name you want here.

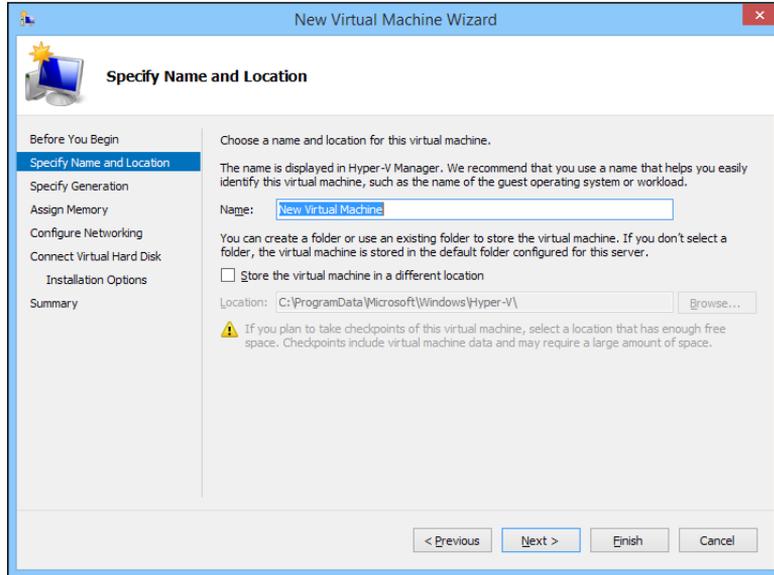


FIGURE 10-11: Specify the name and location of the virtual machine.

4. **Specify the location of the virtual machine's configuration file.**

Every virtual machine has an XML file associated with it that defines the configuration of the virtual machine. You can allow this file to be stored in the default location, or you can override the default and specify a custom location.

5. **Click Next.**

The Specify Generation page appears, as shown in Figure 10-12.

6. **Specify the generation you want to use for the new virtual machine.**

In most cases, you should opt for Generation 2, which uses newer technology than Generation 1 machines. Use Generation 1 only if the guest operating system will be earlier than Windows Server 2012 or Windows 8.

7. **Click Next.**

The Assign Memory page appears, as shown in Figure 10-13.

8. **Indicate the amount of RAM you want to allocate for the new machine.**

The default is 512MB, but you'll almost certainly want to increase that.

I also recommend that you select the Use Dynamic Memory for This Virtual Machine check box, which improves memory performance.

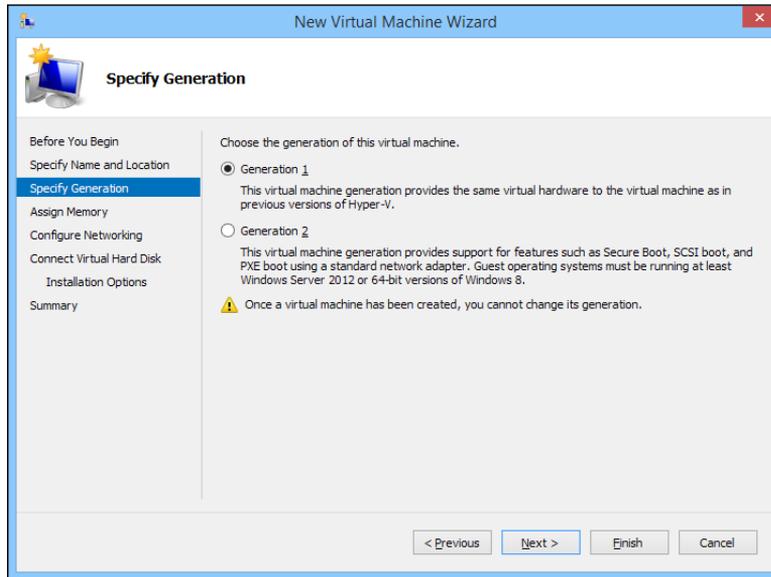


FIGURE 10-12: Specify the generation of the new virtual machine.

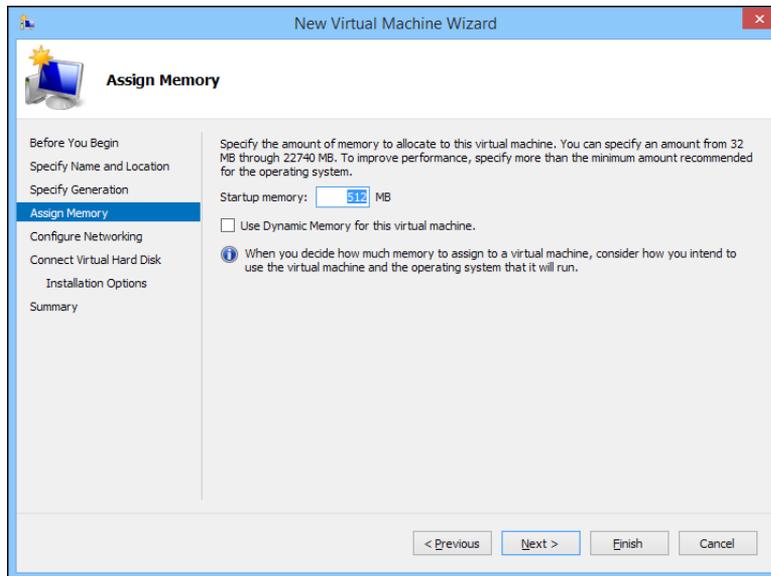


FIGURE 10-13: Specify the memory for the new virtual machine.

9. Click Next.

The Configure Networking page appears, as shown in Figure 10-14.

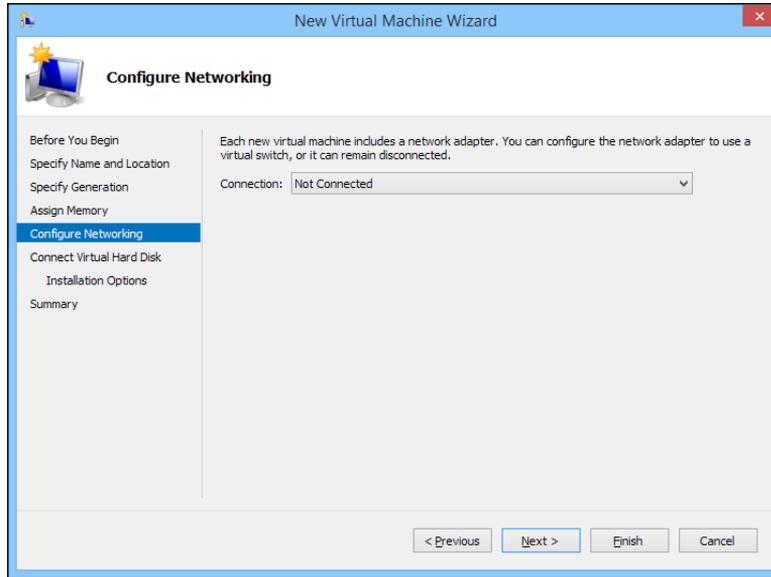


FIGURE 10-14: Configure the networking for the new virtual machine.

10. Select the virtual switch you want to use for the virtual machine.

This is the point where you realize why you needed to create a virtual switch before you started creating virtual machines. Use the Connection drop-down list to select the virtual switch you want to connect to this virtual machine.

11. Click Next.

The Connect Virtual Hard Disk page appears, as shown in Figure 10-15.

12. Assuming you've already created a virtual disk for the virtual machine, choose the Use an Existing Virtual Hard Disk option, click Browse, and locate and select the virtual disk.

If you haven't already created a virtual disk, you can use the Create a Virtual Hard Disk option and create one now.

13. Click Next.

A summary page is displayed indicating the selections you've made.

14. Click Finish.

The virtual machine is created.

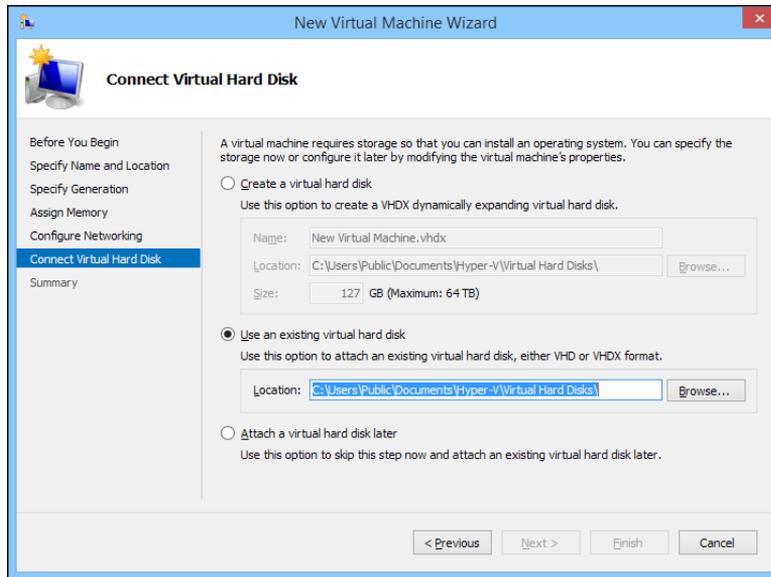


FIGURE 10-15:
Connecting a
virtual disk.

Installing an Operating System

After you've created a virtual machine, the next step is to configure it to install an operating system. First, you need to get the installation media in the form of an .iso file (an .iso file is a disk image of a CD or DVD drive). After you have the .iso file in place, follow these steps:

- 1. From the Hyper-V Manager, choose the new virtual machine and click Settings.**
The Settings dialog box appears, as shown in Figure 10-16.
- 2. Click SCSI Controller in the Hardware list. Then select DVD Drive, and click Add.**
The configuration page shown in Figure 10-17 appears.
- 3. Click the Image File option, click Browse, and select the .iso file that contains the operating system's installation program.**
- 4. Click OK.**

You're returned to the Hyper-V Manager screen.

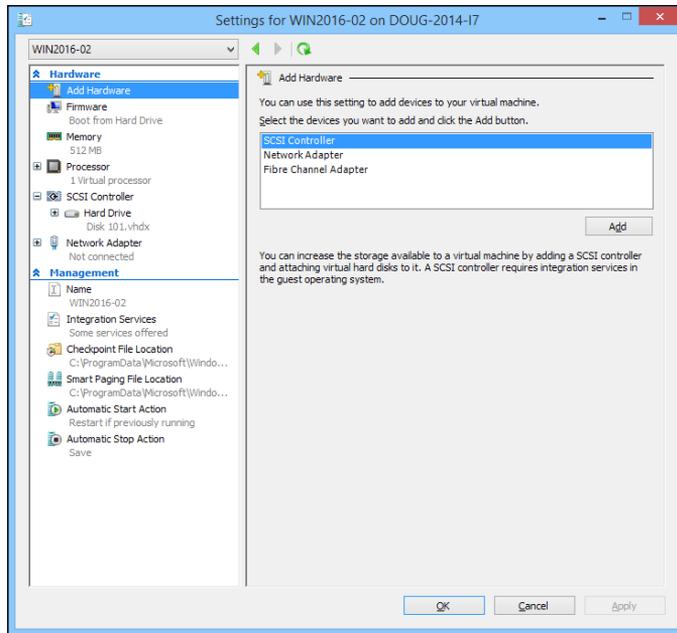


FIGURE 10-16:
Editing the settings for a virtual machine.

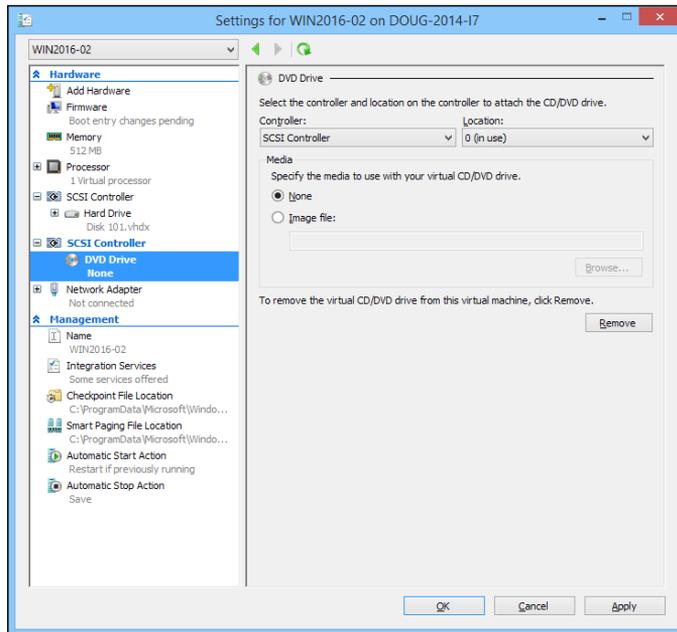


FIGURE 10-17:
Configuring a DVD drive.

5. With the new virtual machine still selected, click Connect.

A console window opens, showing that the virtual machine is currently turned off (see Figure 10-18).

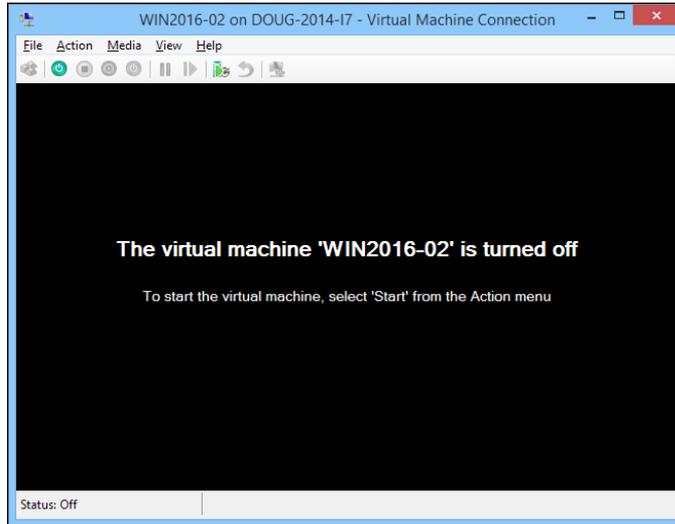


FIGURE 10-18:
Viewing a virtual machine through a console.

6. Click Connect.

7. Click Start.

The virtual machine powers up.

8. When prompted to press a key to boot from the CD or DVD, press any key.

The operating system's installation program starts.

9. Follow the instructions of the installation program to install the operating system.

That's all there is to it. You've now successfully created a Hyper-V virtual machine!



Working with Servers

IN THIS PART . . .

Install and configure Windows Server 2016.

Understand and use Active Director.

Create and manage user accounts.

Manage file servers and email servers.

Set up a company intranet.

IN THIS CHAPTER

Getting a handle on server operating system features

Discovering Windows Server and other server options

Thinking about the different ways to install a server operating system

Getting ready for the installation

Installing a server operating system

Configuring your server roles

Chapter 11

Setting Up a Server

One of the basic choices that you must make before you proceed any further in building your network is to decide which server operating system to use as the foundation for your network. This chapter begins with a description of several important features found in all server operating systems. Next, it provides an overview of the advantages and disadvantages of the most popular server operating systems, including various versions of Windows Server, Linux, and Apple's OS X Server.

Of course, your work doesn't end with the selection of a server operating system. You must then install and configure the operating system to get it working. This chapter provides an overview of what's involved with installing and configuring Microsoft's latest and greatest server operating system, Windows Server 2016.

Server Operating System Features

All server operating systems must provide certain core functions, such as connecting to other computers on the network, sharing files and other resources, and providing for security. In the following sections, I describe some core server operating system features in general terms.

Network support

It goes without saying that a server operating system must provide networking capabilities in order for it to function on a network. If your client computers can't connect to your servers, your network will be useless. For this reason, it's a good idea to make sure your server computers are equipped with more than one network interface. That way, if one of the interfaces fails, the other can pick up the slack and keep your server connected to your network.

In addition to basic network connectivity, one of your servers will typically be responsible for providing some essential software services that are required to keep a network operating in an efficient manner. One of these is called Dynamic Host Configuration Protocol (DHCP); it's the service that recognizes computers and other devices that want to join the network, providing each with a unique address so that all the devices on the network can identify one another. All modern server operating systems are able to provide these services.

File-sharing services

One of the most important functions of a server operating system is to share resources with other network users. The most common resource that's shared is the server's *file system* — organized disk space that a server must be able to share (in whole or in part) with other users. In effect, those users can treat the server's disk space as an extension of their own computers' disk space.

The server operating system allows the system administrator to determine which portions of the server's file system to share.



TIP

Although an entire hard drive can be shared, it isn't commonly done. Instead, individual folders are shared. The administrator can control which users are allowed to access each shared folder.

Because file sharing is the reason why many network servers exist, server operating systems have more sophisticated disk management features than are found in desktop operating systems. For example, most server operating systems can manage two or more hard drives as though they were a single drive. In addition, most can create a *mirror* — an automatic backup copy of a drive — on a second drive.

Multitasking

Only one user at a time uses a desktop computer; however, multiple users simultaneously use server computers. As a result, a server operating system must provide support for multiple users who access the server remotely via the network.

At the heart of multiuser support is *multitasking* — a technique that slices processing time microthin and juggles the pieces lightning fast among running programs. It's how an OS can execute more than one program (a task or a process) at a time. Multitasking operating systems are like the guy who used to spin plates balanced on sticks on the old *Ed Sullivan Show*: running from plate to plate, trying to keep them all spinning so that they don't fall off the sticks. To make it challenging, he'd do it blindfolded or riding on a unicycle. Substitute programs for the plates and file management for the unicycle, and there you are.

Although multitasking creates the appearance that two or more programs execute on the computer at the same time, in reality, a computer with a single processor can execute only one program at a time. The OS switches the CPU from one program to another to create the appearance that several programs execute simultaneously, but at any given moment, only one program processes commands. The others are patiently waiting their turns. (However, if the computer has more than one CPU core, the CPU cores *can* execute programs simultaneously — but that's another kettle of fish.)

Directory services

Directories are everywhere — and were, even in the days when they were all hard copy. When you needed to make a phone call, you looked up the number in a phone directory. When you needed to find the address of a client, you looked her up in your Rolodex. And then there were the nonbook versions: When you needed to find the Sam Goody store at a shopping mall (for example), you looked for the mall directory — usually, a lighted sign showing what was where.

Networks have directories, too, providing information about the resources that are available on the network: users, computers, printers, shared folders, and files. Directories are essential parts of any server operating system.

The most popular modern directory service is called *Active Directory*. Active Directory is a standard component of all Windows operating systems, and because it's so popular, most other operating systems support it as well. Active Directory is a database that organizes information about a network and all its computers and users. It's simple enough to use for networks with just a few computers and users, but powerful enough to work with large networks containing tens of thousands of computers and users. Figure 11-1 shows the Active Directory Users and Computers tool, which manages Active Directory user and computer accounts on Windows Server 2016.

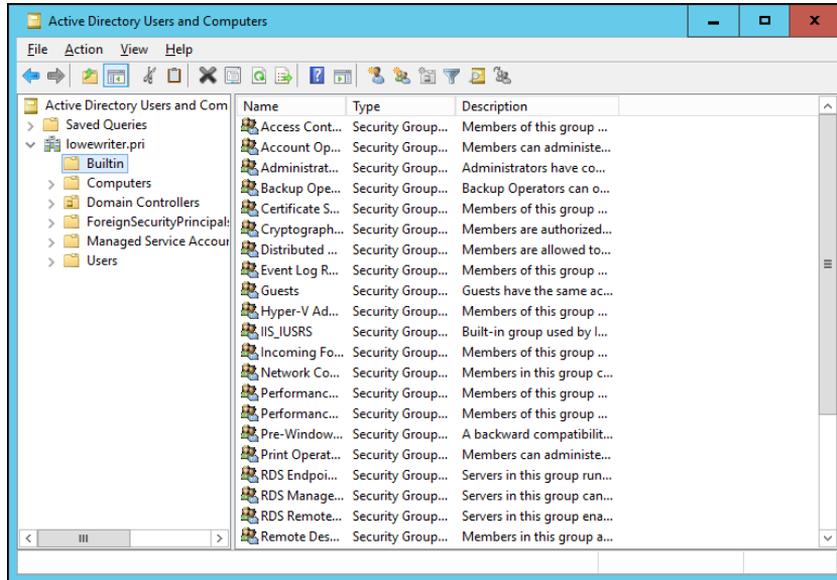


FIGURE 11-1:
Active Directory
Users and
Computers.

Security services

All server operating systems must provide some measure of security to protect the network from unauthorized access. Hacking seems to be the national pastime these days. With most computer networks connected to the Internet, anyone anywhere in the world can — and probably will — try to break into your network.

The most basic type of security is handled through *user accounts*, which grant individual users the right to access the network resources and govern which resources the user can access. User accounts are secured by passwords; therefore, good password policy is a cornerstone of any security system. Most server operating systems give you some standard tools for maintaining network security:

- » **Establish password policies.** For example, you can mandate that passwords have a minimum length and include a mix of letters and numerals.
- » **Set passwords to expire after a certain number of days.** Network users must change their passwords frequently.
- » **Encrypt network data.** A data-encryption capability scrambles data before it's sent over the network or saved on disk, making unauthorized use a lot more difficult.



TIP

Good encryption is the key to setting up a virtual private network (VPN), which enables network users to securely access a network from a remote location by using an Internet connection.

» **Manage digital certificates.** Digital certificates are used to ensure that users are who they say they are and files are what they claim to be.

The overwhelming majority of business networks rely on server versions of Windows, known as Windows Server. Microsoft periodically releases updated versions of Windows Server, so Windows Server is frequently improved, and older versions are occasionally rendered obsolete. Currently, the most commonly used versions are Windows Server 2008, Windows Server 2012, Windows Server 2012 Release 2, and the latest-and-greatest version, known as Windows Server 2016.

But Windows Server is not the only server operating system at your disposal. Many servers — especially those whose primary responsibility is to host websites — use Linux instead of Windows Server. Apple also makes an excellent server operating system, known as OS X Server.

The Many Ways to Install a Network Operating System

Regardless of which server operating system you choose to use for your network servers, you can use any of several common ways to actually install the server operating system software on the server computer. The following sections describe these alternatives.

Full install versus upgrade

One of the basic server operating system installation choices is whether you want to perform a full installation or an upgrade installation. In some cases, you may be better off performing a full installation even if you're installing the server operating system on a computer with an earlier version of the server operating system installed:

» **Full:** If you're installing the server operating system on a new server, you'll be performing a full installation that installs the OS and configures it with default settings.

» **Upgrade:** If you're installing the server operating system on a server computer that already has a server OS installed, you can perform an upgrade installation to replace the existing OS with the new one, yet retain as many of the settings from the existing OS as possible.

Note that you can't upgrade a client version of Windows to a server version. Instead, you must perform a full installation, which deletes the existing Windows OS, or a multiboot installation, which leaves the existing client Windows intact. Either way, however, you can preserve existing data on the Windows computer when you install the server version.

» **Multiboot:** You can also perform a full installation on a computer that already has an OS installed. In that case, you have the option of deleting the existing OS or performing a multiboot installation that installs the new server OS alongside the existing OS. Then, when you restart the computer, you can choose which OS you want to run.

Although multiboot installation may sound like a good idea, it's fraught with peril. Avoid multiboot unless you have a specific reason to use it.



WARNING

Installing over the network

Typically, you install the server operating system directly from the CD-ROM distribution discs on the server's CD-ROM drive. However, you can also install the OS from a shared drive located on another computer — provided that the server computer already has access to the network. You can either use a shared CD-ROM drive or copy the entire contents of the distribution CD-ROM disc onto a shared hard drive.



TIP

If you're going to install the server operating system onto more than one server, save time by first copying the distribution CD onto a shared hard drive. That's because even the fastest CD-ROM drives are slower than the network. Even with a slow 100 Mbps network, access to hard drive data over the network is much faster than access to a local CD-ROM drive.

Gathering Your Stuff

Before you install a server operating system, gather everything you need so you don't have to look for something in the middle of the setup. The following sections describe the items you're most likely to need.

A capable server computer

Obviously, you have to have a server computer on which to install the server operating system. Each server operating system has a list of minimum hardware requirements supported by the OS. For example, Table 11-1 summarizes the minimum requirements for Windows Server 2012.

TABLE 11-1

Minimum Hardware Requirements for Windows Server 2016

Item	Windows Server 2016
CPU	1.4 GHz
RAM	512MB
Free disk space	32GB

My suggestion is that you take these minimums with a grain of salt. Windows Server 2016 will crawl like a snail with 512MB of RAM; I wouldn't bother with less than 8GB, and 16GB is a more appropriate minimum for most purposes.

You should also check your server hardware against the list of compatible hardware published by the maker of your server operating system. For example, Microsoft publishes a list of hardware that it has tested and certified as compatible with Windows servers. This list is called the hardware compatibility list (HCL). You can check the HCL for your specific server by browsing to <https://sysdev.microsoft.com/en-us/hardware/lp1>.

You can also test your computer's compatibility by running the Check System Compatibility option from the Windows distribution CD-ROM.

The server OS

You also need a server OS to install. You need the distribution CDs or DVDs or access to copies of them over the network. In addition to the discs, you should have the following:

- » **The product key:** The installation program will ask you to prove that you have a legal copy of the software. If you have the actual DVDs, the product key should be on a sticker attached to the case.
- » **Your license type:** You can purchase Windows Server on a per-server or a per-user/per-device basis. You need to know which plan you have when you install the server operating system.



TIP

Check the DVD distribution disc for product documentation and additional last-minute information. For example, Windows servers have a `\docs` folder that contains several files that have useful setup information.

Other software

In most cases, the installation program should be able to automatically configure your server's hardware devices and install appropriate drivers. Just in case, though, you should dig out the driver disks/discs that came with your devices, such as network interface cards, SCSI devices, DVD drives, printers, scanners, and so on.

A working Internet connection

Online connectivity isn't an absolute requirement, but the installation will go much smoother if you have a working Internet connection before you start. The installation process may use this Internet connection for several things:

- » **Downloading late-breaking updates or fixes to the OS:** This can eliminate the need to install a Service Pack after you finish installing the server operating system.
- » **Locating drivers for nonstandard devices:** This can be a big plus if you can't find the driver disk for an obscure SCSI card.
- » **Activating the product after you complete the installation (for Windows Server)**

A good book

You'll spend lots of time watching progress bars during installation, so you may as well have something to do while you wait. I recommend Ron Chernow's *Hamilton*, the 800-page tome that was the inspiration for the popular Broadway musical that no one can get tickets for.

Making Informed Decisions

When you install a server operating system, you have to make some decisions about how you want the OS and its servers configured. Most of these decisions aren't cast in stone, so don't worry if you're not 100 percent sure how you want everything configured. You can always go back and reconfigure things. However, you'll save yourself time if you make the right decisions up front rather than just guess when the setup program starts asking you questions.

The following list details most of the decisions that you'll need to make. (This list is for Windows Server 2016 installations. For other server operating systems, the decisions may vary slightly.)

- » **The existing operating system:** If you want to retain the existing operating system, the setup program can perform a multiboot setup, which allows you to choose which operating system to boot to each time you start the computer. This is rarely a good idea for server computers, so I recommend that you elect to delete the existing operating system.
- » **Partition structure:** Most of the time, you'll want to treat the entire server disk as a single partition. However, if you want to divide the disk into two or more partitions, you should do so during setup. (Unlike most of the other setup decisions, this one is hard to change later.)
- » **File system:** Two choices are available for the file system used to format the server's disk: NT File System (NTFS) and Resilient File System (ReFS). NTFS has been around since 1993. ReFS is a relatively new file system (introduced with Windows Server 2012) that offers several important improvements over NTFS. However, because it's still relatively new, most network administrators are reluctant to use it. So NTFS remains the file system of choice.
- » **Computer name:** During setup, you'll be asked to provide the computer name used to identify the server on the network. If your network has only a few servers, you can just pick a name such as Server01 or MyServer. If your network has more than a few servers, you'll want to establish a naming convention you can follow for naming your servers.
- » **Administrator password:** Okay, this one is tough. You don't want to pick something obvious, like Password, Administrator, or your last name. On the other hand, you don't want to type in something random that you'll later forget because you'll find yourself in a big pickle if you forget the administrator password. I suggest that you make up a complex password with uppercase and lowercase letters, some numerals, and a special symbol or two; *then write it down and keep it in a secure location* where you know it won't get lost.
- » **TCP/IP configuration:** You'll need to know what IP address to use for the server. Even if your network has a DHCP server to dynamically assign IP addresses to clients, most servers use static IP addresses.
- » **Domain name:** You'll need to supply the name of the Active Directory domain to which the server will belong.

Final Preparations

Before you begin the actual installation, take a few more steps:

- » **Tidy up.** Clean up the server's disk by uninstalling any software that you don't need and removing any old data that is no longer needed. This cleanup is especially important if you're converting a computer that's been in use as a client computer to a server. You probably don't need Microsoft Office or a bunch of games on the computer after it becomes a server.
- » **Backup.** Do a complete backup of the computer. Operating system setup programs are almost flawless, so the chances of losing data during installation are minimal, but you still face the chance that something may go wrong.
- » **Disconnect serial and USB connection.** If the computer is connected to an uninterruptible power supply (UPS) that has a serial or USB connection to the computer, unplug the serial or USB connection. In some cases, this control connection can confuse the OS's setup program when it tries to determine which devices are attached to the computer.
- » **Chill out.** Light some votive candles, take two acetaminophen, and put on a pot of coffee.

Installing a Server Operating System

The following sections present an overview of a typical installation of Windows Server 2016. Although the details vary, the overall installation process for other server operating systems is similar.

In most cases, the best way to install Windows Server 2016 is to perform a new install directly from the DVD installation media. Although upgrade installs are possible, your server will be more stable if you perform a new install. (For this reason, most network administrators avoid upgrading to newer Windows Server versions until it's time to replace the server hardware.)

To begin the installation, insert the DVD distribution media in the server's DVD drive and then restart the server. This causes the server to boot directly from the distribution media, which initiates the setup program.

As the setup program proceeds, it leads you through two distinct installation phases: Collecting Information and Installing Windows. The following sections describe these installation phases in greater detail.

Phase 1: Collecting Information

In the first installation phase, the setup program asks for the preliminary information that it needs to begin the installation. A setup wizard prompts you for the following information:

- » **Language:** Select your language, time zone, and keyboard type.
- » **Product key:** Enter the 25-character product key that came with the installation media. If setup says that you entered an invalid product key, double-check it carefully. You probably just made a typo.
- » **Operating system type:** Select one of two versions of Windows Server to install:
 - *Server Core:* A streamlined version that does not have a graphical user interface (GUI). You must manage this version of the server remotely using a command-line interface known as PowerShell.
 - *Desktop Experience:* The full version that includes a GUI for simpler management.
- » **License agreement:** The official license agreement is displayed. You have to agree to its terms to proceed.
- » **Install type:** Choose an Upgrade or Clean Install type.
- » **Installation partition:** Choose the disk partition in which you want to install Windows.
- » **Administrator password:** Enter the administrator password.

Phase 2: Installing Windows

In this phase, Windows setup begins the actual process of installing Windows Server. The following steps are performed in sequence:

1. **Copying Files:** Compressed versions of the installation files are copied to the server computer.
2. **Expanding Files:** The compressed installation files are expanded.
3. **Installing Features:** Windows server features are installed.
4. **Installing Updates:** The setup program checks the Microsoft website and downloads any critical updates to Windows Server.
5. **Completing Installation:** When the updates are installed, the setup program reboots so it can complete the installation.

Configuring Your Server

After you install Windows Server 2016, the computer automatically reboots, and you're presented with the Server Manager program as shown in Figure 11-2.

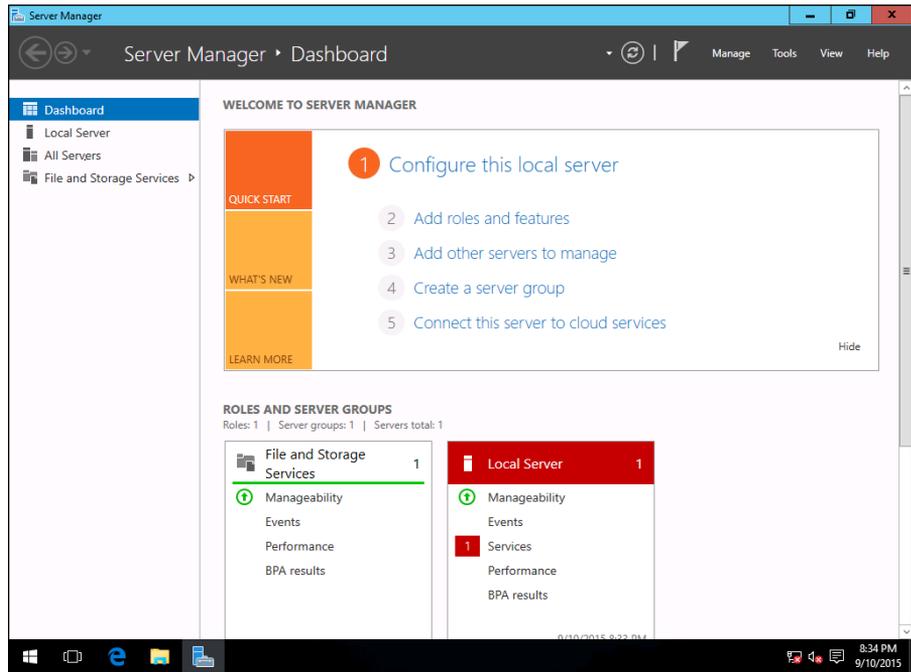


FIGURE 11-2:
The Windows
Server 2016
Server Manager.

From Server Manager, you can perform a number of tasks that are necessary to configure the server for production use. Specifically, you can configure the server roles: the networking features that the server will provide, such as file server, web server, DHCP server, and DNS server.

Chapter 12

Managing Windows User Accounts

Every user who accesses a network must have a user account. User accounts allow you — as network administrator — to control who can access the network and who can't. In addition, user accounts let you specify what network resources each user can use. Without user accounts, all your resources would be open to anyone who casually dropped by your network.

Understanding Windows User Accounts

User accounts are among the basic tools for managing a Windows server. As a network administrator, you'll spend a large percentage of your time dealing with user accounts — creating new ones, deleting expired ones, resetting passwords for forgetful users, granting new access rights, and so on. Before I get into the specific procedures of creating and managing user accounts, this section presents an overview of user accounts and how they work.

Local accounts versus domain accounts

A *local account* is a user account stored on a particular computer, applicable to that computer only. Typically, each computer on your network has a local account for each person who uses that computer.

By contrast, a *domain account* is a user account that's stored by Active Directory (AD) and can be accessed from any computer that's a part of the domain. Domain accounts are centrally managed. This chapter deals primarily with setting up and maintaining domain accounts.

User account properties

Every user account has several important account properties that specify the characteristics of the account. The three most important account properties are

- » **Username:** A unique name that identifies the account. The user must enter the username when logging on to the network. The username is public information. In other words, other network users can (and often should) find out your username.
- » **Password:** A secret word that must be entered to gain access to the account. You can set up Windows so that it enforces password policies, such as the minimum length of the password, whether the password must contain a mixture of letters and numerals, and how long the password remains current before the user must change it.
- » **Group membership:** The group(s) to which the user account belongs. Group memberships are the key to granting access rights to users so that they can access various network resources (such as file shares or printers) or perform certain network tasks (such as creating new user accounts or backing up the server).

Many other account properties record information about the user, such as the user's contact information, whether the user is allowed to access the system only at certain times or from certain computers, and so on.

Creating a New User

To create a new domain user account in Windows Server 2016, follow these steps:

1. **Choose Start ⇨ Administrative Tools ⇨ Active Directory Users and Computers.**

This command fires up the Active Directory Users and Computers management console, as shown in Figure 12-1.

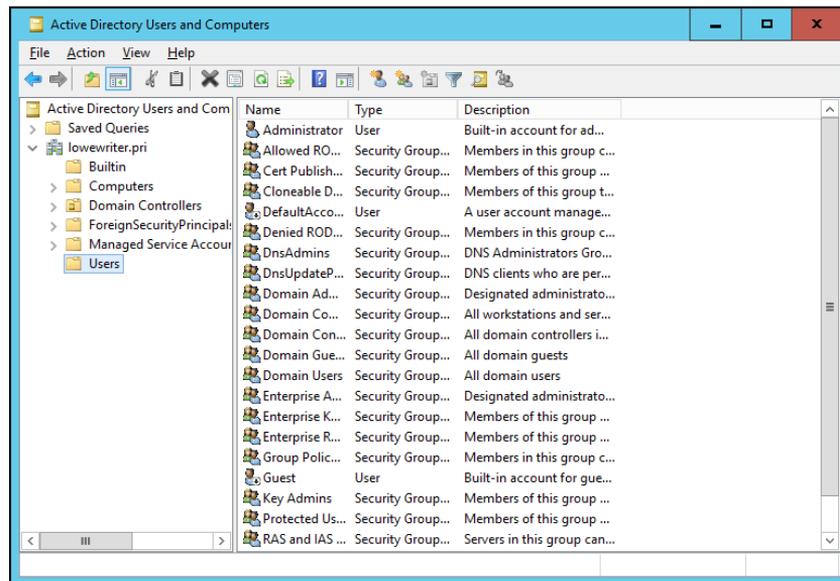


FIGURE 12-1:
The Active Directory Users and Computers management console.

2. **Right-click the domain that you want to add the user to and then choose New ⇨ User from the contextual menu.**

This command summons the New Object – User Wizard, as shown in Figure 12-2.

3. **Enter the user's first name, middle initial, and last name.**

As you fill in these fields, the New Object Wizard automatically fills in the Full Name field.

4. **Change the Full Name field if you want it to appear different from what the wizard proposes.**

You may want to reverse the first and last names so the last name appears first, for example.

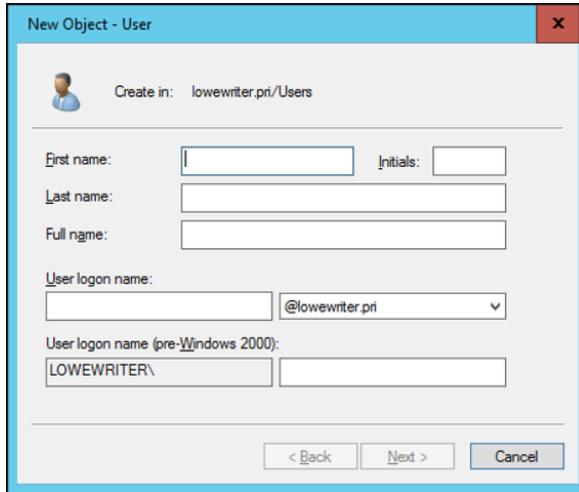


FIGURE 12-2: Use the wizard to create a new user.

5. Enter the user logon name.

This name must be unique within the domain. (Don't worry, if you try to use a name that isn't unique, you'll get an error message.)



TIP

Pick a naming scheme to follow when creating user logon names. You can use the first letter of the first name followed by the complete last name, the complete first name followed by the first letter of the last name, or any other scheme that suits your fancy.

6. Click Next.

The second page of the New Object – User Wizard appears, as shown in Figure 12-3.

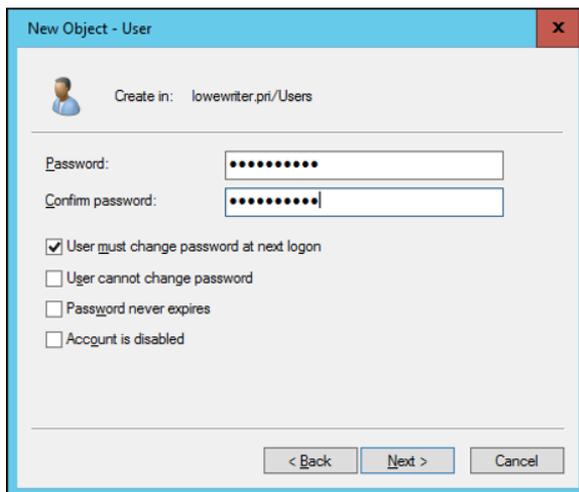


FIGURE 12-3: Set the user's password.

7. Enter the password twice.

You're asked to enter the password and then confirm it, so type it correctly. If you don't enter it identically in both boxes, you're asked to correct your mistake.

8. Specify the password options that you want to apply.

The following password options are available:

- User Must Change Password at Next Logon
- User Cannot Change Password
- Password Never Expires
- Account Is Disabled

For more information about these options, see the section "Setting account options," later in this chapter.

9. Click Next.

You're taken to the final page of the New Object – User Wizard, as shown in Figure 12-4.

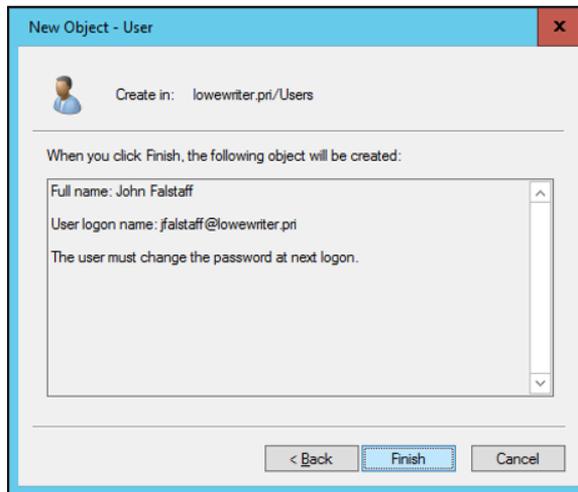


FIGURE 12-4: Verifying the user account information.

10. Verify that the information is correct and then click Finish to create the account.

If the account information isn't correct, click the Back button, and correct the error.

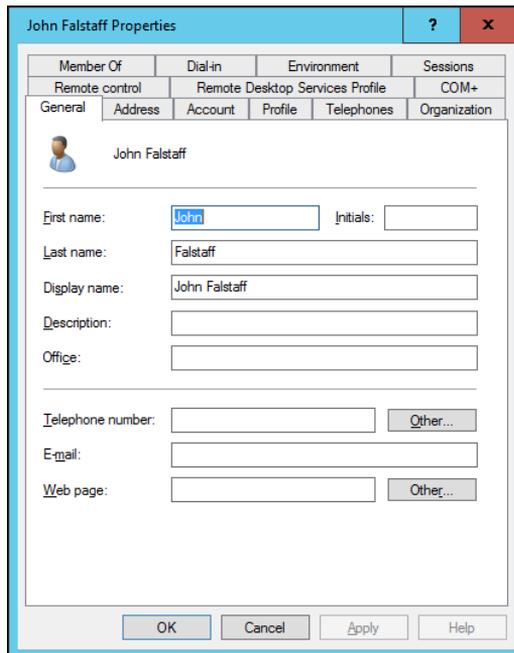
You're done! Now you can customize the user's account settings. At minimum, you'll probably want to add the user to one or more groups. You may also want to add contact information for the user or set up other account options.



An alternative way to create a new user is simply to copy an existing user. When you copy an existing user, you provide a new username and password and Windows copies all the other property settings from the existing user to the new user.

Setting User Properties

After you create a user account, you can set additional properties for the user by right-clicking the new user and choosing Properties from the contextual menu. This command brings up the User Properties dialog box, which has about a million tabs that you can use to set various properties for the user. Figure 12-5 shows the General tab, which lists basic information about the user, such as the user's name, office location, and phone number.



The screenshot shows the 'John Falstaff Properties' dialog box with the 'General' tab selected. The dialog box has a title bar with a question mark and a close button. Below the title bar are several tabs: 'Member Of', 'Dial-in', 'Environment', 'Sessions', 'Remote control', 'Remote Desktop Services Profile', and 'COM+'. The 'General' tab is active and contains the following fields:

- Name:** A small profile picture icon is shown next to the name 'John Falstaff'.
- First name:** A text box containing 'John'.
- Initials:** An empty text box.
- Last name:** A text box containing 'Falstaff'.
- Display name:** A text box containing 'John Falstaff'.
- Description:** An empty text box.
- Office:** An empty text box.
- Telephone number:** An empty text box with an 'Other...' button to its right.
- E-mail:** An empty text box.
- Web page:** An empty text box with an 'Other...' button to its right.

At the bottom of the dialog box are four buttons: 'OK', 'Cancel', 'Apply', and 'Help'.

FIGURE 12-5:
The General tab.

The following sections describe some of the administrative tasks that you can perform via the various tabs of the User Properties dialog box.

Changing the user's contact information

Several tabs of the User Properties dialog box contain contact information for the user, such as

- » **Address:** Change the user's street address, post office box, city, state, zip code, and so on.
- » **Telephones:** Specify the user's phone numbers.
- » **Organization:** Record the user's job title and the name of his boss.

Setting account options

The Account tab of the User Properties dialog box, shown in Figure 12-6, features a variety of interesting options that you can set for the user. You can change the user's logon name, change the password options that you set when you created the account, and set an expiration date for the account.

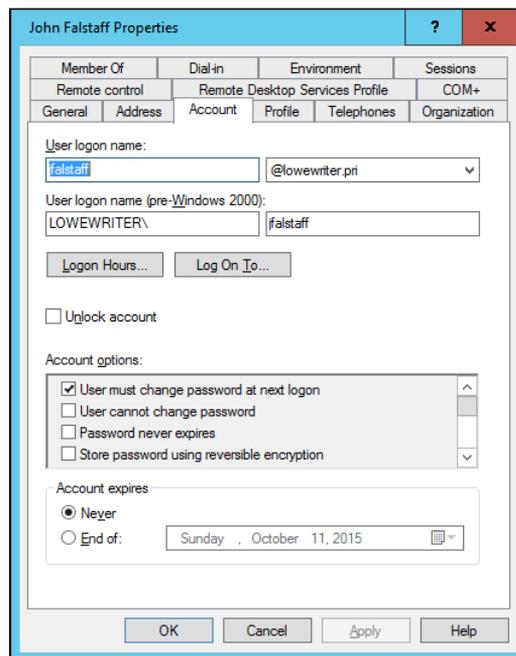


FIGURE 12-6: Set user account info here.

The following account options are available in the Account Options list box:

- » **User Must Change Password at Next Logon:** This default option allows you to create a one-time-only password that can get the user started with the network. The first time the user logs on to the network, he is asked to change the password.
- » **User Cannot Change Password:** Use this option if you don't want to allow users to change their passwords. (Obviously, you can't use this option and the preceding one at the same time.)
- » **Password Never Expires:** Use this option to bypass the password-expiration policy for this user so that the user will never have to change her password.
- » **Store Password Using Reversible Encryption:** This option stores passwords by using an encryption scheme that hackers can easily break, so you should avoid it like the plague.
- » **Account Is Disabled:** This option allows you to create an account that you don't yet need. As long as the account remains disabled, the user won't be able to log on. See the upcoming section, "Disabling and Enabling User Accounts," to find out how to enable a disabled account.
- » **Smart Card Is Required for Interactive Logon:** If the user's computer has a smart card reader to read security cards automatically, select this option to require the user to use it.
- » **Account Is Trusted for Delegation:** This option indicates that the account is trustworthy and can set up delegations. This advanced feature usually is reserved for Administrator accounts.
- » **Account Is Sensitive and Cannot Be Delegated:** This option prevents other users from impersonating this account.
- » **Use DES Encryption Types for This Account:** This option beefs up the encryption for applications that require extra security.
- » **Do Not Require Kerberos Preauthentication:** *Kerberos* refers to a common security protocol used to authenticate users. Select this option only if you are using a different type of security.

Specifying logon hours

You can restrict the hours during which the user is allowed to log on to the system. Click the Logon Hours button on the Account tab of the User Properties dialog box to open the Logon Hours for [User] dialog box, as shown in Figure 12-7.

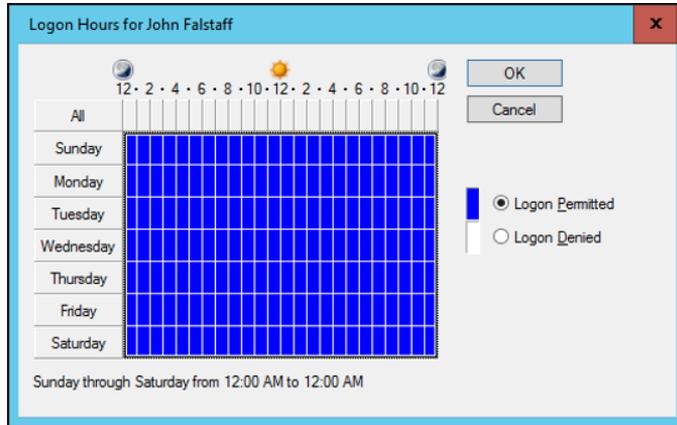


FIGURE 12-7: Restrict a user's logon hours.

Initially, the Logon Hours dialog box is set to allow the user to log on at any time of day or night. To change the hours that you want the user to have access, click a day and time or a range of days and times, select Logon Permitted or Logon Denied, and then click OK.

Restricting access to certain computers

Typically, a user can use his user account to log on to any computer that's part of the user's domain. You can restrict a user to certain computers, however, by clicking the Log On To button on the Account tab of the User Properties dialog box. This button brings up the Logon Workstations dialog box, as shown in Figure 12-8.

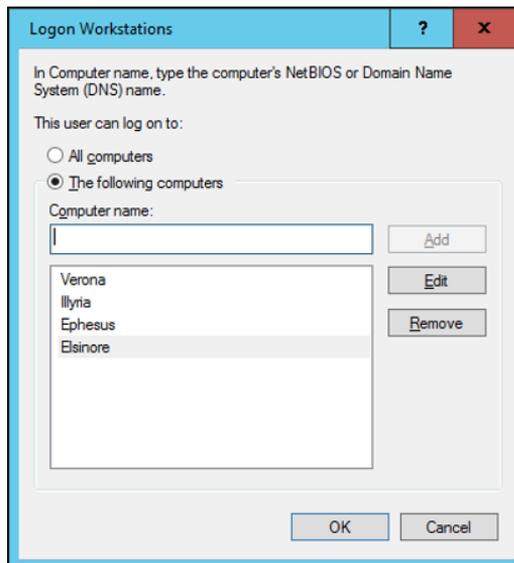


FIGURE 12-8: Restricting the user to certain computers.

To restrict the user to certain computers, select the The Following Computers radio button. Then, for each computer you want to allow the user to log on from, enter the computer's name in the text box and click Add.



TIP

If you make a mistake, you can select the incorrect computer name and then click Edit to change the name. or click Remove to delete the name.

Setting the user's profile information

From the Profile tab, as shown in Figure 12-9, you can configure three bits of information about the user's profile information:

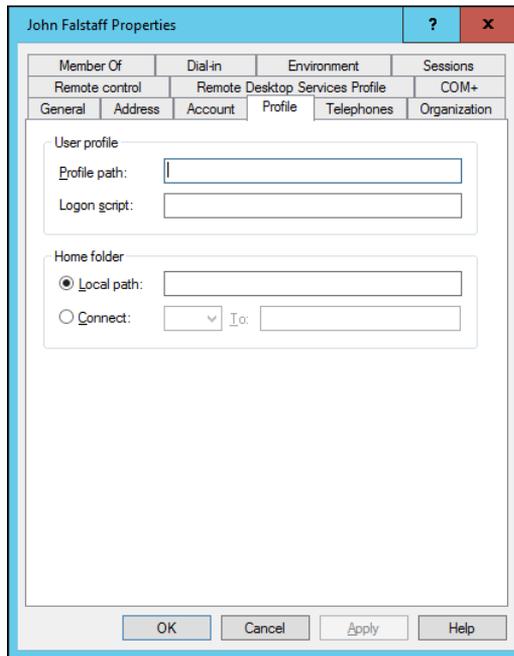


FIGURE 12-9:
The Profile tab.

- » **Profile Path:** This field specifies the location of the user's roaming profile..
- » **Logon Script:** This field is the name of the user's logon script. A *logon script* is a batch file that's run whenever the user logs on. The main purpose of the logon script is to map the network shares that the user requires access to. Logon scripts are carryovers from early versions of Windows NT Server. In Windows Server 2012, profiles are the preferred way to configure the user's computer when the user logs on, including setting up network shares. Many

administrators still like the simplicity of logon scripts, however. For more information, see the section “Creating a Logon Script,” later in this chapter.

» **Home Folder:** This section is where you specify the default storage location for the user.



TIP

From the Profile tab, you can specify the location of an existing profile for the user, but it doesn’t actually let you set up the profile.

Resetting User Passwords

By some estimates, the single most time-consuming task of most network administrators is resetting user passwords. Lest you assume that all users are forgetful idiots, put yourself in their shoes, being made to set their passwords to something incomprehensible (94kD82leL384K) that they have change a week later to something more unmemorable (dJUQ63DWd8331) that they don’t write down. Then network admins get mad when they forget their passwords.

Sooo, when a user calls and says that she forgot her password, the least you can do is (appear to) be cheerful when you reset it. After all, the user probably spent 15 minutes trying to remember it before finally giving up and admitting failure.

Here’s the procedure to reset the password for a user domain account:



REMEMBER

1. **Log on as an administrator.**
You must have administrator privileges to perform this procedure.
2. **Choose Start ⇨ Administrative Tools ⇨ Active Directory Users and Computers.**

The Active Directory Users and Computers management console appears.

3. **In the Active Directory Users and Computers management console, click Users in the console tree.**

Refer to Figure 12-1.

4. **In the Details pane, right-click the user who forgot her password and then choose Reset Password from the contextual menu.**

A dialog box appears allowing you to change the password.

5. **Enter the new password in both password boxes.**

Enter the password twice to ensure that you input it correctly.



TIP

6. (Optional) Select the User Must Change Password at Next Logon option.

If you select this option, the password that you assign will work for only one logon. As soon as the user logs on, she will be required to change the password.

7. Click OK.

That's all there is to it! The user's password is reset.

Disabling and Enabling User Accounts

To temporarily prevent a user from accessing the network, you can disable his account. You can always enable the account later, when you're ready to restore the user to full access. Here's the procedure:

1. Log on as an administrator.

You must have administrator privileges to perform this procedure.

2. From Server Manager, choose Tools ⇨ Active Directory Users and Computers.

3. In the Active Directory Users and Computers management console that appears, click Users in the console tree.

4. In the Details pane, right-click the user that you want to enable or disable; then choose either Enable Account or Disable Account from the contextual menu to enable or disable the user, respectively.

Deleting a User

People come, and people go. And when they go, so should their user account. Deleting a user account is surprisingly easy. Just follow these steps:

1. Log on as an administrator.

You must have administrator privileges to perform this procedure.

2. Choose Start ⇨ Administrative Tools ⇨ Active Directory Users and Computers.

3. In the Active Directory Users and Computers management console that appears, click Users in the console tree.

4. **In the Details pane, right-click the user that you want to delete and then choose Delete from the contextual menu.**

Windows asks whether you really want to delete the user, just in case you're kidding.

5. **Click Yes.**

Poof! The user account is deleted.



WARNING

Deleting a user account is a permanent, nonreversible action. Do it only if you're absolutely sure that you never ever want to restore the user's account. If there's any possibility of restoring the account later, disable the account instead of deleting it. (See the preceding section.)

Working with Groups

A *group* is a special type of account that represents a set of users who have common network access needs. Groups can dramatically simplify the task of assigning network access rights to users. Rather than assign access rights to each user individually, you can assign rights to the group itself. Then those rights automatically extend to any user you add to the group.

The following sections describe some of the key concepts that you need to understand to use groups, along with some of the most common procedures you'll employ when setting up groups for your server.

Creating a group

Here's how to create a group:

1. **Log on as an administrator.**
You must have administrator privileges to perform this procedure.
2. **From Server Manager, choose Tools ⇨ Active Directory Users and Computers.**
The Active Directory Users and Computers management console appears.
3. **Right-click the domain to which you want to add the group and then choose New ⇨ Group from the contextual menu.**
4. **In the New Object – Group dialog box that appears, as shown in Figure 12-10, enter the name for the new group.**

Enter the name in both text boxes.

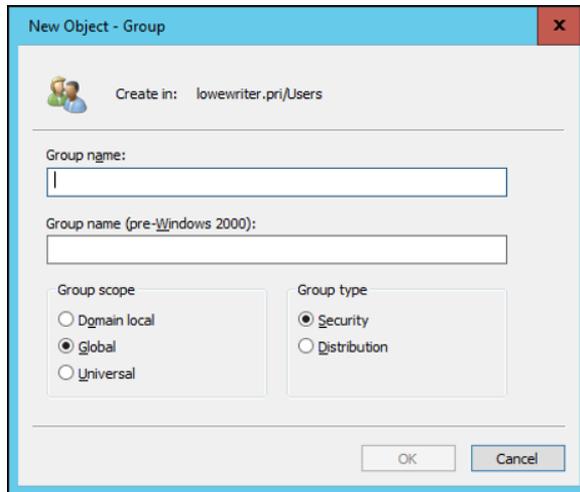


FIGURE 12-10:
Create a new
group.

5. Choose the group scope.

The choices are

- *Domain Local:* For groups that will be granted access rights to network resources
- *Global:* For groups to which you'll add users and Domain Local groups
- *Universal:* If you have a large network with multiple domains

6. Choose the group type.

The choices are Security and Distribution. In most cases, choose Security.

7. Click OK.

The group is created. However, at this point, it has no members. To remedy that, keep reading.

Adding a member to a group

Groups are collections of objects called *members*. The members of a group can be user accounts or other groups. A newly created group (see the preceding section) has no members. As you can see, a group isn't useful until you add at least one member.

Follow these steps to add a member to a group:

1. Log on as an administrator.

You must have administrator privileges to perform this procedure.

2. Choose Start ⇨ Administrative Tools ⇨ Active Directory Users and Computers.

The Active Directory Users and Computers management console appears.

3. Open the folder that contains the group to which you want to add members and then double-click the group.

The Group Properties dialog box appears.

4. Click the Members tab.

The members of the group are displayed, as shown in Figure 12-11.

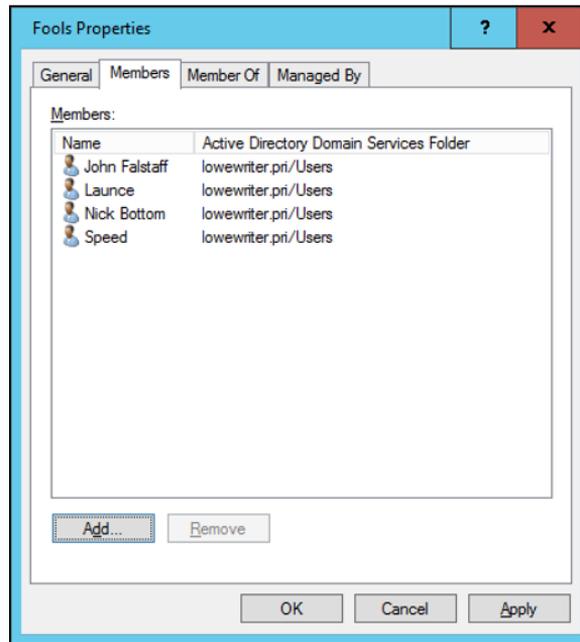


FIGURE 12-11:
Adding members
to a group.

5. Click Add, type the name of a user or other group that you want to add to this group, and then click OK.

The member is added to the list.

6. Repeat Step 5 for each user that you want to add.

Keep going until you add everyone!

7. Click OK.

That's all there is to it.



TIP

On the Member Of tab of the Group Properties dialog box, you can see a list of each group that the current group is a member of.

Creating a Logon Script

A *logon script* is a batch file that's run automatically whenever a user logs on. The most common reason for using a logon script is to map the network shares that the user needs access to. Here's a simple logon script that maps three network shares:

```
echo off
net use m: \\server1\shares\admin
net use n: \\server1\shares\mktg
net use o: \\server2\archives
```

Here, two shares on server1 are mapped to drives M: and N:, and a share on server2 is mapped as drive O:.

If you want, you can use the special variable `%username%` to get the user's username. This variable is useful if you created a folder for each user, and you want to map a drive to each user's folder, as follows:

```
net use u: \\server1\users\%username%
```

If a user logs on with the username `dlowe`, for example, drive U: is mapped to `\\server1\users\dlowe`.



TIP

Scripts should be saved in the `Scripts` folder, which is buried deep in the bowels of the `SYSDVOL` folder — typically, here:

```
c:\Windows\SYSDVOL\Sysvol\domainname\Scripts
```

where *domainname* is your domain name. Because you need to access this folder frequently, I suggest creating a shortcut to it on your desktop.

After you create a logon script, you can assign it to a user by using the Profile tab of the User Properties dialog box. For more information, see the section “Setting the user's profile information,” earlier in this chapter.

Chapter 13

Managing Network Storage

One key purpose of most computer networks is to provide shared access to disk storage. In this chapter, you find out about several ways that a network can provide shared disk storage. Then you discover how to configure Windows Server 2016 to operate as a file server.

Understanding Network Storage

Many network servers exist solely for the purpose of making disk space available to network users. As networks grow to support more users and as users require more disk space, network administrators are continually finding ways to add more storage to their networks. The following sections describe some key concepts for providing network storage.

File servers

A *file server* is simply a network server whose primary role is to share its disk drives. Using a file server is the most common way to provide shared network storage.

A file server can be anything from a simple desktop computer that has been pressed into service as a file server to an expensive (\$25,000 or more) server with redundant components so that the server can continue to run when a component fails. A file server can even consist of advanced disk subsystems with racks of disk drives that can be replaced without shutting down the server.

One of the most common advanced disk subsystems for file servers is Redundant Array of Inexpensive Disks (RAID). A *RAID* system, which is a type of disk storage that hardly ever fails, works by lumping together several disk drives and treating them as though they're a single humongous drive. RAID uses some fancy techniques devised by computer nerds at Berkeley. These techniques ensure that if one of the disk drives in the RAID system fails, no data is lost. The disk drive that failed can be removed and repaired, and the data that was on it can be reconstructed from the other drives.



TIP

Most of this chapter is devoted to showing you how to configure Windows Server 2016 to run as a file server.

Storage appliances

A *storage appliance* is a device specifically designed for providing shared network storage. Also known as Network Attached Storage (NAS), it's a self-contained file server that's preconfigured and ready to run. All you have to do to set it up is take it out of the box, plug it in, and turn it on. Storage appliances are easy to set up and configure, easy to maintain, and less expensive than traditional file servers.

A typical entry-level storage appliance is the Dell PowerVault NX400. This self-contained file server is built into a small rack-mount chassis. It supports up to four hard drives with a total capacity of up to 16 terabytes (TB; that's 16,000 GB). The Dell NX300 runs a special version of Windows Server: Windows Storage Server. This version of Windows, designed specifically for NAS devices, allows you to configure the network storage from any computer on the network by using a web browser.

Note that some storage appliances use customized versions of Linux rather than Windows Storage Server. Also, in some systems, the operating system (OS) resides on a separate hard drive that's isolated from the shared disks so users are prevented from inadvertently damaging the OS.

Understanding Permissions

Before I get into the details of setting up a file server, you need to have a solid understanding of the concept of permissions. *Permissions* allow users to access

shared resources on a network. Simply sharing a resource, such as a disk folder or a printer, doesn't guarantee that a given user is able to access that resource. Windows makes this decision based on the permissions that have been assigned to various groups for the resource and group memberships of the user. For example, if the user belongs to a group that has been granted permission to access the resource, the access is allowed. If not, access is denied.

In theory, permissions sound pretty simple. In practice, however, they can get pretty complicated. The following paragraphs explain some of the nuances of how access control and permissions work:

- » Every object — that is, every file and folder — on an NTFS volume has a set of permissions — the Access Control List (ACL—) associated with it.
- » The ACL identifies which users and groups can access the object and specifies what level of access each user or group has. A folder's ACL may specify that one group of users can read files in the folder, whereas another group can read and write files in the folder, and a third group is denied access to the folder.
- » Container objects — files and volumes — allow their ACLs to be inherited by the objects that they contain. As a result, if you specify permissions for a folder, those permissions extend to the files and child folders that appear within it.

Table 13-1 lists the six permissions that can be applied to files and folders on an NTFS volume.

TABLE 13-1 File and Folder Permissions

Permission	Description
Full Control	The user has unrestricted access to the file or folder.
Modify	The user can change the file or folder's contents, delete the file or folder, read the file or folder, or change the attributes of the file or folder. For a folder, this permission allows you to create new files or subfolders within the folder.
Read & Execute	For a file, this permission grants the right to read or execute the file. For a folder, this permission grants the right to list the contents of the folder or to read or execute any of the files in the folder.
List Folder Contents	This permission applies only to folders; it grants the right to list the contents of the folder.
Read	This permission grants the right to read the contents of a file or folder.
Write	This permission grants the right to change the contents of a file or its attributes. For a folder, this permission grants the right to create new files and subfolders within the folder.

Actually, the six file and folder permissions comprise various combinations of special permissions that grant more detailed access to files or folders. Table 13-2 lists the special permissions that apply to each of the six file and folder permissions.



TIP

Assign permissions to groups rather than to individual users. That way, if a particular user needs access to a particular resource, add that user to a group that has permission to use the resource.

TABLE 13-2 Special Permissions

Special Permission	Full Control	Modify	Read & Execute	List Folder Contents	Read	Write
Traverse Folder/Execute File	*	*	*	*		
List Folder/Read Data	*	*	*	*	*	
Read Extended Attributes	*	*	*	*	*	
Create Files/Write Data	*	*				*
Create Folders/Append Data	*	*				*
Write Attributes	*	*				*
Write Extended Attributes	*	*				*
Delete Subfolders and Files	*					
Delete	*	*				
Read Permissions	*	*	*	*	*	*
Change Permissions	*					
Take Ownership	*					
Synchronize	*	*	*	*	*	*

Understanding Shares

A *share* is simply a folder that is made available to other users via the network. Each share has the following elements:

- » **Share name:** The name by which the share is known over the network
- » **Path:** The path to the folder on the local computer that's being shared, such as C:\Accounting

- » **Description:** A one-line description of the share
- » **Permissions:** A list of users or groups who have been granted access to the share

When you install Windows and configure various server roles, special shared resources are created to support those roles. You shouldn't disturb these special shares unless you know what you're doing. Table 13-3 lists some of the most common special shares.

TABLE 13-3 Special Shares

Share Name	Description
drive\$	The root directory of a drive.
ADMIN\$	Used for remote administration of a computer. This share points to the OS folder (usually, C: \Windows).
IPC\$	Used by named pipes, a programming feature that lets processes communicate with one another.
NETLOGON	Required for domain controllers to function.
SYSVOL	Another required domain controller share.
PRINT\$	Used for remote administration of printers.
FAX\$	Used by fax clients.

Notice that some of the special shares end with a dollar sign (\$). These shares are hidden shares, not visible to users. You can still access them, however, by typing the complete share name (including the dollar sign) when the share is needed. The special share C\$, for example, is created to allow you to connect to the root directory of the C: drive from a network client. You wouldn't want your users to see this share, would you? (Shares such as C\$ are also protected by permissions, of course, so if an ordinary user finds out that C\$ is the root directory of the server's C: drive, he still can't access it.)

Managing Your File Server

To manage shares on a Windows Server 2016 system, open the Server Manager, and select File and Storage Services in the task pane on the left side of the window. Then click Shares to reveal the management console shown in Figure 13-1.

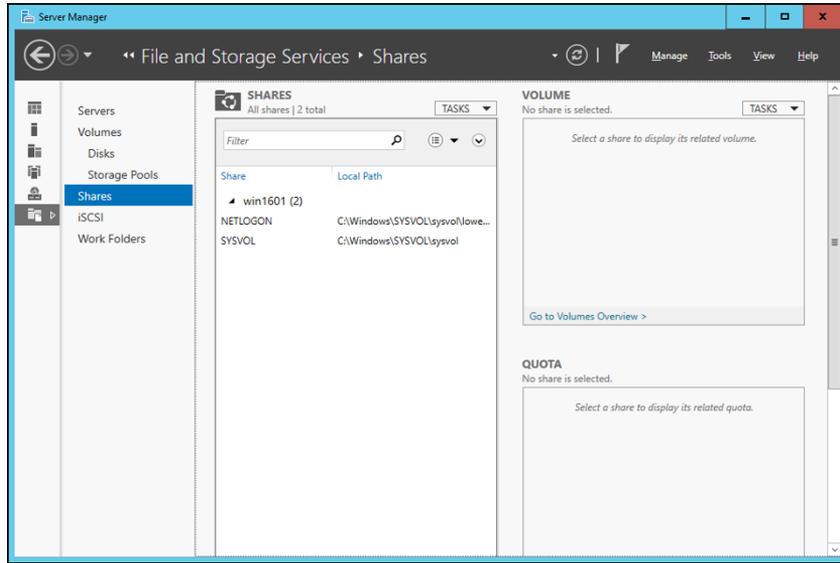


FIGURE 13-1:
Managing shares
in Windows
Server 2016.

The following sections describe some of the most common procedures that you'll use when managing your file server.

Using the New Share Wizard

To be useful, a file server should offer one or more *shares* — folders that have been designated as publicly accessible via the network. To create a new share, use the New Share Wizard:

- 1. In Server Manager, select File and Storage Services, click Shares and then choose New Share from the Tasks drop-down menu.**

The opening screen of the New Share Wizard appears, as shown in Figure 13-2. Here, the wizard asks you what folder you want to share.

- 2. Select SMB Share – Quick in the list of profiles and then click Next.**

The New Share Wizard asks for the location of the share, as shown in Figure 13-3.

- 3. Select the server where you want the share to reside.**

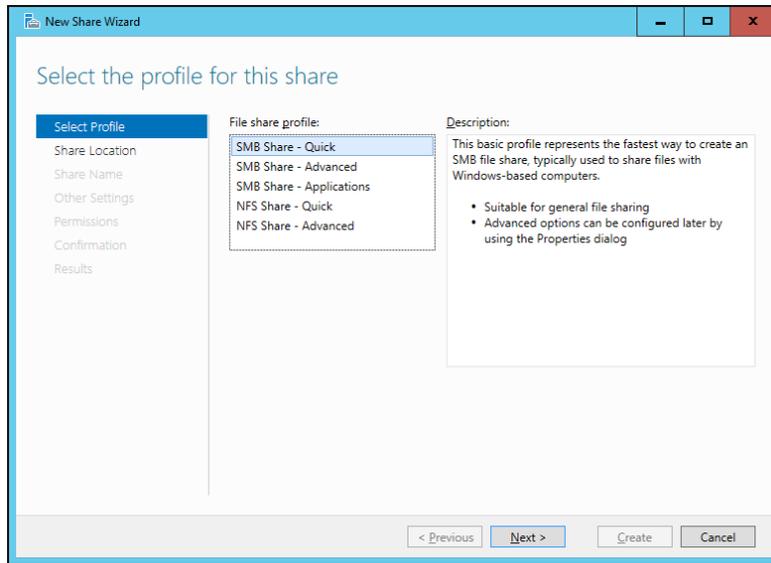


FIGURE 13-2: The New Share Wizard comes to life.

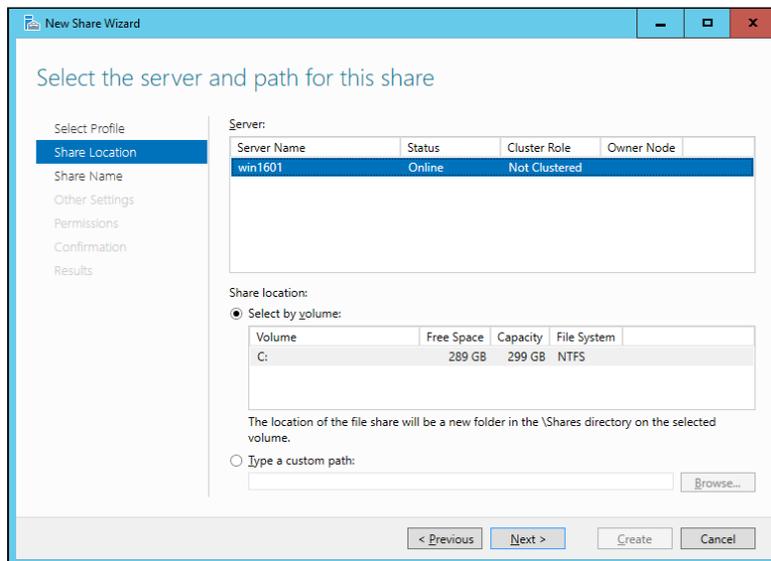


FIGURE 13-3: The wizard asks where you'd like to locate the share.

4. Select the location of the share by choosing one of these two options:

- *Select by Volume:* This option selects the volume on which the shared folder will reside while letting the New Share Wizard create a folder for you. If you select this option, the wizard will create the shared folder on the designated volume. Use this option if the folder doesn't yet exist and you don't

mind Windows placing it in the default location, which is inside a folder called Shares on the volume you specify.

- *Type a Custom Path:* Use this option if the folder exists or if you want to create one in a location other than the Shares folder.

For this example, I chose the Select by Volume example to allow the wizard to create the share in the Shares folder on the C: drive.

5. Click Next.

The dialog box shown in Figure 13-4 appears.

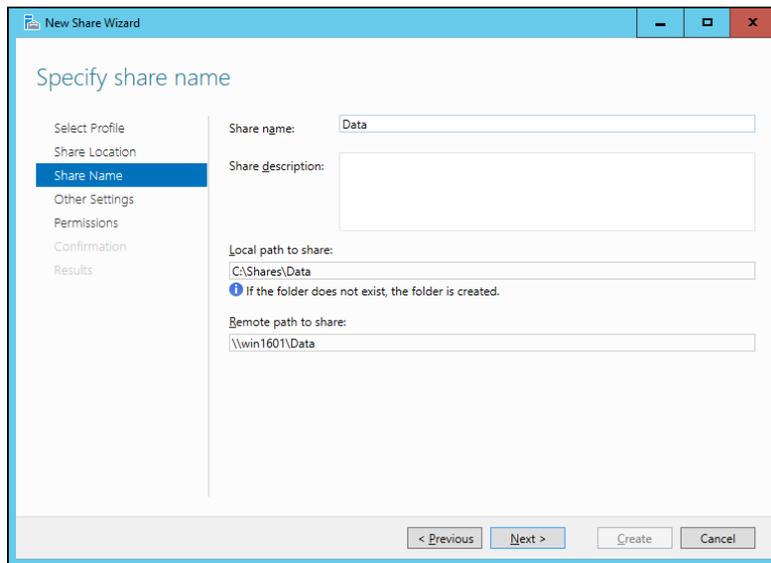


FIGURE 13-4:
The wizard asks for the share name and description.

6. Enter the name that you want to use for the share in the Share Name field.

The default name is the name of the folder being shared. If the folder name is long, you can use a more succinct name here.

For this example, I entered the share name Data.

7. Enter a description for the share.

8. Click Next.

The dialog box shown in Figure 13-5 appears.

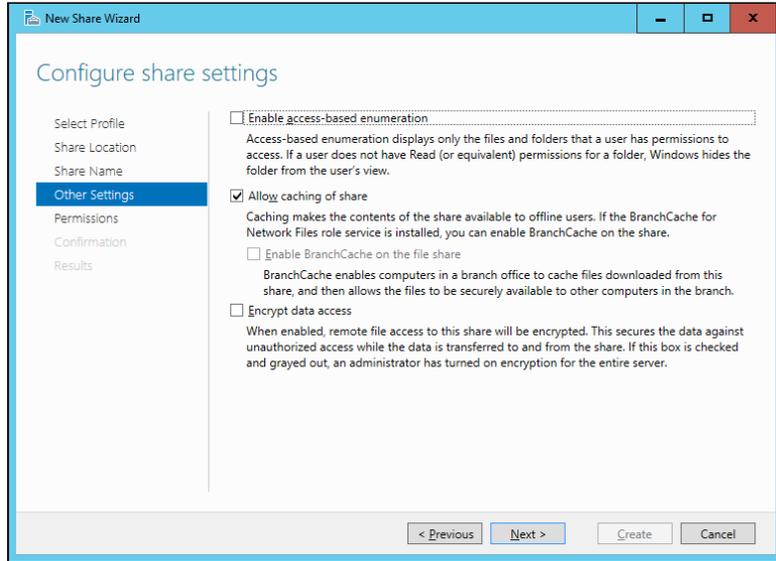


FIGURE 13-5:
Specify the share settings.

9. Select the share settings you'd like to use:

- *Enable Access-Based Enumeration:* Hides files that the user does not have permission to access
- *Allow Caching of Share:* Makes the files available to offline users
- *Encrypt Data Access:* Encrypts files accessed via the share

10. Click Next.

The wizard displays the default permissions that will be used for the new share, as shown in Figure 13-6.

11. (Optional) If you want to customize the permissions, click the Customize Permissions button.

Clicking this button summons the Advanced Security Settings for Data dialog box, where you can customize both the NTFS and the share permissions.

12. Click Next.

The confirmation page appears, as shown in Figure 13-7.

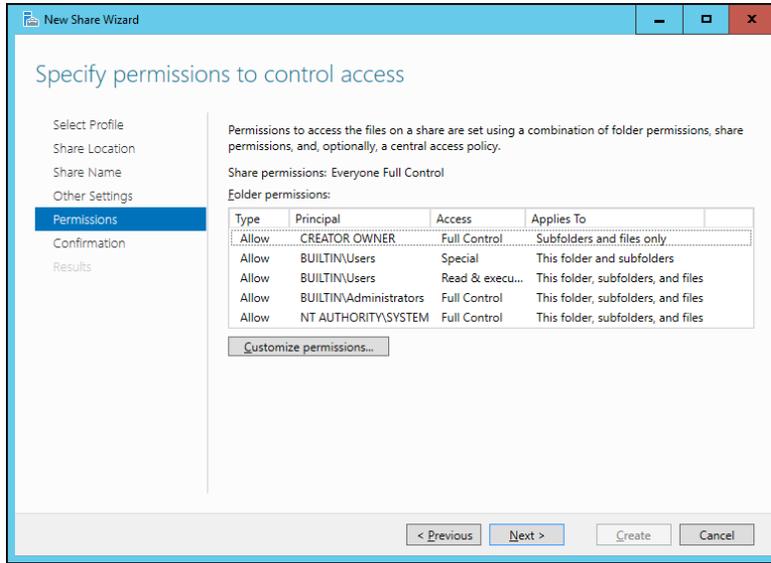


FIGURE 13-6: Setting the share permissions.

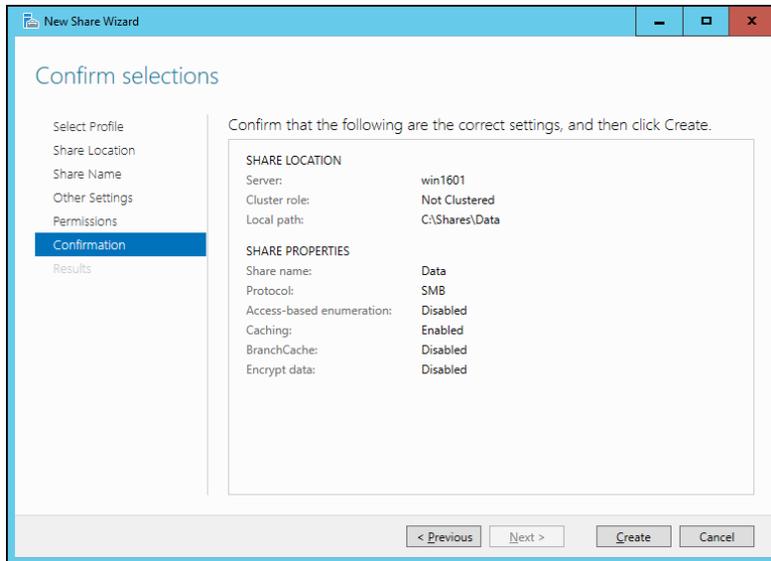


FIGURE 13-7: Confirming your share settings.

13. Verify that all the settings are correct and then click the Create button.

The share is created, and a results dialog box is displayed, as shown in Figure 13-8.

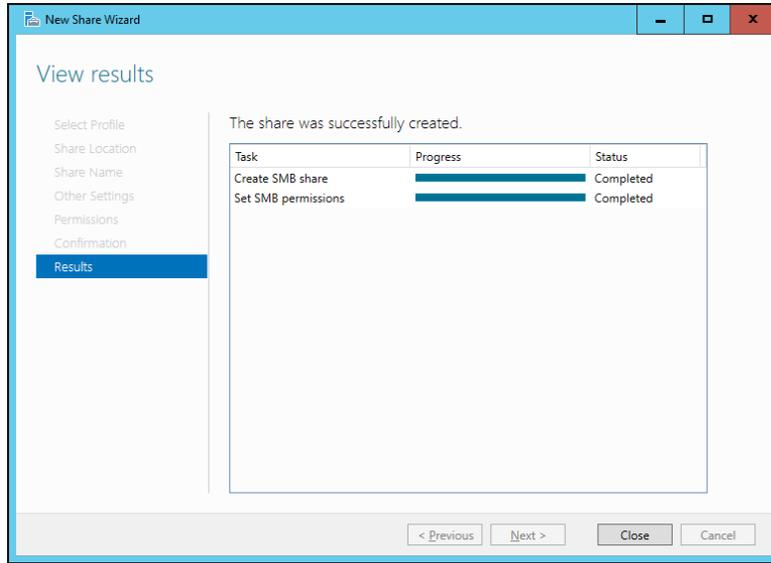


FIGURE 13-8:
You're done!

Sharing a folder without the wizard

If you think wizards should be confined to *Harry Potter* movies, you can set up a share without bothering with the wizard. Just follow these steps:

- 1. Press the Windows key, click Computer, and navigate to the folder that you want to share.**
- 2. Right-click the folder and choose Properties from the contextual menu.**
This action brings up the Properties dialog box for the folder.
- 3. Click the Sharing tab.**
The Sharing tab comes to the front, as shown in Figure 13-9.
- 4. Click the Advanced Sharing button.**
The dialog box shown in Figure 13-10 appears.
- 5. Select the Share This Folder check box to designate the folder as shared.**
The rest of the controls in this dialog box are unavailable until you select this check box.

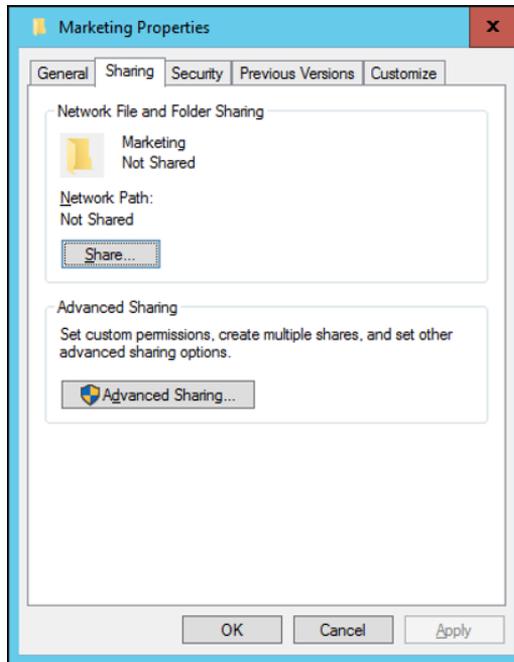


FIGURE 13-9:
Manually share a folder.

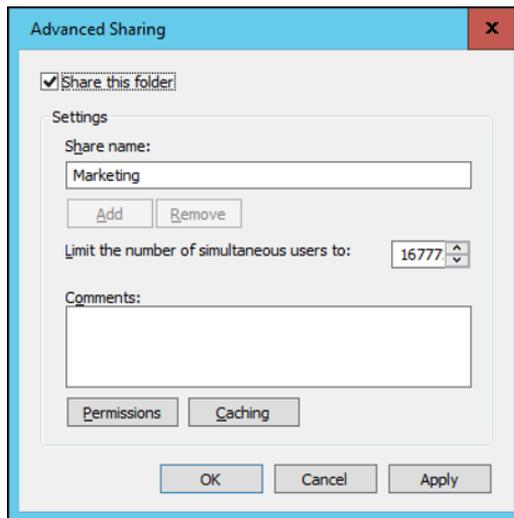


FIGURE 13-10:
Set the share name.

6. Enter the name that you want to use for the share in the Share Name field and then enter a description of the share in the Comments field.

The default name is the name of the folder being shared. If the folder name is long, you can use a more succinct name here.

The description is strictly optional but sometimes helps users determine the intended contents of the folder.

- 7. Click the Permissions button and then set the permissions you want to apply to the share.**

For more information, see the next section.

- 8. Click OK.**

The folder is now shared.

Granting permissions

When you first create a file share, all users are granted read-only access to the share. If you want to allow users to modify files in the share or allow them to create new files, you need to add permissions. Here's how to do this using Windows Explorer:

- 1. Open Windows Explorer by pressing the Windows key and clicking Computer; then browse to the folder whose permissions you want to manage.**
- 2. Right-click the folder you want to manage and then choose Properties from the contextual menu.**

The Properties dialog box for the folder appears.

- 3. Click the Sharing tab; then click Advanced Sharing.**

The Advanced Sharing dialog box appears.

- 4. Click Permissions.**

The dialog box shown in Figure 13-11 appears. This dialog box lists all the users and groups to whom you've granted permission for the folder. Initially, read permissions are granted to a group called Everyone, which means that anyone can view files in the share but no one can create, modify, or delete files in the share.

When you select a user or group from the list, the check boxes at the bottom of the list change to indicate which specific permissions you've assigned to each user or group.

- 5. Click the Add button.**

The dialog box shown in Figure 13-12 appears.

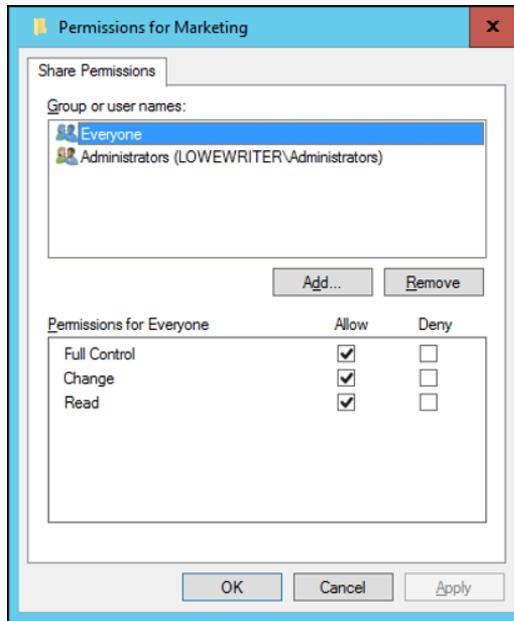


FIGURE 13-11: Set the share permissions.

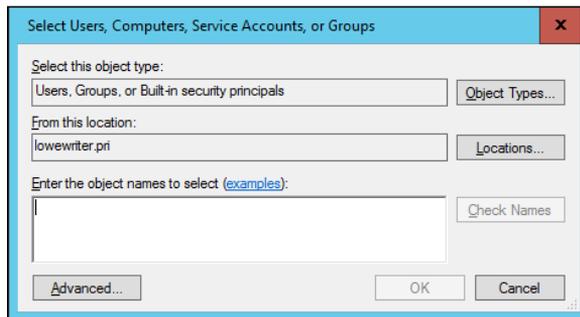


FIGURE 13-12: Adding permissions.

6. Enter the name of the user or group to whom you want to grant permission and then click OK.



TIP

If you're not sure of the name, click the Advanced button. This action brings up a dialog box from which you can search for existing users.

When you click OK, you return to the Share Permissions tab (refer to Figure 13-11), with the new user or group added.

7. Select the appropriate Allow and Deny check boxes to specify which permissions to allow for the user or group.

8. Repeat Steps 5–7 for any other permissions that you want to add.

9. When you're done, click OK.

Here are a few other thoughts to ponder concerning adding permissions:

- » If you want to grant full access to everyone for this folder, don't bother adding another permission. Instead, select the Everyone group and then select the Allow check box for each permission type.
- » You can remove a permission by selecting the permission and then clicking the Remove button.
- » If you'd rather not fuss with the Share and Storage Management console, you can set the permissions from My Computer. Right-click the shared folder, choose Sharing and Security from the contextual menu, and then click Permissions. Then you can follow the preceding procedure, picking up at Step 5.
- » The permissions assigned in this procedure apply only to the share itself. The underlying folder can also have permissions assigned to it. If that's the case, whichever of the restrictions is most restrictive always applies. If the share permissions grant a user Full Control permission but the folder permission grants the user only Read permission, for example, the user has only Read permission for the folder.



REMEMBER

Chapter 14

Managing Exchange Server 2016

Although not strictly a part of Windows Server, Exchange Server 2016 is the mail server software that's used on most Windows networks. Yes, I know Microsoft doesn't call Exchange Server a "mail server." It's a "messaging and collaboration server." But the basic reason for Exchange Server's existence is email. The other messaging and collaboration features are just icing on the cake.

In this chapter, you discover how to perform the most commonly requested maintenance chores for Exchange Server, such as how to create a new mailbox, grant a user access to an additional mailbox, and deal with mailbox size limits.

Exchange Server can be installed within your own network, or you can have Microsoft host Exchange Server for you as part of its cloud-based Office 365. In this chapter, I assume that Exchange Server 2016 has been installed locally within your network. However, the procedures for managing the hosted version of Exchange 2016 are similar.

Creating a Mailbox

When you create a new mailbox, you can either specify an existing Active Directory user or use the Exchange Admin Center to create a new Active Directory user and mailbox at the same time. The following procedure describes the steps you should follow to create a new Active Directory user with a mailbox:

1. **Open the Exchange Admin Center by choosing Start ⇨ All Apps ⇨ Microsoft Exchange Server 2016 ⇨ Exchange Admin Center.**
2. **When prompted, enter your username and password.**

The Exchange Admin Center comes to life, as shown in Figure 14-1.

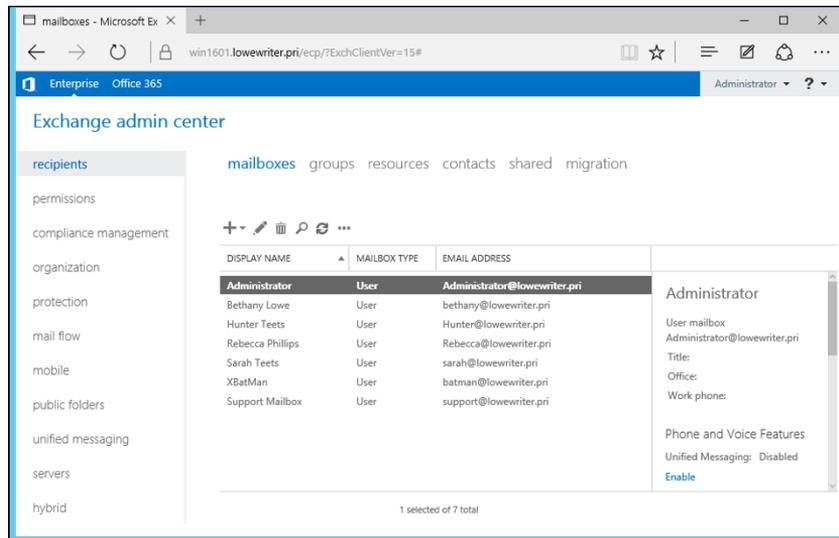


FIGURE 14-1:
The Exchange
Admin Center.

3. **Click the plus sign (+) icon above the list of mailboxes, and then choose User Mailbox.**

The New User Mailbox page appears, as shown in Figure 14-2.

4. **Enter the alias for the new mailbox.**

The alias is the portion of the email address that appears before the at (@) sign. For example, in the email address mitchell@lowewriter.com, the alias is mitchell.

5. **Choose New User.**

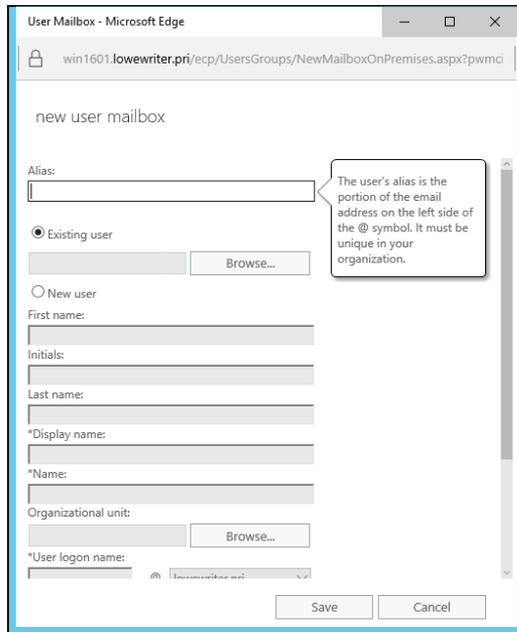


FIGURE 14-2:
Creating a user mailbox.

6. Type the user's first name, middle initial, and last name.

As you type the name, the Display Name field is automatically filled in.

7. Change the Display Name field if you want it to appear different from what was proposed.

For example, you may want to reverse the first and last names so the last name appears first.

8. Enter the user logon name.

This name must be unique within the domain and will be used to form the user's email address.

9. Enter the password twice.

You're asked to type the password twice, so type it correctly. If you don't type it identically in both boxes, you're asked to correct your mistake.

10. If the password is temporary, select the User Must Change Password at Next Logon check box.

This setting requires the user to change the temporary password the first time he or she logs on.

11. Click Save.

The user's mailbox is created and appears on the Recipients page, as shown in Figure 14-3.

12. Pat yourself on the back; then click Finish.

You're done!

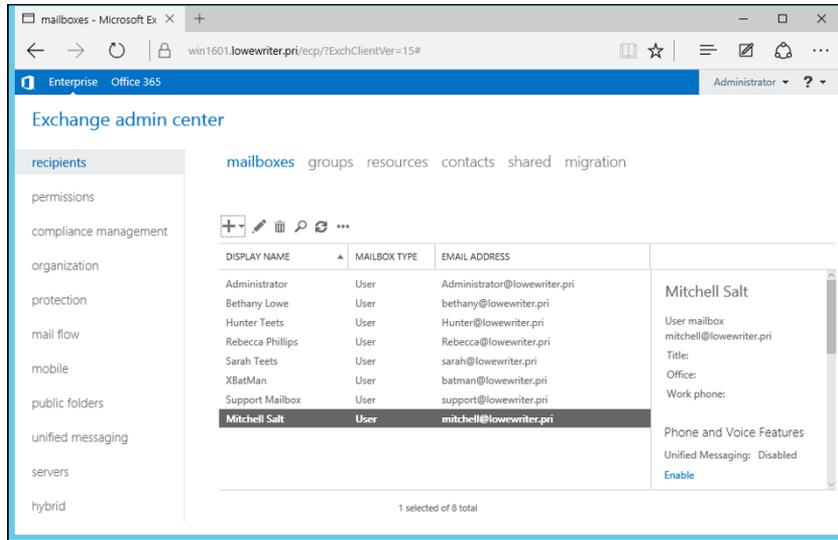


FIGURE 14-3:
The new user's mailbox appears on the Recipients page.

Managing Mailboxes

After you set up a mailbox, you can use the Exchange Admin Center to manage the basic settings of the mailbox. To do that, select the mailbox you want to manage and click the Edit icon. This action brings up the Mailbox page, which is the portal that grants access to many of the most frequently used features of Exchange.

The following sections describe several commonly used features that are controlled via this page.

Enabling mailbox features

Exchange Mailbox Features refers to several features of Exchange mailboxes that are controlled via the Mailbox Features tab of the mailbox User Mailbox page. This tab is shown in Figure 14-4.

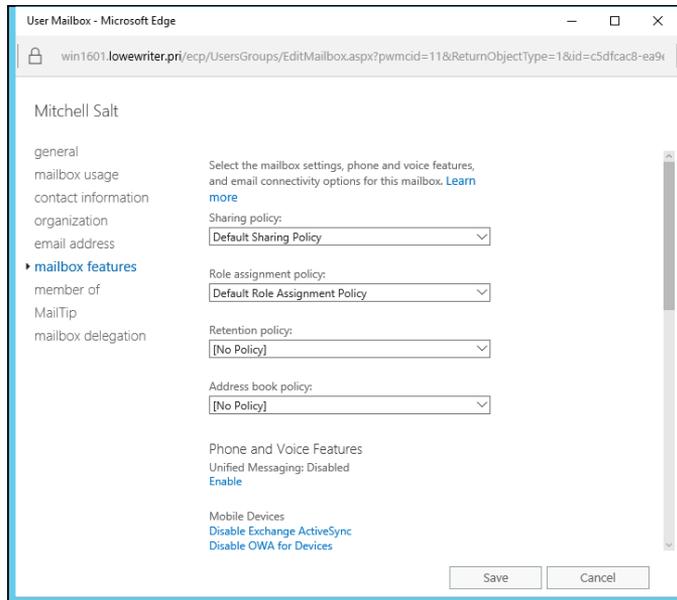


FIGURE 14-4:
The Mailbox
Features tab.

The following paragraphs describe the features that are controlled from this tab:

- » **Outlook Web App:** Lets the user access her Exchange mailbox from a web browser rather than from an Outlook client. With this feature enabled, the user can read email from any computer that has an Internet connection. This feature used to be called Outlook Web Access.
- » **Exchange ActiveSync:** Activates the ActiveSync feature, which allows Exchange data to synchronize with mobile devices, such as iPhones or Windows Mobile phones.
- » **MAPI:** Enables email using the MAPI protocol.
- » **POP3:** Enables email using the POP3 protocol.
- » **IMAP4:** Enables email using the IMAP4 protocol.
- » **Archiving:** Enables the Exchange Archive feature, which is available only with the Enterprise edition of Exchange.
- » **Mail Flow:** Lets you set delivery options such as creating an automatic forwarder, as described in the next section.
- » **Message Size Restrictions:** Lets you set the maximum size of incoming or outgoing messages.

Creating a forwarder

A *forwarder* is a feature that automatically forwards any incoming email to another email address. This feature is most often used when an employee is on vacation or leave, and the employee's manager requests that someone else temporarily handle the absent employee's email.

1. In the Exchange Admin Center, open the User Mailbox page for the user.
2. Select the Mailbox Features tab, and then click View Details in the Mail Flow section.

The delivery options are displayed, as shown in Figure 14-5.

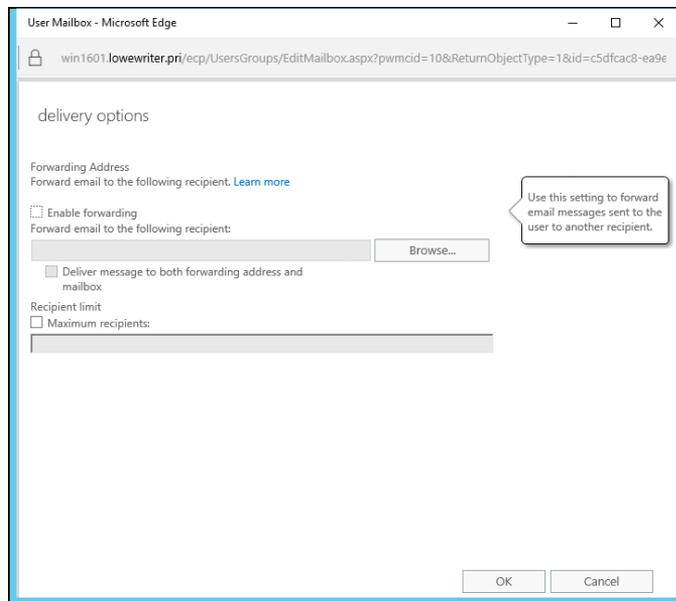


FIGURE 14-5:
Setting the delivery options.

3. Select the **Enable Forwarding** check box.
 4. Click the **Browse** button.
- The Select Recipient window appears, as shown in Figure 14-6.
5. Select the recipient you want to forward the email to and then click **OK**.
- The name you selected is displayed in the text box next to the Browse button in the Delivery Options page (refer to Figure 14-5).
6. If you want the email to be delivered to this user's mailbox in addition to the forwarding address, select the **Deliver Message to Both Forwarding Address and Mailbox** check box.

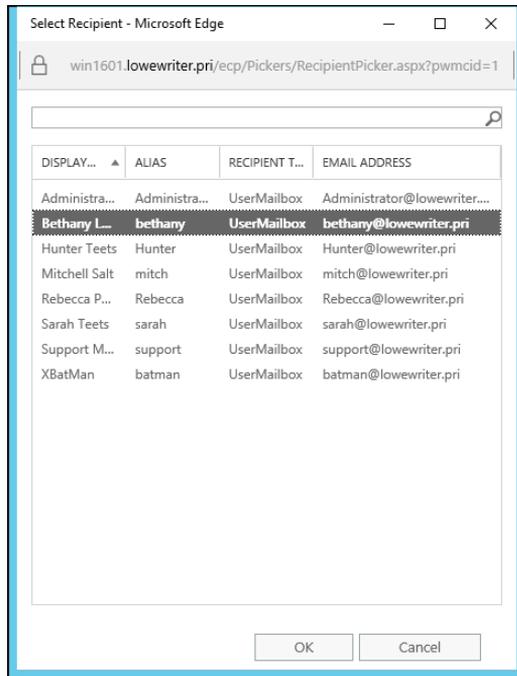


FIGURE 14-6:
Select the recipient for the forwarder.

If you leave this option deselected, only the forwarding address will receive the email; the mail won't be delivered to this user's mailbox.

7. Click OK to close the Delivery Options page.

You return to the User Mailbox page.

8. Click Save to save your changes.

Setting mailbox storage limits

Exchange lets you set a limit on the size of each user's mailbox. In a very small organization, you can probably get away without imposing strict mailbox size limits. If your organization has 20 or more users, though, you need to limit the size of each user's mailbox to prevent the Exchange private mail store from getting out of hand.

Exchange provides three kinds of storage limits for user mailboxes:

- » **Issue Warning At:** When this limit is exceeded, an email warning is sent to the user to let him know that his mailbox is getting large.
- » **Prohibit Send At:** When this limit is reached, the user can't send email, but the mailbox continues to receive email. The user won't be able to send emails again until she deletes enough emails to reduce the mailbox size below the limit.

- » **Prohibit Send and Receive At:** When this limit is reached, the mailbox shuts down and can neither send nor receive emails.
- » **Keep Deleted Items for (Days):** Most users don't realize it, but when they permanently delete an item from their mailbox, the item isn't really permanently deleted. Instead, Exchange retains the item for the period specified by this limit. The default is 14 days.
- » **Keep Deleted Mailboxes for (Days):** This limit specifies how long Exchange should retain mailboxes that you delete.

You can (and should) set a default storage limit that applies to all mailboxes in your organization. You can also override these limits for specific users. The limits you set will depend on many factors, including the number of users in your organization, the type of email they typically use (for example, do they require large attachments?), and the amount of disk space available on your Exchange server.

1. In the Exchange Admin Center, select Databases.

A list of databases for your Exchange environment appears, as shown in Figure 14-7. In this example, there is just one database. For larger Exchange environments, you may see several databases.

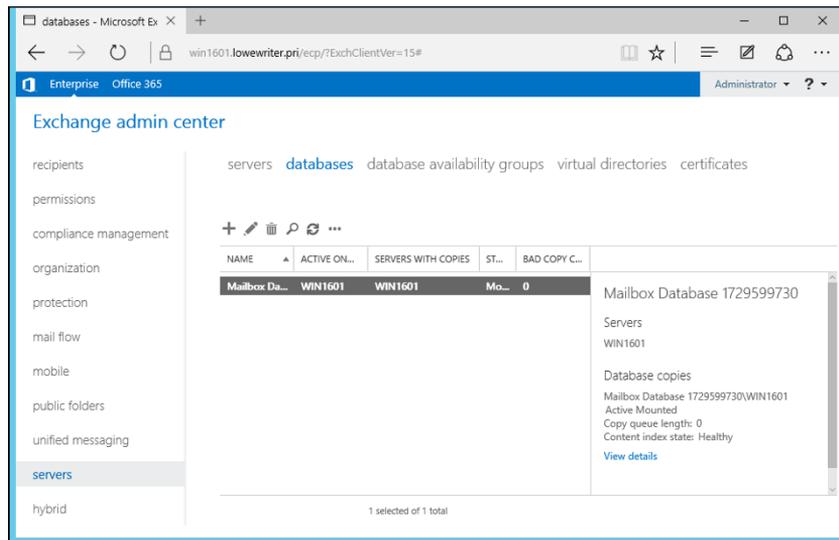


FIGURE 14-7: The Exchange Admin Center Databases page.

2. Double-click the database whose limits you want to change.

The Mailbox Database page appears, as shown in Figure 14-8.

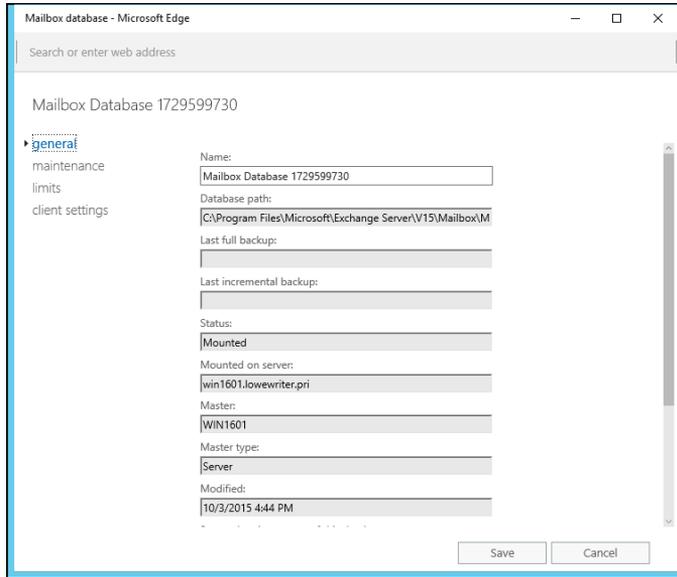


FIGURE 14-8:
The Mailbox
Database page.

3. Select Limits.

The Mailbox Database Limits page is displayed, as shown in Figure 14-9.

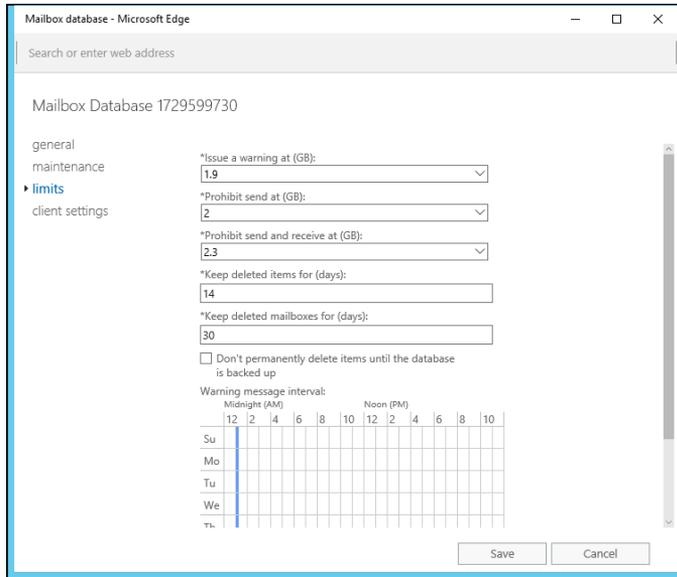


FIGURE 14-9:
The Mailbox
Database Limits
page.

4. Change the Storage Limits settings to meet your needs.

By default, the storage limits are set quite high: Warnings are issued at about 1.9GB, send permission is revoked at 2GB, and both send and receive permissions are revoked at 2.3GB. These limits are generous, but bear in mind that if you have 100 users, your mailbox database may grow to 200GB. You may want to set lower limits.

5. Click Save.

The limits you set take effect immediately.

If you impose restrictive default storage limits for your users, you may want to relax the limits on a case-by-case basis. Some users may require a larger mailbox because of the type of work they do, and you probably don't want to impose a tight limit on your boss.

Fortunately, overriding the default limits for a specific user is easy. Here are the steps:

1. In the Exchange Admin Center, choose Recipients, and then double-click the mailbox you want to edit.

The user's mailbox page appears.

2. Select Mailbox Usage.

The Mailbox Usage page appears, as shown in Figure 14-10.

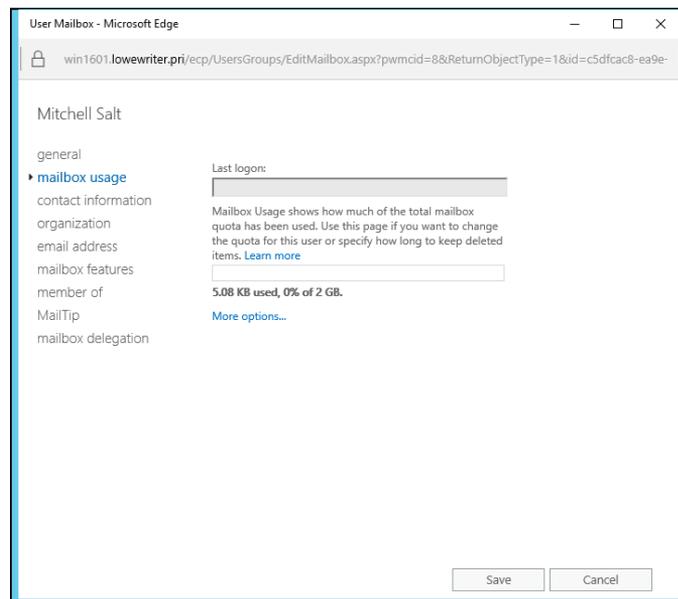


FIGURE 14-10:
The Mailbox Usage page.

3. Click More Options.

The mailbox limits options are displayed, as shown in Figure 14-11.

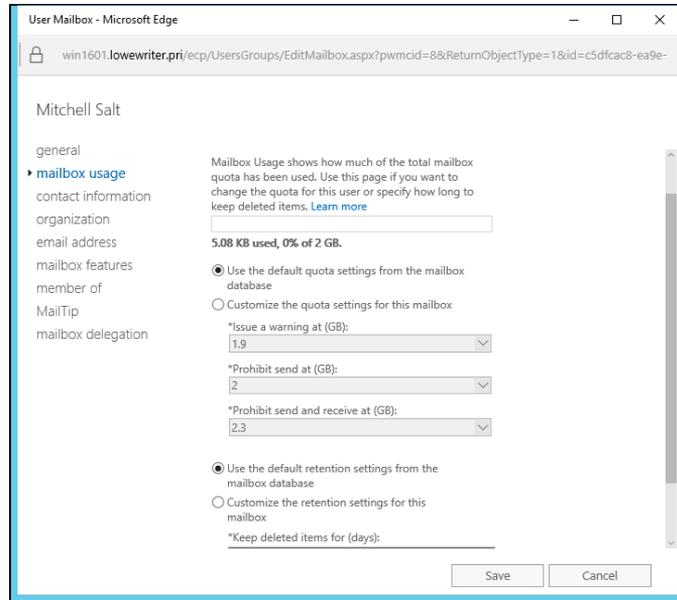


FIGURE 14-11:
Setting the
default storage
limits.

4. Change the Storage Limits settings to meet your needs.

You can adjust any of the mailbox limits up or down to create limits for the user that are either more or less restrictive than the defaults for the database.

5. Click Save.

The limits you set take effect immediately.

Configuring Outlook for Exchange

After you create an Exchange mailbox for a user, you can configure that user's Outlook client software to connect to the user's account. Although you can do this configuration directly within Outlook, it's better to do it outside Outlook, using the Control Main Mail applet. Here are the steps:

1. Open Control Panel and then open the Mail applet.

The dialog box shown in Figure 14-12 appears.

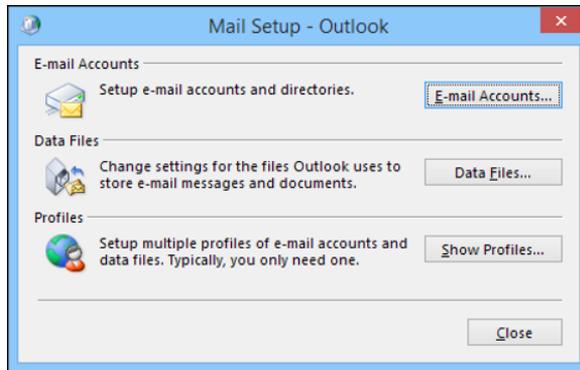


FIGURE 14-12:
The Mail Setup dialog box.

2. Click the Show Profiles button.

The dialog box shown in Figure 14-13 appears, listing the mail profiles that already exist on the computer.

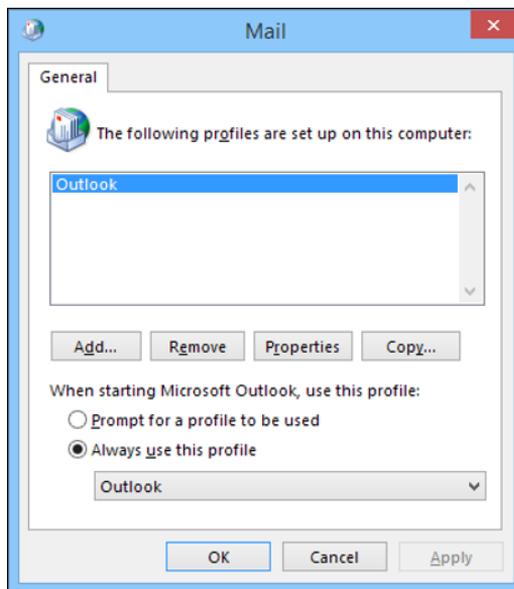


FIGURE 14-13:
Viewing mail profiles.

3. Double-click the user's profile.

The Mail Setup dialog box shown in Figure 14-14 appears.

4. Click the E-mail Accounts button.

The Account Settings dialog box appears, as shown in Figure 14-15.

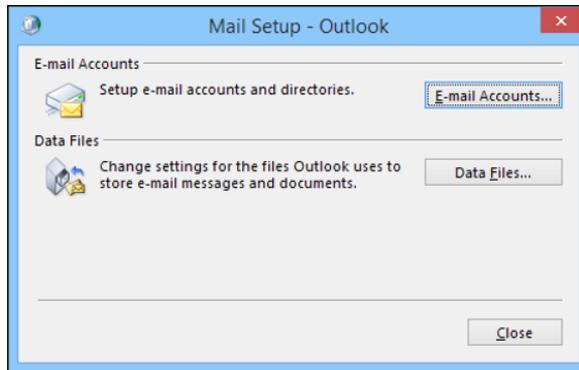


FIGURE 14-14:
The Mail Setup dialog box.

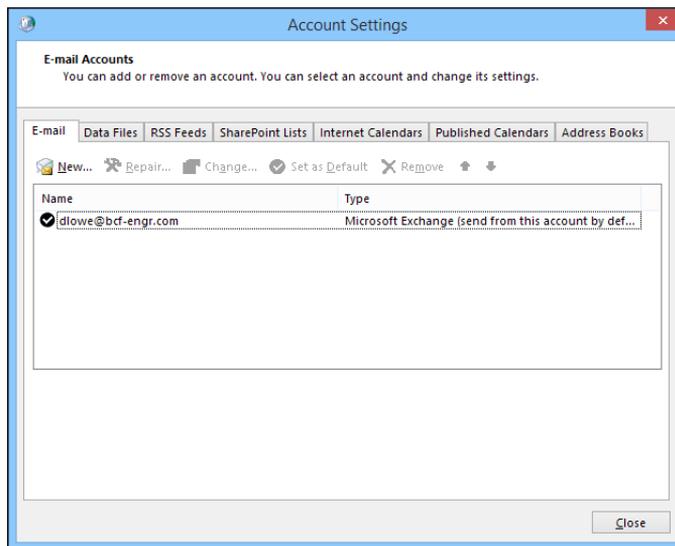


FIGURE 14-15:
The Account Settings dialog box.

5. Click the New icon.

An Add E-mail Account dialog box appears.

Don't enter your email address as prompted in this dialog box; instead, proceed to Step 6.



WARNING

6. Click the Manually Configure Server Settings or Additional Server Types option and then click Next.

A dialog box asks you what type of email account you want to create. The choices are:

- Outlook.com or Exchange ActiveSync Compatible Service
- POP3 or IMAP

7. Select the Outlook.com or Exchange ActiveSync Compatible option and then click Next.

The dialog box shown in Figure 14-16 appears.

The screenshot shows a Windows-style dialog box titled "Add Account". At the top, it says "Server Settings" and "Enter the information that is required to connect to an Exchange ActiveSync service." Below this are several sections: "User Information" with "Your Name:" and "E-mail Address:" fields; "Server Information" with "Mail server:" field; "Logon Information" with "User Name:" and "Password:" fields, and a checked "Remember password" checkbox; and "Offline Settings" with a "Mail to keep offline:" slider set to "All". At the bottom are buttons for "< Back", "Next >", and "Cancel".

FIGURE 14-16: You must identify the Exchange server and provide a username.

8. Enter the name of the Exchange server and the username in the appropriate text boxes; then click Next.

After your account information is verified, a confirmation message is displayed.

9. Click OK.

The confirmation message goes away, and the last page of the E-Mail Accounts Wizard appears.

10. Click the Finish button.

The wizard is dismissed.

11. Choose File ⇨ Exit to close Outlook.

12. Restart Outlook.

The mailbox should be configured.

IN THIS CHAPTER

Getting acquainted with intranets

Finding good uses for intranets

Figuring out what you need to set up an intranet

Setting up an IIS web server

Setting up a simple intranet

Managing IIS

Chapter 15

Creating an Intranet

No, I'm not mispronouncing *Internet*, although an intranet is similar to the Internet, but with a twist. Rather than connect your computer to millions of other computers around the world, an intranet connects your computer to other computers in your company or organization. How is an intranet different from your ordinary, run-of-the-mill network? Read on, and I'll explain.

Defining an Intranet

Everyone knows that the Internet is the best thing since fried Twinkies. Millions upon millions of computer users surf the web every day, searching for information, sharing stuff, downloading files, and doing all the things that people do these days on the Internet.

Companies realized long ago that such a platform would help employees share and access data and documents. But understandably, not many companies wanted their proprietary and private info all over the Internet. So ingenious network managers at large companies figured out that although “the” web is a great conduit for distributing public information around the world, “a” web — a private internal web — is even better for distributing sensitive information within a company. Thus, the idea of intranets was born. An *intranet* is simply a network that's

built by using the same tools and protocols that are used by the global Internet but applied instead to an organization's internal network.



TIP

Think of an intranet as a small, private version of the World Wide Web. Anyone who connects to your local area network (LAN) can access your intranet. Like the Internet, an intranet is accessed by using a web browser, such as Internet Explorer or Chrome. However, an Intranet doesn't cross the boundary of your company's firewall onto the public Internet because the information on the intranet is stored on the company's server computers rather than on a computer outside of your organization.

More important, an intranet cannot be accessed by people *outside* of your organization — or, more specifically, by people on the other side of your firewall, unless your firewall is specifically configured to enable outside access. Thus, an intranet can be viewed by people within your organization, but not by people on the outside.



REMEMBER

An intranet is analogous to a closed-circuit television system that can be viewed only by people within the organization that owns the system. In contrast, the Internet is more like cable television in that anyone who's willing to pay a monthly fee can watch.

Here are two interesting but contradictory points of view about the significance of intranets:

- » Some computer industry pundits say that intranets are more popular than the Internet. For example, many companies that sell web development tools make more money selling software used for intranets than for the Internet.
- » On the other hand, other industry pundits think that the intranet phenomenon is merely a fad that some other promising new technology, such as pet rocks or hula hoops, will replace in a few years. Only time will tell.

Identifying Intranet Uses

Intranets can distribute just about any type of information within a company. Intranets use three basic types of applications:

- » **Publishing:** Information is posted in the form of pages that you can view from any computer with access to the intranet. This type of intranet application is commonly used for company newsletters, policy manuals, and price lists, for example.



TIP

Publishing applications are simple to set up. In fact, you may be able to set up one without a lot of outside help from highly paid computer consultants.

- » **Transactional:** Information is gathered from users of the intranet who file online expense reports, report problems to the help desk, or enroll in employee benefit programs, for example.



WARNING

Expect to spend big bucks on computer consulting to get an intranet transaction application set up.

- » **Social:** Social intranet applications are designed to mimic the features of social Internet applications like Facebook or Twitter. They provide features that enable social interaction such as chatting, engaging in conversations, sharing pictures or other documents, and other forms of collaboration.

Setting Up an Intranet

To properly set up an intranet, you need the right tools. Here's a list of requirements:

- » **A network:** An intranet doesn't require its own cabling; it can operate on your existing network.
- » **A server computer that's dedicated to the intranet:** Make sure that this computer has plenty of RAM (at least 4GB) and gigabytes of disk space (at least 100GB). Of course, the more users your network has and the more information you intend to place on the server, the more RAM and disk storage you need.
- » **Windows Server or a Linux operating system:** Web server software requires one or the other.
- » **Web server software for the server computer:** You need to install a web server, such as IIS (for Windows servers) or Apache (for Linux servers).
- » **Programs to help you create web pages:** If you're the type who dreams in binary, you can create web pages by typing HTML codes directly into text files. In that case, the only program you need is Notepad. Alternatively, you can use a program designed specifically for creating web pages, such as Microsoft FrontPage, or perhaps something fancier, such as Adobe Dreamweaver. If you're going to develop transaction-based applications, you need additional tools.

A WEBLESS INTRANET

The correct way to set up a proper intranet is to set up a Windows-based server running IIS or a Linux-based server running Apache or some other web server. However, you can create a rudimentary intranet without going to the trouble of setting up an actual web server. Here's how:

- 1. Set up a share on a file server that will hold the HTML files that make up your intranet.**
- 2. Create an HTML file for the home page of your intranet, and save the file in the location you create in Step 1.**

I recommend that you name it `index.html`.

- 3. Create any other HTML files that your intranet needs.**

The `index.html` file should include links to these pages.

- 4. Point your web browser to the `index.html` file at the shared network location.**

For example, if the server is named `iserver` and the share is named `intranet`, enter this information into your browser's address box: `\\iserver\intranet\index.html`. *Voilà!* — you have an instant intranet without the fuss of a web server.

This rudimentary intranet works without a web server because a web browser can display HTML files directly, without the need for a web server. However, without a web server, your intranet is limited in what it can do. In particular, all its pages must be *static* (their content is fixed). For *dynamic* content, which users interact with, you need to set up a web server.

A common way to set up an intranet is to use Microsoft's SharePoint, which is web-based software specifically designed for creating a company intranet. For more information about SharePoint, check out *SharePoint 2013 For Dummies*, by Ken Withee (Wiley). You can also find a chapter about using SharePoint in my *Networking All-In-One For Dummies*, 6th Edition (Wiley).

Setting Up an IIS Web Server

IIS is a free component of Windows Server 2016, but it's not installed by default. After you complete installing Windows Server, you must add the Web Server role to enable IIS. The following procedure is for Windows Server 2016, but the procedure for Windows 2008 Server (or 2003, for that matter) is similar:

1. Open the Server Manager; then choose Add Roles and Features.

The Add Roles and Features Wizard comes to life.

2. Follow the steps of the Add Roles and Features Wizard up to the Select Server Roles step.

The Select Server Roles page is shown in Figure 15-1.

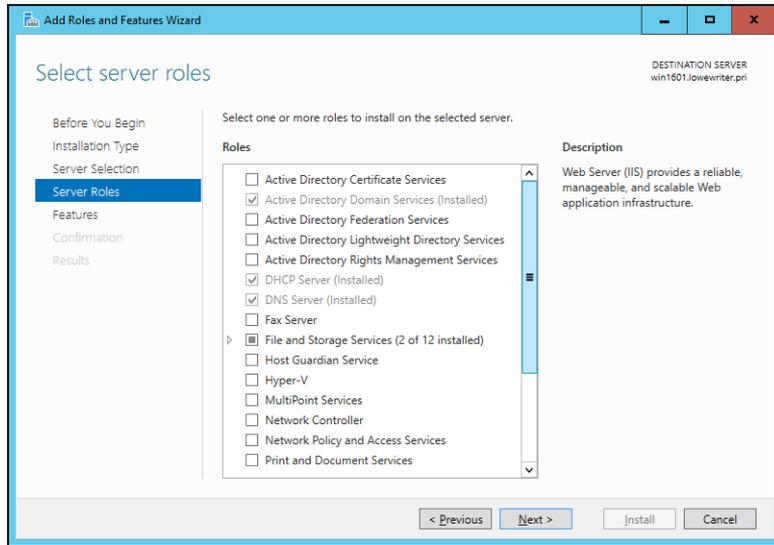


FIGURE 15-1:
The Select Server Roles page of the Add Roles and Features Wizard.

3. Select the Web Server (IIS) check box and then click Next.

The Add Roles and Features Wizard asks whether you want to install the related IIS Management Console, as shown in Figure 15-2.

4. Click the Add Features button; then click Next.

The Select Features page appears.

5. Click Next.

The Features page appears.

6. Click Next.

The Web Server Role (IIS) page appears, as shown in Figure 15-3.

7. Click Next.

The Select Role Services page appears, as shown in Figure 15-4. This page lists a variety of optional services that can be configured for IIS.

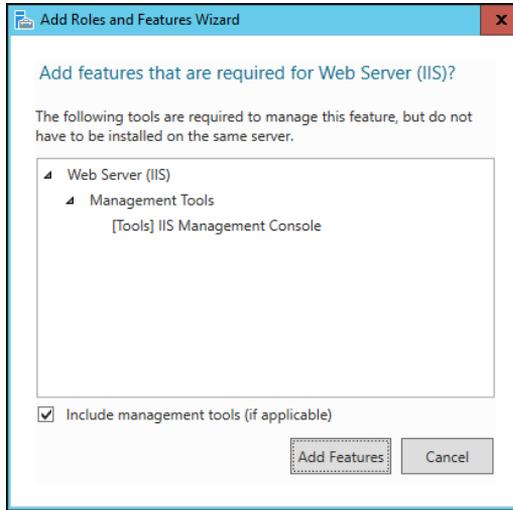


FIGURE 15-2: Installing the IIS Management Console.

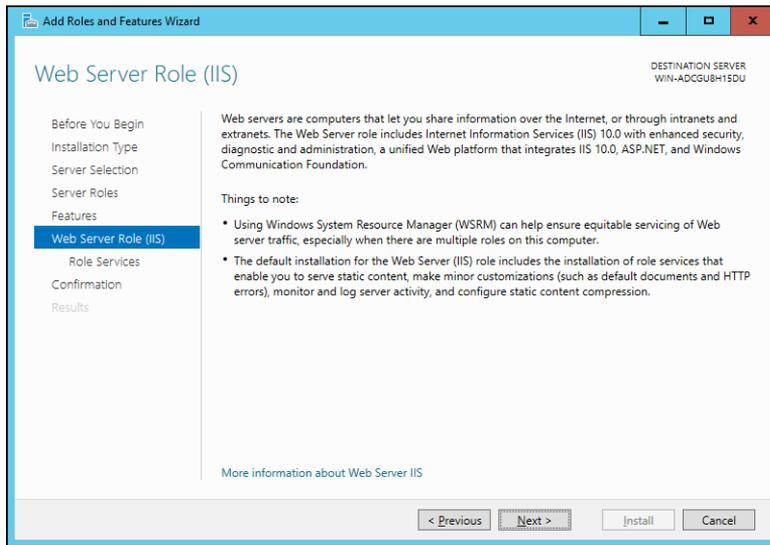


FIGURE 15-3: The Web Server Role (IIS) page of the Add Roles and Features Wizard.

8. Select the services you want to configure for IIS.

If you want, you can study this list and try to anticipate which features you think you'll need. Or you can just leave the default options selected.



TIP

You can always return to the Add Roles and Features Wizard to add features you leave out here.

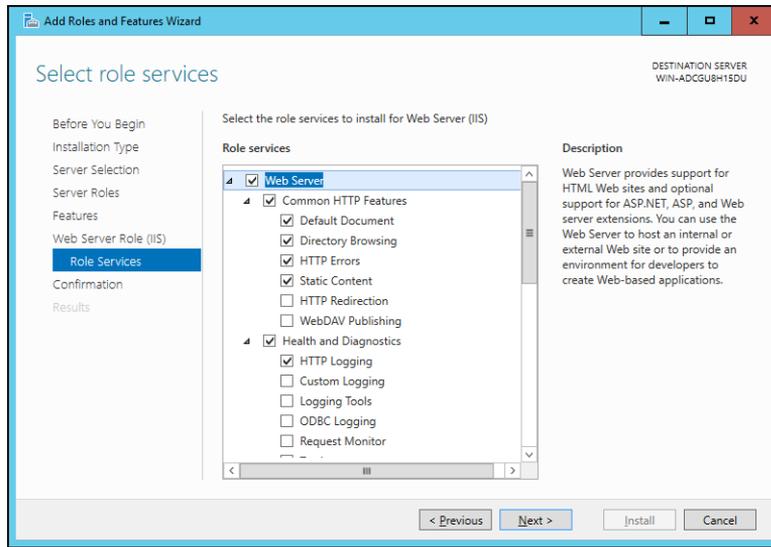


FIGURE 15-4: The Select Role Services page of the Add Roles and Features Wizard.

9. Click Next.

A confirmation page appears.

10. Click Install.

The features you selected are installed. This may take a few minutes, so now would be a good time to take a walk.

When the installation finishes, an Installation Results page is displayed to verify that IIS was properly installed.

11. Click Close.

IIS is now installed and ready to use!

Understanding the Default Website

Initially, IIS is configured with a single website: the *default website*. You can test that IIS is up and running by opening a browser window on the server and typing **localhost** in the address bar. You can also reach this page by entering your local domain name in the address bar, such as **lowewriter.pri**. Figure 15-5 shows the standard welcome page that appears when you browse to the default site.

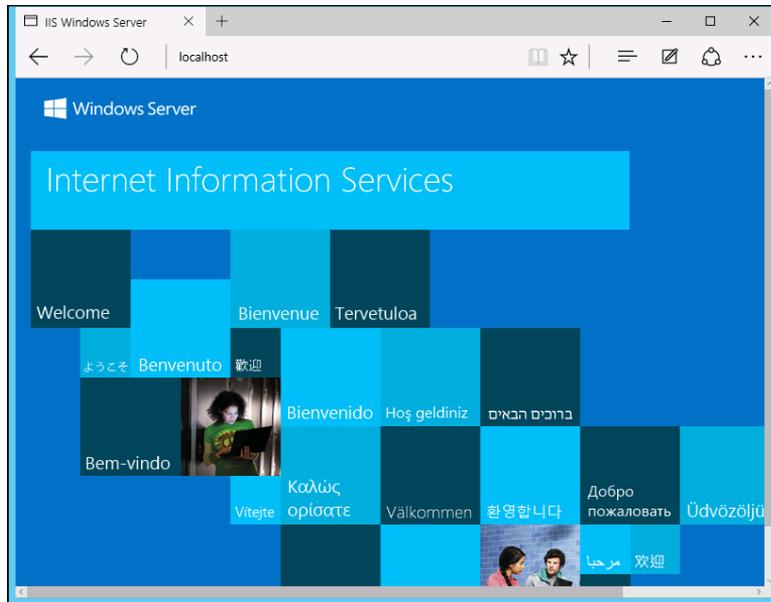


FIGURE 15-5:
The default website.

The actual files that make up the default website are stored on the server's C: drive in a folder named `\inetpub\wwwroot`. When you browse to the default website without requesting a specific file (simply by typing `localhost` in the address bar, for example), IIS looks for the following files, in this order:

- » default.htm
- » default.asp
- » index.htm
- » index.html
- » iisstart.htm
- » default.aspx

Initially, `c:\inetpub\wwwroot` contains just two files: `iisstart.htm` and `welcome.png`. The `iisstart.htm` file is the file that's displayed when you browse to the website; it contains the HTML markup necessary to display the image contained in the `welcome.png` file, which is the image you actually see on the page.

You can preempt the standard page for the default website by providing your own file one of the preceding names. You can follow these steps, for example, to create a simple `default.htm` file that displays the words *Hello World!* as the start page for the default website (you must be logged in as an administrator):

1. **Open an Explorer window, and browse to `c:\inetpub\wwwroot`.**
2. **Choose `File` → `New` → `Text Document`, type `default.htm` for the filename, and press `Enter`.**
3. **Right-click the `default.htm` file you just created and choose `Open With` → `Notepad` from the contextual menu.**
4. **Enter the following text in the Notepad window:**

```
<HTML>
<BODY>
<H1>Hello World!</H1>
</BODY>
</HTML>
```

5. **Choose `File` → `Save` to save the file and then choose `File` → `Exit` to quit Notepad.**
6. **Open a browser window.**
7. **Type `localhost` in the address bar, and press `Enter`.**

The page shown in Figure 15-6 appears.

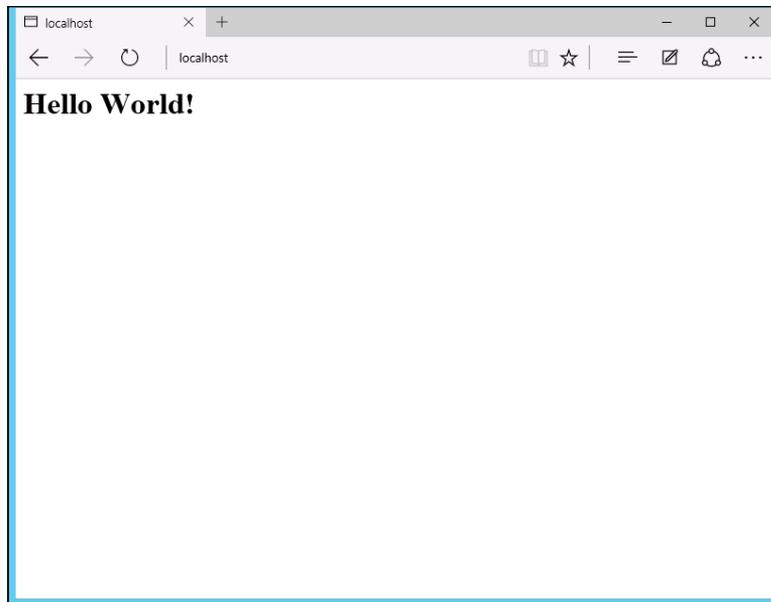


FIGURE 15-6:
Hello World!

Creating Websites

IIS has the ability to host multiple websites. This is an extremely useful feature not only for web servers that host public sites, but also for web servers that host internal (intranet) sites. You might create a separate intranet website for Human Resources and assign it the website name `hr`. Then, assuming that the domain name is `lowewriter.pri`, users can browse to the website by using the address `hr.lowewriter.pri`.

Here are the steps:

1. **Using Windows Explorer, create a folder in which you will save the files for the new website.**

For this example, I created a folder named `c:\HR-Web-Site`.

2. **In Server Manager, choose Tools ⇨ Internet Information Services (IIS) Manager.**

The IIS Manager springs to life, as shown in Figure 15-7.

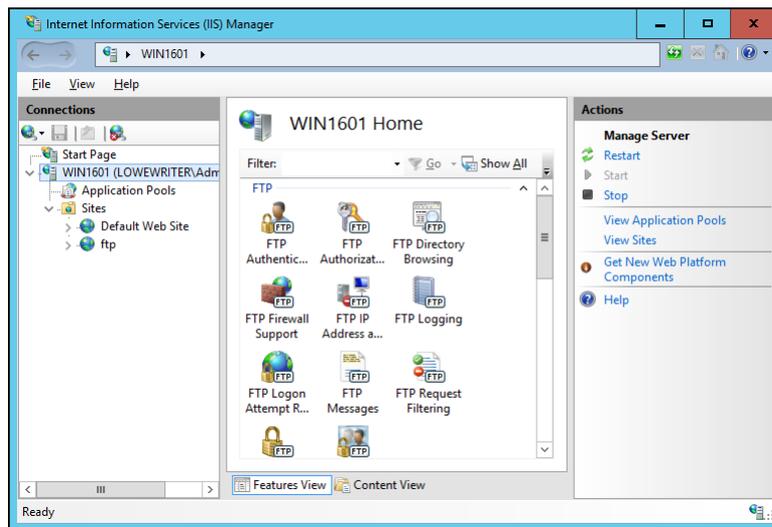


FIGURE 15-7:
The IIS Manager.

3. **Right-click Sites and then choose Add Website from the contextual menu.**

The Add Website dialog box appears, as shown in Figure 15-8.

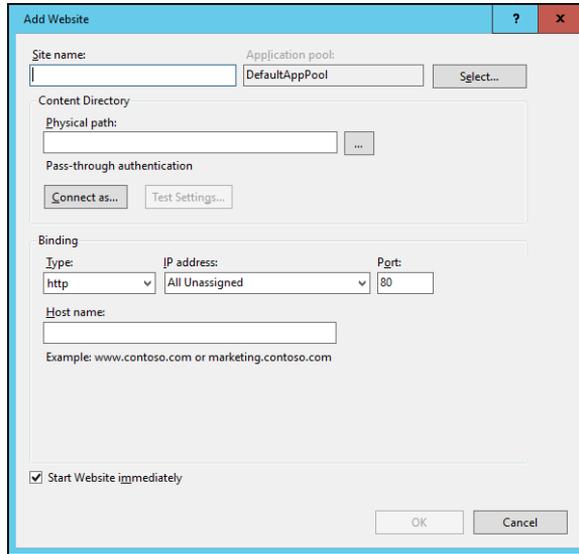


FIGURE 15-8:
The Add Website dialog box.

4. Enter a name for the website in the Site Name text box.

For this example, I used HR.

5. Click the Browse button (the one with the ellipsis), browse to the folder you created in Step 1, and then click OK.

For this example, I browsed to C:\HR-Web-Site.

6. In the Host Name text box, enter the exact DNS name you want to use for the site.

For this example, I entered **hr.lowewriter.pri**.

7. Click OK.

The newly created website appears below the Sites node in the IIS Manager, as shown in Figure 15-9.

8. Close the IIS Manager.

9. Create a web page to display in the folder you created in Step 1.

For this example, I used Notepad to create a text file named `default.htm`, with the following text:

```
<HTML>
<BODY>
<H1>Welcome to the HR Web Site!</H1>
</BODY>
</HTML>
```

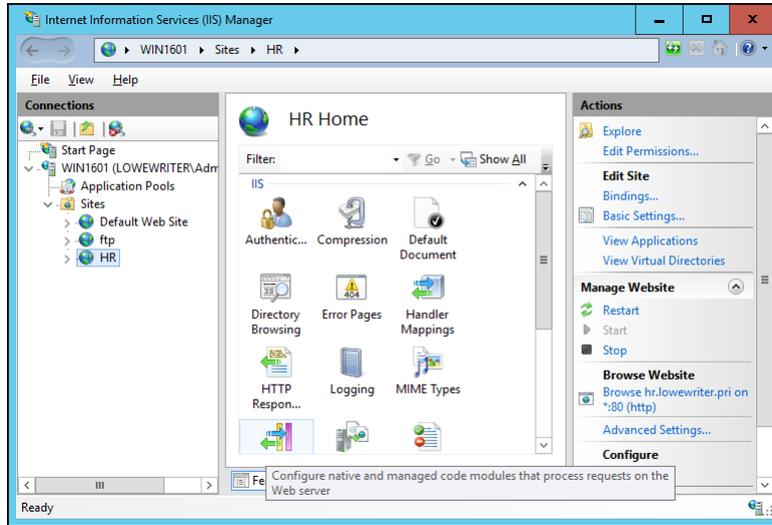


FIGURE 15-9:
The HR website
appears in the IIS
Manager.

10. In Server Manager, choose Tools ⇄ DNS.

This brings up the DNS Manager console, as shown in Figure 15-10.

11. In the navigation pane, navigate to the node for your domain.

In this example, I navigated to `lowewriter.pri`.

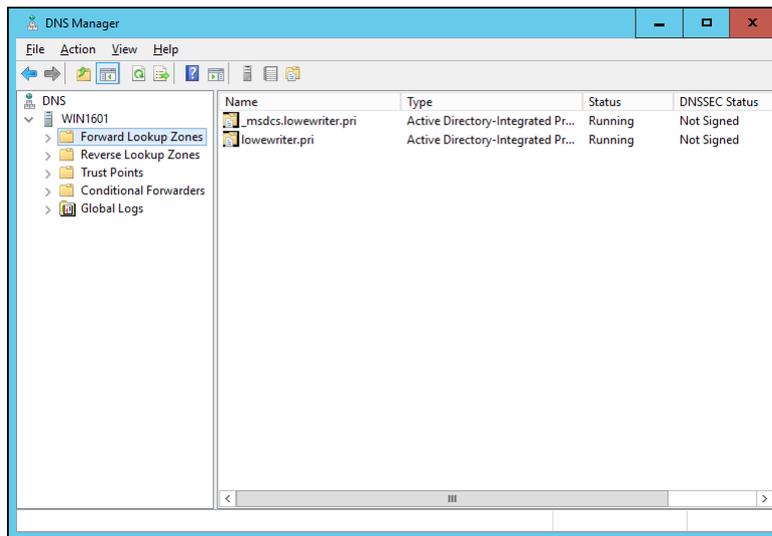


FIGURE 15-10:
The DNS Manager
console.

12. Choose Action ⇨ New Alias (CNAME).

The New Resource Record dialog box appears, as shown in Figure 15-11.

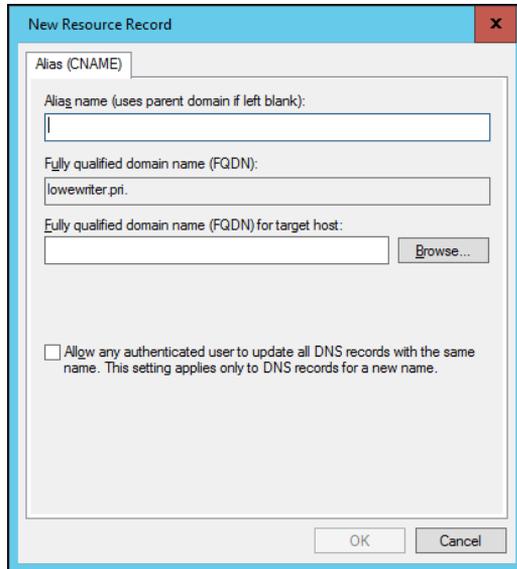


FIGURE 15-11:
Creating a
CNAME record.

13. Enter the alias name you want to use in the Alias Name text box.

For this example, I entered simply hr.

14. Enter the computer name of your web server in the Fully Qualified Domain Name (FQDN) for Target Host text box.

For this example, I entered lserver01.

15. Click OK.

The DNS alias is created.

16. Close the DNS Manager console.

17. Open a browser window.

18. Browse to the alias address you just created.

For this example, I browsed to `hr.lowewriter.pri`. Figure 15-12 shows the resulting page.

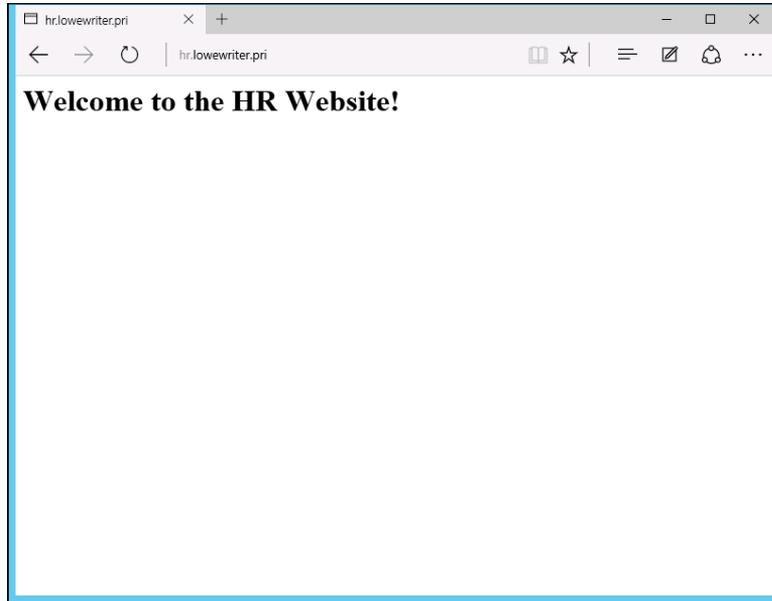


FIGURE 15-12:
Viewing a
website.

4 Managing and Protecting Your Network

IN THIS PART . . .

Learn what network management is all about.

Solve pesky networking problems.

Protect your network data by backing it up.

Secure your network by enforcing user account policies.

Make your network even more secure by installing a firewall and use antivirus software.

IN THIS CHAPTER

Exploring the many jobs of the network administrator

Documenting the network

Dusting, vacuuming, and mopping

Managing network users

Choosing the right tools

Building your library

Getting certified

Chapter 16

Welcome to Network Management

Help wanted. Network administrator to help small business get control of a network run amok. Must have sound organizational and management skills. Only moderate computer experience required. Part-time only.

Does this sound like an ad that your company should run? Every network needs a network administrator, whether the network has 2 computers or 2,000. Of course, managing a 2,000-computer network is a full-time job, whereas managing a 2-computer network isn't. At least, it shouldn't be.

This chapter introduces you to the boring and tedious job of network administration. Oops — you're probably reading this chapter because you've been elected to be the network manager, so I'd better rephrase that:

This chapter introduces you to the wonderful, exciting world of network management! Oh, boy! This is going to be fun!

What a Network Administrator Does

A network administrator “administers” a network: installing, configuring, expanding, protecting, upgrading, tuning, and repairing the network.

A network administrator takes care of the network hardware (such as cables, hubs, switches, routers, servers, and clients) and the network software (such as network operating systems, email servers, backup software, database servers, and application software). Most important, the administrator takes care of network users by answering their questions, listening to their troubles, and solving their problems.

On a big network, these responsibilities constitute a full-time job. Large networks tend to be volatile: Users come and go, equipment fails, software chokes, and life in general seems to be one crisis after another.

Smaller networks are much more stable. After you get your network up and running, you probably won’t have to spend much time managing its hardware and software. An occasional problem may pop up, but with only a few computers on the network, problems should be few and far between.

Regardless of the network’s size, the administrator attends to common chores:

- » **Get involved in every decision to purchase new computers, printers, or other equipment.**
 - » **Put on the pocket protector whenever a new computer is added to the network.** The network administrator’s job includes considering changes in the cabling configuration, assigning a computer name to the new computer, integrating the new user into the security system, and granting user rights.
 - » **Whenever a software vendor releases a new version of its software, read about the new version and decide whether its new features warrant an upgrade.** In most cases, the hardest part of upgrading to new software is determining the *migration path* — that is, upgrading your entire network to the new version while disrupting the network and its users as little as possible. This statement is especially true if the software in question happens to be your network operating system because any change to the network operating system can potentially impact the entire network.
- Between upgrades, software vendors periodically release patches and service packs that fix minor problems. For more information, see Chapter 20.
- » **Perform routine chores, such as backing up the servers, archiving old data, and freeing up server disk space.** Much of the task of network



REMEMBER

administration involves making sure that things keep working by finding and correcting problems before users notice that something is wrong. In this sense, network administration can be a thankless job.

- » **Gather, organize, and track the entire network's software inventory.** You never know when something will go haywire on the ancient Windows Vista computer that Joe in Marketing uses, and you have to reinstall that old copy of Lotus Approach. Do you have any idea where the installation disks are?

Choosing the Part-Time Administrator

The larger the network, the more technical support it needs. Most small networks — with just a dozen or so computers — can get by with a part-time network administrator. Ideally, this person should be a closet computer geek: someone who has a secret interest in computers but doesn't like to admit it. Someone who will take books home and read them over the weekend. Someone who enjoys solving computer problems just for the sake of solving them.

The job of managing a network definitely requires computer skills, but it isn't entirely a technical job. Much of the work that the network administrator does is routine housework. Basically, the network administrator dusts, vacuums, and mops the network periodically to keep it from becoming a mess.

Here are some additional ideas on picking a part-time network administrator:

- » The network administrator needs to be an organized person. Conduct a surprise office inspection and place the person with the neatest desk in charge of the network. (Don't warn them in advance, or everyone may mess up their desks intentionally the night before the inspection.)
- » Allow enough time for network administration. For a small network (say, no more than 20 or so computers), an hour or two each week is enough. More time is needed upfront as the network administrator settles into the job and discovers the ins and outs of the network. After an initial settling-in period, though, network administration for a small office network doesn't take more than an hour or two per week. (Of course, larger networks take more time to manage.)
- » Make sure that everyone knows who the network administrator is and that the network administrator has the authority to make decisions about the network, such as what access rights each user has, what files can and can't be stored on the server, how often backups are done, and so on.

- » Pick someone who is assertive and willing to irritate people. A good network administrator should make sure that backups are working *before* a hard drive fails and make sure that antivirus protection is in place *before* a virus wipes out the entire network. This policing will irritate people, but it's for their own good.
- » In most cases, the person who installs the network is also the network administrator. This is appropriate because no one understands the network better than the person who designs and installs it.
- » The network administrator needs an understudy — someone who knows almost as much about the network, is eager to make a mark, and smiles when the worst network jobs are delegated.
- » The network administrator has some sort of official title, such as Network Boss, Network Czar, Vice President in Charge of Network Operations, or Dr. Network. A badge, a personalized pocket protector, or a set of Spock ears helps, too.

The Three “Ups” of Network Management

Much of the network manager's job is routine stuff — the equivalent of vacuuming, dusting, and mopping, or changing your car's oil and rotating the tires.

Three of the most important routine tasks that a network administrator must do vigilantly are what I call the “Three Ups of Network Management.” They are

- » **Back up:** The network manager must ensure that the network is properly backed up. If something goes wrong and the network isn't backed up, guess who gets the blame? On the other hand, if disaster strikes yet you're able to recover everything from yesterday's backup with only a small amount of work lost, who gets the pat on the back, the fat bonus, and the vacation in the Bahamas? Chapter 18 describes the options for network backups. Read it *soon*.
- » **Lock-up:** Another major task for a network administrator is sheltering the network from the evils of the outside world. These evils come in many forms, including hackers trying to break into your network and virus programs arriving through email or untrustworthy websites. Chapter 19 describes this task in more detail.

- » **Clean-up:** Users think that the network server is like the attic: They want to throw files up there and leave them forever. No matter how much disk storage your network has, your users will fill it up sooner than you think, so the network manager gets the fun job of cleaning up the attic once in a while. The best advice I can offer is to continually complain about how messy it is up there and warn your users that spring cleaning is on the to-do list.

Managing Network Users

Managing network technology is the easiest part of network management. Computer technology can be confusing at first, but computers aren't as confusing as people. The real challenge of managing a network is managing the network's users.

The difference between managing technology and managing users is obvious: You can figure out computers, but who can ever really figure out people? The people who use the network are much less predictable than the network itself. Here are some tips for dealing with users:

- » **Make user training a key part of the network manager's job.** Make sure that everyone who uses the network understands how it works and how to use it. If the network users don't understand how the network works, they may unintentionally do all kinds of weird things to it.
- » **Treat network users respectfully.** If users don't understand how to use the network, it's not their fault. Explain it to them. Offer a class. Buy each one a copy of this book, and tell them to read it during the lunch hour. Hold their hands. Just don't treat them like idiots.
- » **Create a network cheat sheet.** It should contain everything users need to know about using the network — on one page. Everyone needs a copy.
- » **Be as responsive as possible.** If you don't quickly fix a network user's problem, he may try to fix it. You don't want that to happen.



TIP

The better you understand the psychology of network users, the more prepared you are for the strangeness they often serve up. Toward that end, I recommend that you read the *Diagnostic and Statistical Manual of Mental Disorders* (also known as *DSM-5*) from cover to cover.

Acquiring Software Tools for Network Administrators

Network managers need certain tools to get their jobs done. Managers of big, complicated, expensive networks need big, complicated, expensive tools. Managers of small networks need small tools.

Some tools that a manager needs are hardware tools, such as screwdrivers, cable crimpers, and hammers. The tools I'm talking about, however, are software tools. I mention a couple of them earlier in this chapter: Visio (to help you draw network diagrams) and a network-discovery tool to help you map your network. Here are a few others:

» **Built-in TCP/IP commands:** Many of the software tools that you need in order to manage a network come with the network itself. As the network manager, you should read through the manuals that come with your network software to see which management tools are available. For example, Windows includes a `net diag` command that you can use to make sure that all the computers on a network can communicate with each other. (You can run `net diag` from an MS-DOS prompt.) For TCP/IP networks, you can use the TCP/IP diagnostic commands that I summarize in Table 16-1.

TABLE 16-1 TCP/IP Diagnostic Commands

Command	What It Displays
<code>arp</code>	Address resolution information used by the Address Resolution Protocol (ARP)
<code>hostname</code>	Your computer's host name
<code>ipconfig</code>	Current TCP/IP settings
<code>nbtstat</code>	The status of NetBIOS over TCP/IP connections
<code>netstat</code>	Statistics for TCP/IP
<code>nslookup</code>	DNS information
<code>ping</code>	Verification that a specified computer can be reached
<code>route</code>	The PC's routing tables
<code>tracert</code>	The route from your computer to a specified host

- » **System Information:** This program, which comes with Windows, is a useful utility for network managers.
- » **Hotfix Checker:** This handy tool from Microsoft scans your computers to see which patches need to be applied. You can download the Hotfix Checker for free from the Microsoft website. Just go to www.microsoft.com and search for **hotfix**.
- » **Baseline Security Analyzer:** If you prefer GUI-based tools, check out this program, which you can download for free from the Microsoft website. To find it, go to www.microsoft.com and search for **Microsoft Baseline Security Analyzer**.
- » **Protocol analyzer:** A *protocol analyzer* (or *packet sniffer*) can monitor and log the individual packets that travel along your network. You can configure the protocol analyzer to filter specific types of packets, watch for specific types of problems, and provide statistical analysis of the captured packets.

Most network administrators agree that Sniffer, by NetScout Systems, Inc. (www.netscout.com) is the best protocol analyzer available. However, it's also one of the most expensive. If you prefer a free alternative, check out *Wireshark*, which you can download from www.wireshark.org.
- » **Network Monitor:** All current versions of Windows include a program called Network Monitor, which provides basic protocol analysis and can often help solve pesky network problems.



TIP

Building a Library

Scotty delivered one of his best lines in the original *Star Trek* series when he refused to take shore leave so that he could get caught up on his technical journals. “Don’t you ever relax?” asked Kirk. “I am relaxing!” Scotty replied.

To be a good network administrator, you need to read computer books — lots of them. And you need to enjoy doing it. If you’re the type who takes computer books with you to the beach, you’ll make a great network administrator.

Read books on a variety of topics. I don’t recommend specific titles, but I do recommend that you get a good, comprehensive book on each of these topics:

- » Network security and hacking
- » Wireless networking
- » Network cabling and hardware

- » Ethernet
- » Windows Server 2008, 2012, and 2016
- » Windows 7, 8, 8.1, and 10
- » Linux
- » TCP/IP
- » DNS
- » Sendmail or Microsoft Exchange Server, depending on which email server you use

In addition to books, you may also want to subscribe to some magazines to keep up with what's happening in the networking industry. Here are a few you should probably consider, along with their web addresses:

- » *InformationWeek*: www.informationweek.com
- » *InfoWorld*: www.infoworld.com
- » *Network Computing*: www.networkcomputing.com
- » *Network World*: www.networkworld.com
- » *2600 The Hacker Quarterly* (a great magazine on computer hacking and security): www.2600.com



TIP

The Internet is one of the best sources of technical information for network administrators. You'll want to stock your browser's Favorites menu with plenty of websites that contain useful networking information. In addition, you may want to subscribe to one of the many online newsletters that deliver fresh information on a regular basis via email.

Pursuing Certification

Remember the scene near the end of *The Wizard of Oz* when the Wizard grants the Scarecrow a diploma, the Cowardly Lion a medal, and the Tin Man a testimonial?

Network certifications are kind of like that. I can picture the scene now:

The Wizard: "And as for you, my network-burdened friend, any geek with thick glasses can administer a network. Back where I come from, there are people who do nothing but configure Cisco routers all day long. And they don't have any more

brains than you do. But they do have one thing you don't have: certification. And so, by the authority vested in me by the Universita Committeeatum E Pluribus Unum, I hereby confer upon you the coveted certification of CND."

You: "CND?"

The Wizard: "Yes, that's, uh, *Certified Network Dummy.*"

You: "The Seven Layers of the OSI Reference Model are equal to the Sum of the Layers on the Opposite Side. Oh, joy, rapture! I feel like a network administrator already!"

My point is that certification in and of itself doesn't guarantee that you really know how to administer a network. That ability comes from real-world experience — not exam crams.

Nevertheless, certification is becoming increasingly important in today's competitive job market. So you may want to pursue certification, not just to improve your skills, but also to improve your resume. Certification is an expensive proposition. Each test can cost several hundred dollars, and depending on your technical skills, you may need to buy books to study or enroll in training courses before you take the tests.

You can pursue two basic types of certification: vendor-specific certification and vendor-neutral certification. The major software vendors such as Microsoft and Cisco provide certification programs for their own equipment and software. CompTIA, a nonprofit industry trade association, provides the best-known vendor-neutral certification.

Helpful Bluffs and Excuses

As network administrator, you just won't be able to solve a problem sometimes, at least not immediately. You can do two things in this situation. The first is to explain that the problem is particularly difficult and that you'll have a solution as soon as possible. The second solution is to look the user in the eyes and, with a straight face, try one of these phony explanations:

- » Blame it on the version of whatever software you're using. "Oh, they fixed that with version 39."
- » Blame it on cheap, imported memory chips.
- » Blame it on Democrats. Or Republicans. Doesn't matter.

- » Blame it on oil company executives.
- » Blame it on the drought.
- » Blame it on the rain.
- » Blame it on climate change.
- » Hope that the problem wasn't caused by stray static electricity. Those types of problems are very difficult to track down. Tell your users that not properly discharging themselves before using their computers can cause all kinds of problems.
- » You need more memory.
- » You need a bigger hard drive.
- » You need a faster processor.
- » Blame it on Jar-Jar Binks.
- » You can't do that in Windows 10.
- » You can only do that in Windows 10.
- » Could be a virus.
- » Or sunspots.
- » No beer and no TV make Homer something something something.

IN THIS CHAPTER

Checking the obvious things

Fixing computers that have expired

Pinpointing the cause of trouble

Restarting client and server computers

Reviewing network event logs

Keeping a record of network woes

Chapter 17

Solving Network Problems

Face it: Networks are prone to breaking.

They have too many parts. Cables. Connectors. Cards. Switches. Routers. All these parts must be held together in a delicate balance, and the network equilibrium is all too easy to disturb. Even the best-designed computer networks sometimes act as if they're held together with baling wire, chewing gum, and duct tape.

To make matters worse, networks breed suspicion. After your computer is attached to a network, users begin to blame the network every time something goes wrong, regardless of whether the problem has anything to do with the network. You can't get columns to line up in a Word document? Must be the network. Your spreadsheet doesn't add up? The @#\$% network's acting up again. The stock market's down? Arghhh!!!!!!

The worst thing about network failures is that sometimes they can shut down an entire company. It's not so bad if just one user can't access a particular shared

folder on a file server. If a critical server goes down, however, your network users may be locked out of their files, applications, email, and everything else they need to conduct business as usual. When that happens, they'll be beating down your doors and won't stop until you get the network back up and running.

In this chapter, I review some of the most likely causes of network trouble and suggest some basic troubleshooting techniques that you can employ when your network goes on the fritz.

When Bad Things Happen to Good Computers

Here are some basic troubleshooting steps explaining what you should examine at the first sign of network trouble. In many (if not most) of the cases, one of the following steps can get your network back up and running:

- 1. Make sure that your computer and everything attached to it is plugged in.**



Computer geeks love it when a user calls for help, and they get to tell the user that the computer isn't plugged in or that its power strip is turned off. They write it down in their geek logs so that they can tell their geek friends about it later. They may even want to take your picture so that they can show it to their geek friends. (Most "accidents" involving computer geeks are a direct result of this kind of behavior. So try to be tactful when you ask a user whether he or she is sure the computer is actually turned on.)

- 2. Make sure that your computer is properly connected to the network.**
- 3. Note any error messages that appear on-screen.**
- 4. Try restarting the computer.**



An amazing number of computer problems are cleared up by a simple restart of the computer. Of course, in many cases, the problem recurs, so you'll have to eventually isolate the cause and fix the problem. Some problems are only intermittent, and a simple reboot is all that's needed.

- 5. Try the built-in Windows network troubleshooter.**

For more information, see the section, "Using the Windows Networking Troubleshooter," later in this chapter.

6. Check the free disk space on your computer and on the server.

When a computer runs out of disk space or comes close to it, strange things can happen. Sometimes you get a clear error message indicating such a situation, but not always. Sometimes the computer just grinds to a halt; operations that used to take a few seconds now take a few minutes.

7. Do a little experimenting to find out whether the problem is indeed a network problem or just a problem with the computer itself.

See the section, “Time to Experiment,” later in this chapter, for some simple things that you can do to isolate a network problem.

8. Try restarting the network server.

See the section, “Restarting a Network Server,” later in this chapter.

Fixing Dead Computers

If a computer seems totally dead, here are some things to check:

- » **Make sure that the computer is plugged in.**
- » **If the computer is plugged into a surge protector or a power strip, make sure that the surge protector or power strip is plugged in and turned on.** If the surge protector or power strip has a light, it should be glowing. Also, the surge protector may have a reset button that needs to be pressed.
- » **Make sure that the computer’s On/Off switch is turned on.** This advice sounds too basic to even include here, but many computers have two power switches: an on/off switch on the back of the computer and a push-button on the front that actually starts the computer. If you push the front button and nothing happens, check the switch on the back to make sure it’s in the ON position.



REMEMBER

To complicate matters, newer computers have a Sleep feature, in which they appear to be turned off but really they’re just sleeping. All you have to do to wake such a computer is jiggle the mouse a little. (I used to have an uncle like that.) It’s easy to assume that the computer is turned off, press the power button, wonder why nothing happened, and then press the power button and hold it down, hoping it will take. If you hold down the power button long enough, the computer will actually turn itself off. Then, when you turn the computer back on, you get a message saying the computer wasn’t shut down properly. Arghhh! The moral of the story is to jiggle the mouse if the computer seems to have nodded off.

- » **If you think the computer isn't plugged in but it looks like it is, listen for the fan.** If the fan is running, the computer is getting power, and the problem is more serious than an unplugged power cord. (If the fan isn't running but the computer is plugged in and the power is on, the fan may be out to lunch. Also, if the computer is making an annoying squealing sound, the problem is most likely a bad fan. Fortunately, fans are inexpensive and easy to replace.)
- » **If the computer is plugged in and turned on but still not running, plug a lamp into the outlet to make sure that power is getting to the outlet.** You may need to reset a tripped circuit breaker or replace a bad surge protector. Or you may need to call the power company. (If you live in California, don't bother calling the power company. It probably won't do any good.)
- » **Check the surge protector.** Surge protectors have a limited life span. After a few years of use, many surge protectors continue to provide electrical power for your computer, but the components that protect your computer from power surges no longer work. If you're using a surge protector that is more than two or three years old, replace the old surge protector with a new one.
- » **Make sure that the monitor is plugged in and turned on.** The monitor has a separate power cord and switch. (The monitor actually has two cables that must be plugged in. One runs from the back of the monitor to the back of the computer; the other is a power cord that comes from the back of the monitor and must be plugged into an electrical outlet.)
- » **Make sure that all cables are plugged in securely.** Your keyboard, monitor, mouse, and printer are all connected to the back of your computer by cables. Make sure that the other ends of the monitor and printer cables are plugged in properly, too.
- » **If the computer is running but the display is dark, try adjusting the monitor's contrast and brightness.** Some monitors have knobs that you can use to adjust the contrast and brightness of the monitor's display. They may have been turned down all the way.



REMEMBER

Ways to Check a Network Connection

The cables that connect client computers to the rest of the network are finicky beasts. They can break at a moment's notice, and by "break," I don't necessarily mean "to physically break in two." Although some broken cables look like someone got to the cable with pruning shears, most cable problems aren't visible to the naked eye.



TIP

» **Twisted-pair cable:** If your network uses twisted-pair cable, you can quickly tell whether the cable connection to the network is good by looking at the back of your computer. Look for a small light located near where the cable plugs in; if this light is glowing steadily, the cable is good. If the light is dark or it's flashing intermittently, you have a cable problem (or a problem with the network card or the hub or switch that the other end of the cable is plugged in to).

If the light isn't glowing steadily, try removing the cable from your computer and reinserting it. This action may cure the weak connection.

» **Patch cable:** Hopefully, your network is wired so that each computer is connected to the network with a short (six feet or so) patch cable. One end of the patch cable plugs into the computer, and the other end plugs into a cable connector mounted on the wall. Try quickly disconnecting and reconnecting the patch cable. If that doesn't do the trick, try to find a spare patch cable that you can use.

» **Switches:** Switches are prone to having cable problems, too — especially switches that are wired in a “professional manner,” involving a rat's nest of patch cables. Be careful whenever you enter the lair of the rat's nest. If you need to replace a patch cable, be very careful when you disconnect the suspected bad cable and reconnect the good cable in its place.

A Bunch of Error Messages Just Flew By!

Error messages that display when your computer boots can provide invaluable clues to determine the source of the problem.

If you see error messages when you start up the computer, keep the following points in mind:

» **Don't panic if you see a lot of error messages.** Sometimes, a simple problem that's easy to correct can cause a plethora of error messages when you start your computer. The messages may look as if your computer is falling to pieces, but the fix may be very simple.

» **If the messages fly by so fast that you can't see them, press your computer's Pause key.** Your computer comes to a screeching halt, giving you a chance to catch up on your error-message reading. After you've read enough, press the Pause key again to get things moving. (On keyboards that don't have a Pause key, pressing Ctrl+Num Lock or Ctrl+S does the same thing.)



TIP

- » If you miss the error messages the first time, restart the computer and watch them again.
- » Better yet, press F8 when you see the Starting Windows message. This displays a menu that allows you to select from several startup options. (Note that this won't work on Windows 8 or later.)

Double-Checking Your Network Settings

I swear that there are little green men who sneak into offices at night, turn on computers, and mess up TCP/IP configuration settings just for kicks. These little green men are affectionately known as *networchons*.

Remarkably, network configuration settings sometimes get inadvertently changed so that a computer, which enjoyed the network for months or even years, one day finds itself unable to access the network. So one of the first things you do, after making sure that the computers are actually on and that the cables aren't broken, is a basic review of the computer's network settings. Check the following:

- » At a command prompt, run `ipconfig` to make sure that TCP/IP is up and running on the computer and that the IP addresses, subnet masks, and default gateway settings look right.
- » Call up the network connection's Properties dialog box and make sure that the necessary protocols are installed correctly.
- » Open the System Properties dialog box (double-click System in Control Panel) and check the Computer Name tab.

Make sure that the computer name is unique and also that the domain or workgroup name is spelled properly.

- » Double-check the user account to make sure that the user really has permission to access the resources he needs.

Using a Windows Troubleshooter

Windows comes with several built-in troubleshooters that can often help you to pin down the cause of a network problem. These troubleshooters are organized into four broad categories:

- » **Programs:** Solve compatibility issues with older programs.
- » **Hardware and Sound:** Solve issues such as hardware compatibility problems, incorrect device configuration, printing problems, and audio playback difficulties.
- » **Network and Internet:** Solve problems connecting to the Internet or working with shared files and folders. Figure 17-1 shows the Network Adapter Troubleshooter, which helps you solve problems with a network adapter.
- » **System and Security:** Fix problems related to Windows Update or other system-related problems.

The Windows Troubleshooters work by asking questions and making suggestions for fixes, and often offers to fix incorrectly configured devices for you. Answer the questions asked by the troubleshooter and click Next to move from screen to screen. The troubleshooters can't solve all of your problems, but they can point out the causes of the most common problems.

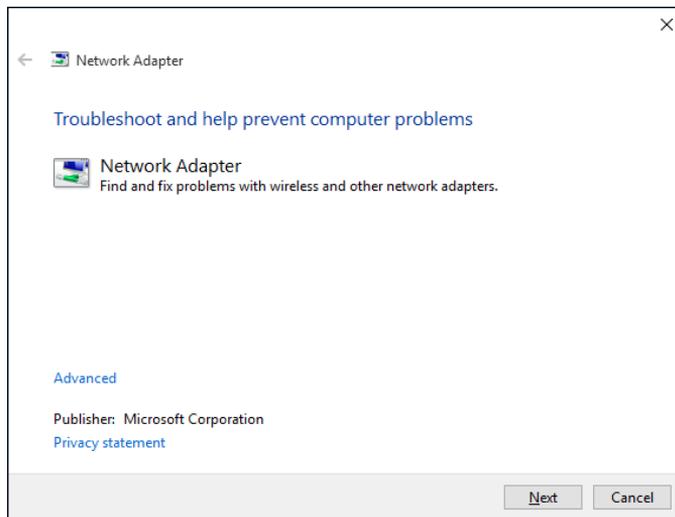


FIGURE 17-1:
The Windows 10
Network Adapter
Troubleshooter.

Here are the steps for starting a troubleshooter for Windows versions 8 and later:

- 1. Open the Control Panel.**
- 2. Click Troubleshooting.**
This displays a list of the troubleshooting categories.
- 3. Click the category you want to troubleshoot.**
A list of helpful troubleshooters is displayed.

4. **Right-click the troubleshooter that seems most directly related to the problem you're having and choose Run As Administrator.**

Running the troubleshooter with Administrator privileges will help it better find the source of your problems.

Time to Experiment

If you can't find some obvious explanation for your troubles — say, the computer is unplugged — you need to do some experimenting to narrow down the possibilities. Design your experiments to answer one basic question: Is it a network problem or a local computer problem?

Here are some ways you can narrow down the cause of the problem:

- » **Try performing the same operation on someone else's computer.** If no one on the network can access a network drive or printer, something is probably wrong with the network. On the other hand, if the error occurs on only one computer, the problem is likely with that computer. The wayward computer may not be reliably communicating with the network or configured properly for the network, or the problem may have nothing to do with the network at all.
- » **If you can perform the operation on another computer without problems, try logging on to the network with another computer using your own username.** Then see whether you can perform the operation without error. If you can, the problem is probably on your computer. If you can't, the problem may be with the way your user account is configured.
- » **If you can't log on at another computer, try waiting for a bit.** Your account may be temporarily locked out. This can happen for a variety of reasons — the most common of which is trying to log on with the wrong password several times in a row. If you're still locked out an hour later, call the network administrator and offer a doughnut.

Who's on First?

When troubleshooting a networking problem, find out who is actually logged on to a network server. For example, if a user can't access a file on the server, you can check whether the user is logged on. If so, you know that the user's account is valid. Even so, the user may not have permission to access the particular file or folder that

he's attempting to access. On the other hand, if the user isn't logged on, the problem may lie with the account itself or how the user is attempting to connect to the server.

I also recommend finding out who's logged on in case you need to restart the server. For more information about restarting a server, see the section, "Restarting a Network Server," later in this chapter.

To find out who is currently logged on to a Windows server, right-click the Computer icon on the desktop and choose Manage from the menu that appears. This brings up the Computer Management window. Open System Tools in the tree list and then open Shared Folders and select Sessions. A list of users who are logged on appears.



TIP

You can immediately disconnect all users by right-clicking Sessions in the Computer Management window and choosing All Tasks ⇨ Disconnect All. Be warned, however, that this can cause users to lose data.

Restarting a Client Computer

Sometimes, trouble gets a computer so tied up in knots that the only thing you can do is reboot. In some cases, the computer just starts acting weird. Strange characters appear on the screen, or Windows goes haywire and doesn't let you exit a program. Sometimes, the computer gets so confused that it can't even move. It just sits there, like a deer staring at oncoming headlights. It won't move, no matter how hard you press Esc or Enter. You can move the mouse all over your desktop, or you can even throw it across the room, but the mouse pointer on the screen stays perfectly still.

When a computer starts acting strange, you need to reboot. If you must reboot, you should do so as cleanly as possible. I know this procedure may seem elementary, but the technique for safely restarting a client computer is worth repeating, even if it is basic:

1. Save your work if you can.

Use the File ⇨ Save command to save any documents or files that you were editing when things started to go haywire. If you can't use the menus, try clicking the Save button on the toolbar. If that doesn't work, try pressing Ctrl+S (the standard keyboard shortcut for the Save command).

2. Close any running programs if you can.

Use the File ⇨ Exit command or click the Close button in the upper-right corner of the program window. Or press Alt+F4.

If a program refuses to close, you can usually shut it down by using Windows Task Manager. Right-click the Windows task bar and choose Start Task Manager. Then select the program you want to close and click the End Task button.

3. Restart the computer.

- *Windows 7:* Click the Start button, click the right arrow that appears at the bottom-right corner of the Start menu, and then click Restart.
- *Windows 8:* Oddly enough, shutting down Windows 8 is a bit challenging. You can stare at the desktop all day and not find an intuitive way to shut down your computer. The secret lies in the Charms Bar, which you can find by hovering the mouse over the lower-right corner of the screen. Next, click the Settings icon, and then click the Shut Down icon.
- *Windows 8.1 and 10:* Mercifully, with Windows 8.1, Microsoft re-introduced sanity into some of the most basic parts of Windows. Shutting down is one of them: To shut down a Windows 8.1 or Windows 10 computer, click the Start button, choose Power, and then choose Shut Down.

If restarting your computer doesn't seem to fix the problem, you may need to turn off your computer and then turn it on again. To do so, follow the previous procedure but choose Shut Down instead of Restart.

Here are a few things to try if you have trouble restarting your computer:

1. If your computer refuses to respond to the Start ⇄ Shut Down command, try pressing Ctrl+Alt+Delete.

This is called the “three-finger salute.” It's appropriate to say, “Queueue” while you do it.

When you press Ctrl+Alt+Delete, Windows displays a dialog box that enables you to close any running programs or shut down your computer entirely.

2. If pressing Ctrl+Alt+Delete doesn't do anything, you've reached the last resort. The only thing left to do is turn off the computer by pressing the power On/Off button and holding it down for a few seconds.



WARNING

Turning off your computer by pressing the power button is a drastic action that you should take only after your computer becomes completely unresponsive. Any work you haven't yet saved to disk is lost. (Sniff.) (If your computer doesn't have a Reset button, turn off the computer, wait a few moments, and then turn the computer back on again.)



REMEMBER

If at all possible, save your work before restarting your computer. Any work you haven't saved is lost. Unfortunately, if your computer is totally tied up in knots, you probably can't save your work. In that case, you have no choice but to push your computer off the digital cliff.

Booting in Safe Mode

Windows provides a special start-up mode called *Safe Mode* that's designed to help fix misbehaving computers. When you start your computer in Safe Mode, Windows loads only the most essential parts of itself into memory — the bare minimum required for Windows to work. Safe Mode is especially useful when your computer has developed a problem that prevents you from using the computer at all.

To boot into Safe Mode on a Windows 7 or earlier computer, first restart the computer. Then, as soon as the computer begins to restart, start pressing the F8 key — just tap away at it until a menu titled Advanced Boot Options appears. One of the options on this menu is Safe Mode; use the up- or down-arrow keys to select that option, and then press Enter to boot in Safe Mode.

On a Windows 8, 8.1, or 10 computer, you can reboot into Safe Mode by holding down the Shift key when you choose the Restart command.

Using System Restore

System Restore is a Windows feature that periodically saves important Windows configuration information and allows you to later return your system to a previously saved configuration. This can often fix problems by reverting your computer to a time when it was working.

By default, Windows saves *restore points* whenever you install new software on your computer or apply a system update. Restore points are also saved automatically every seven days.

Although System Restore is turned on by default, you should verify that System Restore is active and running to make sure that System Restore points are being created. To do that, right-click Computer from the Start menu, choose Properties, and then click the System Protection tab. The dialog box shown in Figure 17-2 is displayed. Verify that the Protection status for your computer's C: drive is On. If it isn't, select the C: drive and click the Configure button to configure System Restore for the drive.

If your computer develops a problem, you can restore it to a previously saved restore point by clicking the System Protection tab. This brings up the System Restore Wizard, as shown in Figure 17-3. This wizard allows you to select the restore point you want to use.

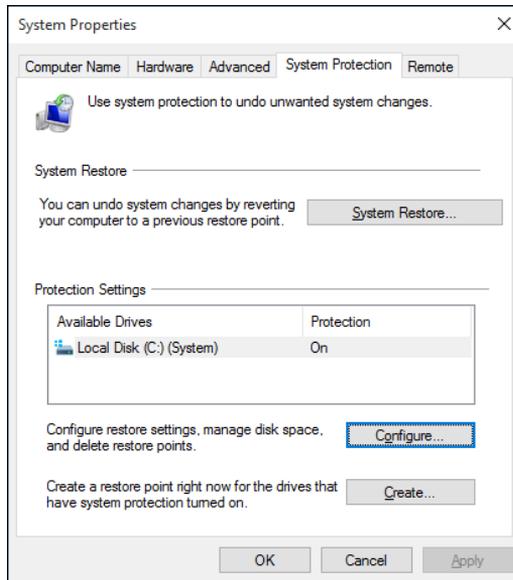


FIGURE 17-2:
The System Protection tab of the System Properties dialog box.

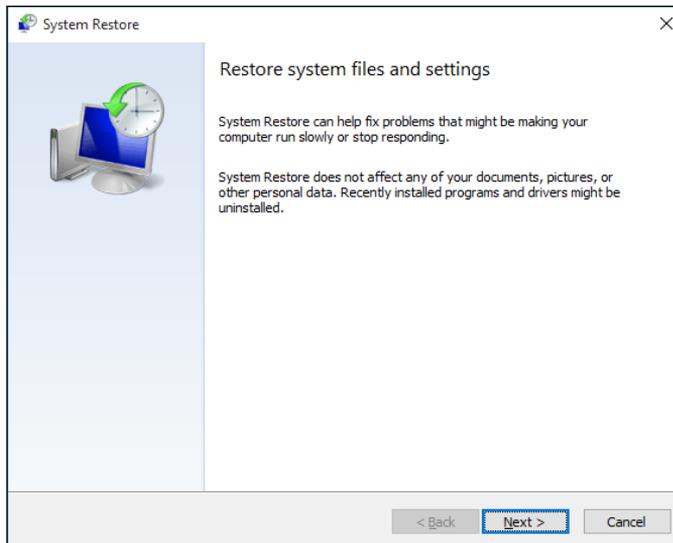


FIGURE 17-3:
Use System Restore to restore your system to an earlier configuration.

Here are a few additional thoughts to remember about System Restore:

- » System Restore *does not* delete data files from your system. Thus, files in your Documents folder won't be lost.
- » System Restore *does* remove any applications or system updates you've installed since the restore point was made. Thus, you need to reinstall those



WARNING

applications or system updates — unless, of course, you determine that an application or system update was the cause of your problem in the first place.

- » System Restore automatically restarts your computer. The restart may be slow because some of the changes made by System Restore happen after the restart.
- » Do *not* turn off or cut power to your computer during System Restore. Doing so may leave your computer in an unrecoverable state.

Restarting Network Services

Once in a while, the operating system (OS) service that supports the task that's causing you trouble inexplicably stops or gets stuck. If users can't access a server, it may be because one of the key network services has stopped or is stuck.

You can review the status of services by using the Services tool, the Windows 10 version of which is shown in Figure 17-4. To display it, click the Start button and type **Services**. Then choose the Services desktop app. Review the list of services to make sure that all key services are running. If an important service is paused or stopped, restart it.

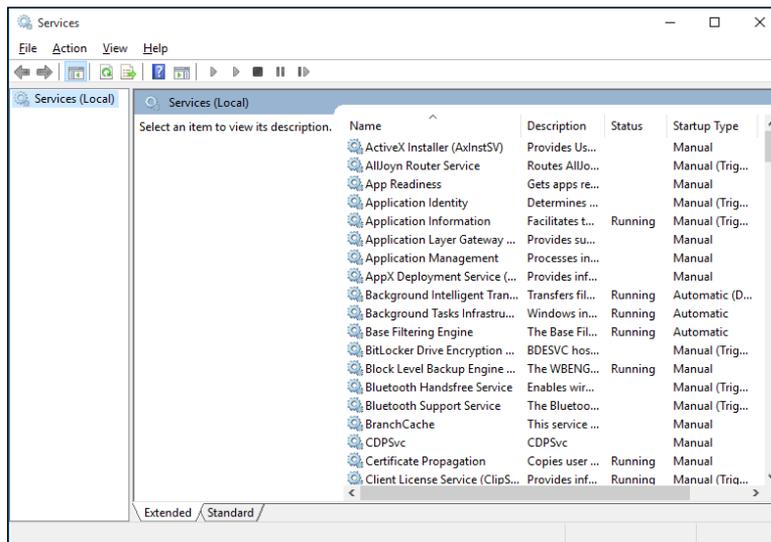


FIGURE 17-4:
Looking at which
services are
running.

Which services qualify as “important” depends on what roles you define for the server. Table 17-1 lists a few important services that are common to most Windows network operating systems. However, many servers require additional services besides these. In fact, a typical server will have many dozens of services running simultaneously.

TABLE 17-1 Key Windows Services

Service	Description
Computer Browser	Maintains a list of computers on the network that can be accessed. If this service is disabled, the computer won't be able to use browsing services, such as My Network Places.
DHCP Client	Enables the computer to obtain its IP address from a Dynamic Host Configuration Protocol (DHCP) server. If this service is disabled, the computer's IP address won't be configured properly.
DNS Client	Enables the computer to access a Domain Name Server (DNS) server to resolve DNS names. If this service is disabled, the computer won't be able to handle DNS names, including Internet addresses and Active Directory (AD) names.
Server	Provides basic file- and printer-sharing services for the server. If this service is stopped, clients won't be able to connect to the server to access files or printers.
Workstation	Enables the computer to establish client connections with other servers. If this service is disabled, the computer won't be able to connect to other servers.



WARNING

Key services usually stop for a reason, so simply restarting a stopped service probably won't solve your network's problem — at least, not for long. You should review the System log to look for any error messages that may explain why the service stopped in the first place.

Restarting a Network Server

Sometimes, the only way to flush out a network problem is to restart the network server that's experiencing trouble.



WARNING

Restarting a network server is something you should do only as a last resort. Windows Server is designed to run for months or even years at a time without rebooting. Restarting a server invariably results in a temporary shutdown of the network. If you must restart a server, try to do it during off hours if possible.



TIP

Before you restart a server, check whether a specific service that's required has been paused or stopped. You may be able to just restart the individual service rather than the entire server. For more information, see the section, "Restarting Network Services," earlier in this chapter.

Here's the basic procedure for restarting a network server:

1. Make sure that everyone is logged off the server.

The easiest way to do that is to restart the server after normal business hours, when everyone has gone home for the day. Then, you can just shut down the server and let the shutdown process forcibly log off any remaining users.

To find out who's logged on, refer to the earlier section, "Who's on First?"

2. After you're sure the users have logged off, shut down the network server.

You want to do this step behaving like a good citizen if possible — decently, and in order. Use the Start ⇨ Shut Down command to shut down the server. This summons a dialog box that requires you to indicate the reason for the shutdown. The information you supply here is entered into the server's System log, which you can review by using Event Viewer. See the next section for more on Event Viewer.

3. Reboot the server computer or turn it off and then on again.

Watch the server start up to make sure that no error messages appear.

4. Tell everyone to log back on and make sure that everyone can now access the network.



WARNING

Heed the following when you consider restarting the network server:

- » **Restarting the network server is more drastic than restarting a client computer.** Make sure that everyone saves his or her work and logs off the network before you do it! You can cause major problems if you blindly turn off the server computer while users are logged on.
- » **Obviously, restarting a network server is a major inconvenience to every network user.** Better offer treats.

Looking at Event Logs

One of the most useful troubleshooting techniques for diagnosing network problems is to review the network operating system's built-in event logs. These logs contain information about interesting and potentially troublesome events that

occur during the daily operation of your network. Ordinarily, these logs run in the background, quietly gathering information about network events. When something goes wrong, you can check the logs to see whether the problem generated a noteworthy event. In many cases, the event logs contain an entry that pinpoints the exact cause of the problem and suggests a solution.

To display the event logs in a Windows server, use Event Viewer, which is available from the Administrative Tools menu. For example, Figure 17-5 shows an Event Viewer from a Windows Server 2016 system. The tree listing on the left side of Event Viewer lists five categories of events that are tracked: Application, Security, Setup, System, and Forwarded Events. Select one of these options to see the log that you want to view. For details about a particular event, double-click the event to display a dialog box with detailed information about the event.

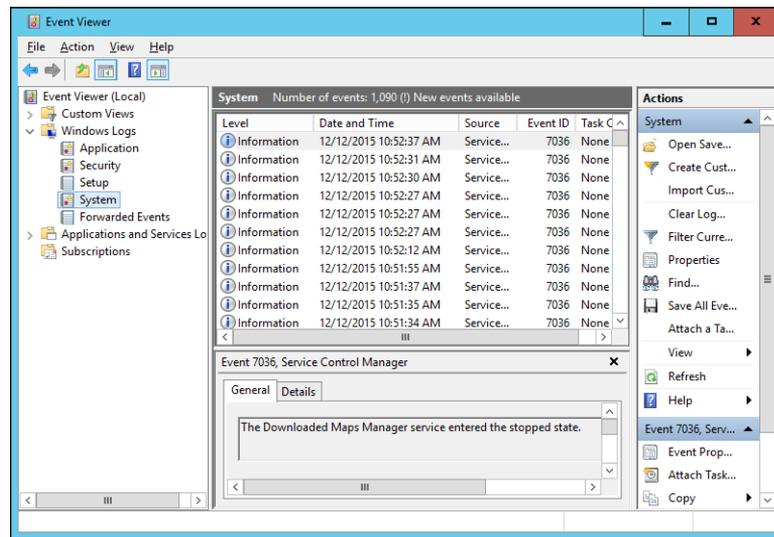


FIGURE 17-5:
View event logs here.

Documenting Your Trials and Tribulations

For a large network, you probably want to invest in problem-management software that tracks each problem through the entire process of troubleshooting, from initial report to final resolution. For small- and medium-sized networks, it's probably sufficient to put together a three-ring binder with pre-printed forms. Or record your log in a Word document or an Excel spreadsheet.

Regardless of how you track your network problems, the tracking log should include the following information:

»» **The real name and the network username of the person reporting the problem**

»» **The date the problem was first reported**

»» **An indication of the severity of the problem**

Is it merely an inconvenience, or is a user unable to complete his or her work because of the problem? Does a workaround exist?

»» **The name of the person assigned to resolve the problem**

»» **A description of the problem**

»» **A list of the software involved, including versions**

»» **A description of the steps taken to solve the problem**

»» **A description of any intermediate steps that were taken to try to solve the problem, along with an indication of whether those steps were “undone” when they didn’t help solve the problem**

»» **The date the problem was finally resolved**

Understanding the need for backups

Working with tape drives and other backup media

Understanding the different types of backups

Mastering tape rotation and other details

Chapter 18

Backing Up Your Data

If you're the hapless network manager, the safety of the data on your network is your responsibility. In fact, it's your primary responsibility. You get paid to lie awake at night worrying about your data. Will it be there tomorrow? If it's not, can you get it back? And — most important — if you can't get it back, will you have a job tomorrow?

This chapter covers the ins and outs of being a good, responsible, trustworthy network manager. No one gives out merit badges for this stuff, but someone should.

Backing Up Your Data

Having data backed up is the cornerstone of any disaster recovery plan. Without backups, a simple hard drive failure can set your company back days or even weeks while it tries to reconstruct lost data. In fact, without backups, your company's very existence is in jeopardy.



REMEMBER

The fundamental goal of backing up is simple: Keep a spare copy of your network's critical data so that no matter what happens, you never lose more than one day's work. The stock market may crash, Earth may be hit by a giant asteroid, or the Cleveland Browns might win the Super Bowl. But as long as you're on top of your backups, you'll survive.

The way to meet the primary goal of backups is, naturally, to make sure that data is reliably backed up every day. For many networks, you can back up all the network hard drives every night. And even if full nightly backups aren't possible, you can still use techniques that can ensure that every file on the network has a backup copy that's no more than one day old.

Choosing Where to Back Up Your Data

If you plan on backing up the data on your network server's hard drives, you obviously need some type of media on which to back up the data. You could copy the data onto CDs, but a 500GB hard drive would need more than 750 CDs for a full backup. That's a few more discs than most people want to keep in the closet. You could use DVDs, but you'll still need about a dozen of them, as well as an hour or so to fill each one. Sigh. That means devoting every Saturday to creating your backups.



TIP

Because of the limitations of CDs and DVDs, most network administrators back up network data to some other type of storage device. The three most common options are

- » **Tape:** Magnetic tape, the oldest storage medium for backups, is still one of the most widely used types. One of the biggest advantages of tape backups is that tape cartridges are small and can thus be easily transported to an offsite location.
- » **Network Attached Storage (NAS):** A *Network Attached Storage* device connects directly to your network. NAS devices are often used as backup devices because they are inexpensive. Depending on your needs, you can acquire enterprise-grade NAS storage that is rack-mounted, or you can purchase inexpensive consumer-grade NAS storage that is portable and so, like tape, can be transported offsite.
- » **Cloud backup:** An increasingly popular option is to use a third-party service to back up your data to a remote location via the Internet. Cloud backup has the advantage of already being offsite.

Backing Up to Tape

Another benefit of using a tape backup is that you can run it unattended. In fact, you can schedule a tape backup to run automatically during off hours when no one is using the network. For unattended backups to work, though, you must ensure

that you have enough tape capacity to back up your entire network server's hard drive without having to manually switch tapes. If your network server has only 100GB of data, you can easily back it up onto a single tape. If you have 1,000GB of data, however, invest in a tape drive that features a magazine changer that can hold several tapes and automatically cycle them in and out of the drive. That way, you can run your backups unattended.

You have several distinct types of tape backup systems to choose from:

- » **Travan drives:** A popular style of tape backup for small servers is a Travan drive, which comes in a variety of models with tape capacities ranging from 20GB to 40GB. You can purchase a 20GB drive for less than \$200.
- » **DAT, DLT, and LTO units:** For larger networks, you can get tape backup units that offer higher capacity and faster backup speed than Travan drives — for more money, of course. Digital audio tape (DAT) units can back up as much as 80GB on a single tape, and DLT (digital linear tape) drives can store up to 800GB on one tape. The newest generation of Linear Tape Open (LTO) drives can store 6TB on a single tape. DAT, DLT, and LTO drives can cost \$1,000 or more, depending on the capacity.
- » **Robotic units:** If you're really up the backup creek, with hundreds of gigabytes to back up, you can get robotic tape backup units that automatically fetch and load tape cartridges from a library. That way, you can do complete backups without having to load tapes manually. As you can likely guess, these units aren't inexpensive: Small ones, which have a library of about eight tapes and a total backup capacity of more than 5,000GB, start at about \$4,000.

Understanding Backup Software

All versions of Windows come with a built-in backup program. In addition, most tape drives come with backup programs that are often faster or more flexible than the standard Windows backup.

You can also purchase sophisticated backup programs that are specially designed for networks that have multiple servers with data that must be backed up. For a basic Windows file server, you can use the backup program that comes with Windows Server. Server versions of Windows come with a decent backup program that can run scheduled, unattended tape backups.

Backup programs do more than just copy data from your hard drive to tape. Backup programs use special compression techniques to squeeze your data so that you can cram more data onto fewer tapes. Compression factors of 2:1 are common, so you

can usually squeeze 100GB of data onto a tape that would hold only 50GB of data without compression. (Tape drive manufacturers tend to state the capacity of their drives by using compressed data, assuming a 2:1 compression ratio. Thus, a 200GB tape has an uncompressed capacity of 100GB.)



WARNING

Whether you achieve a compression factor of 2:1 depends on the nature of the data you're backing up:

- » **Documents:** If your network is used primarily for Microsoft Office applications and is filled with Word and Excel documents, you'll probably get better than 2:1 compression.
- » **Graphics:** If your network data consists primarily of graphic image files, you probably won't get much compression. Most graphic image file formats are already compressed, so they can't be compressed much more by the backup software's compression methods.

Backup programs also help you keep track of which data has been backed up and which hasn't. They also offer options, such as incremental or differential backups that can streamline the backup process, as I describe in the next section.



REMEMBER

If your network has more than one server, invest in good backup software. One popular program is Yosemite Server Backup, made by Barracuda (www.barracuda.com). Besides being able to handle multiple servers, one of the main advantages of backup software is that it can properly back up Microsoft Exchange server data.

Comparing Types of Backups

You can perform five different types of backups. Many backup schemes rely on full daily backups, but for some networks, using a scheme that relies on two or more of these backup types is more practical.

The differences among the five types of backups involve a little technical detail known as the “archive bit,” which indicates whether a file has been modified since it was backed up. The archive bit is a little flag stored along with the filename, creation date, and other directory information. Any time a program modifies a file, the archive bit is set to the On position. That way, backup programs know that the file has been modified and needs to be backed up.

The differences among the various types of backups center on whether they use the archive bit to determine which files to back up, as well as whether they flip the archive bit to the Off position after they back up a file. Table 18-1 summarizes these differences, which I explain in the following sections.



TIP

Backup programs allow you to select any combination of drives and folders to back up. As a result, you can customize the file selection for a backup operation to suit your needs. For example, you can set up one backup plan that backs up all a server’s shared folders and drives, plus its mail server stores, but then leaves out folders that rarely change, such as the operating system folders or installed program folders. You can then back up those folders on a less-regular basis. The drives and folders that you select for a backup operation are collectively called the *backup selection*.

TABLE 18-1 How Backup Types Use the Archive Bit

Backup Type	Selects Files Based on Archive Bit?	Resets Archive Bits After Backing Up?
Normal	No	Yes
Copy	No	No
Daily	No*	No
Incremental	Yes	Yes
Differential	Yes	No

*Selects files based on the Last Modified date.

The archive bit would have made a good Abbott and Costello routine. (“All right, I wanna know who modified the archive bit.” “What.” “Who?” “No, What.” “Wait a minute . . . just tell me what’s the name of the guy who modified the archive bit!” “Right.”)

Normal backups

A *normal backup* — also called a *full backup* — is the basic type of backup. In a normal backup, all files in the backup selection are backed up regardless of whether the archive bit has been set. In other words, the files are backed up even if they haven’t been modified since the last time they were backed up. When each file is backed up, its archive bit is reset, so backups that select files based on the archive bit setting won’t back up the files.

When a normal backup finishes, none of the files in the backup selection has its archive bit set. As a result, if you immediately follow a normal backup with an incremental backup or a differential backup, files won’t be selected for backup by the incremental or differential backup because no file will have its archive bit set.

The easiest backup scheme is to simply schedule a normal backup every night. That way, all your data is backed up on a daily basis. Then, if the need arises, you

can restore files from a single tape or set of tapes. Restoring files is more complicated when other types of backups are involved.



REMEMBER

Do normal backups nightly if you have the tape capacity to do them unattended — that is, without having to swap tapes. If you can't do an unattended normal backup because the amount of data to be backed up is greater than the capacity of your tape drive(s), you have to use other types of backups in combination with normal backups.



TIP

If you can't get a normal backup on a single tape, and you can't afford a second tape drive or a tape changer, take a hard look at the data that's being included in the backup selection. I recently worked on a network that was difficult to back up onto a single tape. When I examined the data that was being backed up, I discovered a large amount of static data that was essentially an online archive of old projects. This data was necessary because network users needed it for research purposes, but the data was read-only. Even though the data never changed, it was being backed up to tape every night, and the backups required two tapes. After I removed this data from the cycle of nightly backups, the backups were able to squeeze onto a single tape again.

If you remove static data from the nightly backup, make sure that you have a secure backup of the static data on tape, CD-RW, or some other media.

Copy backups

A *copy backup* is similar to a normal backup except that the archive bit isn't reset when each file is copied. As a result, copy backups don't disrupt the cycle of normal and incremental or differential backups.

Copy backups usually aren't incorporated into regular, scheduled backups. Instead, you use a copy backup when you want to do an occasional one-shot backup. If you're about to perform an operating system upgrade, for example, you should back up the server before proceeding. If you do a full backup, the archive bits are reset, and your regular backups are disrupted. If you do a copy backup, however, the archive bits of any modified files remain unchanged. As a result, your regular normal and incremental or differential backups are unaffected.

If you don't incorporate incremental or differential backups into your backup routine, the difference between a copy backup and a normal backup is moot.

Daily backups

A *daily backup* backs up just those files that changed the same day when the backup was performed. A daily backup examines the modification date stored with each

file's directory entry to determine whether a file should be backed up. Daily backups don't reset the archive bit.



WARNING

I'm not a big fan of this option because of the small possibility that some files may slip through the cracks. Someone may be working late one night and modify a file after the evening's backups have completed — but before midnight — meaning that those files won't be included in the following night's backups. Incremental or differential backups, which rely on the archive bit rather than the modification date, are more reliable.

Incremental backups

An *incremental backup* backs up only those files that were modified since the last time you did a backup. Incremental backups are a lot faster than full backups because your network users probably modify only a small portion of the files on the server on any given day. As a result, if a full backup takes three tapes, you can probably fit an entire week's worth of incremental backups on a single tape.

When an incremental backup copies each file, it resets the file's archive bit. That way, the file will be backed up again before your next normal backup only when a user modifies the file again.

Here are some thoughts about using incremental backups:



TIP

»» **The easiest way to use incremental backups is the following:**

- A *normal* backup every Monday

If your full backup takes more than 12 hours, you may want to do it on Friday so that it can run over the weekend.

- An *incremental* backup on each remaining normal business day (for example, Tuesday, Wednesday, Thursday, and Friday)

»» **When you use incremental backups, the complete backup consists of the full backup tapes and all the incremental backup tapes that you've made since you did the full backup.**

If the hard drive crashes, and you have to restore the data onto a new drive, you first restore Monday's normal backup and then restore each of the subsequent incremental backups.

»» **Incremental backups complicate restoring individual files because the most recent copy of the file may be on the full backup tape or on any of the incremental backups.**



TECHNICAL
STUFF

Backup programs keep track of the location of the most recent version of each file to simplify the process.

»» **When you use incremental backups, you can choose whether you want to**

- Store each incremental backup on its own tape.
- Append each backup to the end of an existing tape.

Often, you can use a single tape for a week of incremental backups.



TIP

Differential backups

A *differential backup* is similar to an incremental backup except that it doesn't reset the archive bit when files are backed up. As a result, each differential backup represents the difference between the last normal backup and the current state of the hard drive.

To do a full restore from a differential backup, you first restore the last normal backup and then restore the most recent differential backup.

Suppose that you do a normal backup on Monday and differential backups on Tuesday, Wednesday, and Thursday, and your hard drive crashes Friday morning. On Friday afternoon, you install a new hard drive. To restore the data, you first restore the normal backup from Monday. Then you restore the differential backup from Thursday. The Tuesday and Wednesday differential backups aren't needed.

The main difference between incremental and differential backups is that

- »» *Incremental* backups result in smaller and faster backups.
- »» *Differential* backups are easier to restore.

If your users often ask you to restore individual files, consider using differential backups.



TIP

Choosing between Local and Network Backups

When you back up network data, you have two basic approaches to running the backup software:

- » You can perform a *local backup*, in which the backup software runs on the file server itself and backs up data to a tape drive that's installed in the server.
- » Or you can perform a *network backup*, in which you use one network computer to back up data from another network computer. In a network backup, the data has to travel over the network to get to the computer that's running the backup.

If you run the backups from the file server, you'll tie up the server while the backup is running, and users will complain that their server access has slowed to a snail's pace. On the other hand, if you run the backup over the network from a client computer or a dedicated backup server, you'll flood the network with gigabytes of data being backed up. Then your users will complain that the entire network has slowed to a snail's pace.

Network performance is one of the main reasons why you should try to run your backups during off hours, when other users aren't accessing the network. Another reason to run backups during off hours is so that you can perform a more thorough backup. If you run your backup while other users are accessing files, the backup program is likely to skip any files that are being accessed by users at the time the backup runs. As a result, your backup won't include those files. Ironically, the files most likely to get left out of the backup are often the files that need backing up the most, because they're the files that are being used and modified.

Here are some extra thoughts on client and server backups:



TIP

- » **Backing up directly from the server isn't necessarily more efficient than backing up from a client because data doesn't have to travel over the network.** The network may well be faster than the tape drive. The network probably won't slow down backups unless you back up during the busiest time of the day, when hordes of network users are storming the network gates.
- » **To improve network backup speed and to minimize the effect that network backups have on the rest of the network, consider using a 1,000 Mbps switch instead of a normal 100 Mbps switch to connect the servers and the backup client.** That way, network traffic between the server and the backup client won't bog down the rest of the network.
- » **Any files that are open while the backups are running won't get backed up.** That's usually not a problem, because backups are run at off hours when people have gone home. If someone leaves his computer on with a Word document open, however, that Word document won't be backed up. One way to solve this problem is to set up the server so that it automatically logs everyone off the network before the backups begin.

» **Some backup programs have special features that enable them to back up open files.** The backup programs that come with Windows Server (versions 2003 and later) do this by creating a snapshot of the volume when it begins, thus making temporary copies of any files that are modified during the backup. The backup backs up the temporary copies rather than the versions being modified. When the backup finishes, the temporary copies are deleted.

Deciding How Many Sets of Backups to Keep

Don't try to cut costs by purchasing one backup tape and reusing it every day. What happens if you accidentally delete an important file on Tuesday and don't discover your mistake until Thursday? Because the file didn't exist on Wednesday, it won't be on Wednesday's backup tape. If you have only one tape that's reused every day, you're outta luck.

The safest scheme is to use a new backup tape every day and keep all your old tapes in a vault. Pretty soon, though, your tape vault can start looking like the warehouse where they stored the Ark of the Covenant at the end of *Raiders of the Lost Ark*.



TIP

As a compromise between these two extremes, most users purchase several tapes and rotate them. That way, you always have several backup tapes to fall back on, just in case the file you need isn't on the most recent backup tape. This technique is *tape rotation*, and several variations are commonly used:

- » **The simplest approach is to purchase three tapes and label them A, B, and C.** You use the tapes on a daily basis in sequence: A the first day, B the second day, and C the third day; then A the fourth day, B the fifth day, C the sixth day, and so on. On any given day, you have three generations of backups: today's, yesterday's, and the day-before-yesterday's. Computer geeks like to call these the *grandfather*, *father*, and *son* tapes.
- » **Another simple approach is to purchase five tapes and use one each day of the workweek.**
- » **A variation of the preceding bullet is to buy eight tapes.** Take four of them, and write *Tuesday* on one label, *Wednesday* on the second, *Thursday* on the third, and *Friday* on the fourth label. On the other four tapes, write *Monday 1*, *Monday 2*, *Monday 3*, and *Monday 4*. Now tack up a calendar on the wall near

the computer, and number all the Mondays in the year: 1, 2, 3, 4, 1, 2, 3, 4, and so on.

On Tuesday through Friday, you use the appropriate daily backup tape. When you run a full backup on Monday, consult the calendar to decide which Monday tape to use. With this scheme, you always have four weeks' worth of Monday backup tapes, plus individual backup tapes for the rest of the week.

» **If bookkeeping data lives on the network, make a backup copy of all your files (or at least all your accounting files) immediately before closing the books each month; then retain those backups for each month of the year.** This doesn't necessarily mean that you should purchase 12 additional tapes. If you back up just your accounting files, you can probably fit all 12 months on a single tape. Just make sure that you back up with the Append to Tape option rather than the Erase Tape option so that the previous contents of the tape aren't destroyed. Also, treat this accounting backup as completely separate from your normal daily backup routine.



WARNING

Keep at least one recent full backup at another location. That way, if your office should fall victim to an errant Scud missile or a rogue asteroid, you can re-create your data from the backup copy that you stored offsite. Make sure that the person entrusted with the task of taking the backups to this offsite location is trustworthy.

Verifying Tape Reliability

From experience, I've found that although tape drives are very reliable, they do run amok once in a while. The problem is that they don't always tell you when they're not working. A tape drive (especially one of the less-expensive Travan drives; refer to "Backing Up to Tape," earlier in this chapter) can spin along for hours, pretending to back up your data — but in reality, your data isn't being written reliably to the tape. In other words, a tape drive can trick you into thinking that your backups are working just fine. Then, when disaster strikes and you need your backup tapes, you may just discover that the tapes are worthless.



TIP

Don't panic! Here's a simple way to assure yourself that your tape drive is working. Just activate the "compare after backup" feature of your backup software. As soon as your backup program finishes backing up your data, it rewinds the tape, reads each backed-up file, and compares it with the original version on the hard drive. If all files compare, you know that your backups are trustworthy.

Here are some additional thoughts about the reliability of tapes:

- » The compare-after-backup feature doubles the time required to do a backup, but that doesn't matter if your entire backup fits on one tape. You can just run the backup after hours. Whether the backup and repair operation takes one hour or ten doesn't matter, as long as it's finished by the time the network users arrive at work the next morning.
- » If your backups require more than one tape, you may not want to run the compare-after-backup feature every day. Be sure to run it periodically, however, to check that your tape drive is working.
- » If your backup program reports errors, throw away the tape, and use a new tape.
- » Actually, you should ignore that last comment about waiting for your backup program to report errors. You should discard tapes *before* your backup program reports errors. Most experts recommend that you should use a tape only about 20 times before discarding it. If you use the same tape every day, replace it monthly. If you have tapes for each day of the week, replace them twice yearly. If you have more tapes than that, figure out a cycle that replaces tapes after about 20 uses.

Keeping Backup Equipment Clean and Reliable

An important aspect of backup reliability is proper maintenance of your tape drives. Every time you back up to tape, little bits and specks of the tape rub off onto the read and write heads inside the tape drive. Eventually, the heads become too dirty to read or write data reliably.

To counteract this problem, clean the tape heads regularly. The easiest way to clean them is to use a cleaning cartridge for the tape drive. The drive automatically recognizes when you insert a cleaning cartridge and then performs a routine that wipes the cleaning tape back and forth over the heads to clean them. When the cleaning routine is done, the tape is ejected. The whole process takes only about 30 seconds.

Because the maintenance requirements of drives differ, check each drive's user's manual to find out how and how often to clean the drive. As a general rule, clean drives once weekly.

The most annoying aspect of tape drive cleaning is that the cleaning cartridges have a limited life span, and unfortunately, if you insert a used-up cleaning cartridge, the drive accepts it and pretends to clean the drive. For this reason, keep track of how many times you use a cleaning cartridge and replace it as recommended by the manufacturer.

Setting Backup Security

Backups create an often-overlooked security exposure for your network: No matter how carefully you set up user accounts and enforce password policies, if any user (including a guest) can perform a backup of the system, that user may make an unauthorized backup. In addition, your backup tapes themselves are vulnerable to theft. As a result, make sure that your backup policies and procedures are secure by taking the following measures:

- » **Set up a user account for the user who does backups.** Because this user account has backup permission for the entire server, guard its password carefully. Anyone who knows the username and password of the backup account can log on and bypass any security restrictions that you place on that user's normal user ID.
- » **Counter potential security problems by restricting the backup user ID to a certain client and a certain time of the day.** If you're really clever (and paranoid), you can probably set up the backup user's account so that the only program it can run is the backup program.
- » **Use encryption to protect the contents of your backup tapes.**
- » **Secure the backup tapes in a safe location, such as . . . um, a safe.**

IN THIS CHAPTER

Assessing the risk for security

Determining your basic security philosophy

Physically securing your network equipment

Considering user account security

Looking at other network security techniques

Making sure your users are secure

Chapter 19

Securing Your Network

Before you had a network, computer security was easy. You simply locked your door when you left work for the day. You could rest easy, secure in the knowledge that the bad guys would have to break down the door to get to your computer.

The network changes all that. Now, anyone with access to any computer on the network can break into the network and steal *your* files. Not only do you have to lock your door, but you have to make sure that other people lock their doors, too.

Fortunately, network operating systems have built-in provisions for network security, deterring someone from stealing your files even if he does break down the door. All modern network operating systems have security features that are more than adequate for all but the most paranoid users.



TIP

When I say *more than adequate*, I mean it. Most networks have security features that would make even Maxwell Smart happy. Using all these security features is kind of like Smart insisting that the Chief lower the “Cone of Silence” (which worked so well that Max and the Chief couldn’t hear each other!). Don’t make your system so secure that even the good guys can’t get their work done.



WARNING

If any computer on your network is connected to the Internet, you must harden your network against intrusion via the Internet. For more information, see Chapter 20. Also, if your network supports wireless devices, you have wireless security issues. For information about security for wireless networks, see Chapter 9.

Do You Need Security?

Most small networks are in small businesses or departments where everyone knows and trusts everyone else. Folks don’t lock up their desks when they take a coffee break, and although everyone knows where the petty cash box is, money never disappears.

Network security isn’t necessary in an idyllic setting like this one, is it? You bet it is. Here’s why any network should be set up with at least some concern for security:

- » Even in the friendliest office environment, some information is and should be confidential. If this information is stored on the network, you want to store it in a directory that’s available only to authorized users.
- » Not all security breaches are malicious. A network user may be routinely scanning through files and come across a filename that isn’t familiar. The user may then call up the file, only to discover that it contains confidential personnel information, juicy office gossip, or your résumé. Curiosity, rather than malice, is often the source of security breaches.
- » Sure, everyone at the office is trustworthy now. However, what if someone becomes disgruntled, a screw pops loose, and he decides to trash the network files before jumping out the window? What if someone decides to print a few \$1,000 checks before packing off to Tahiti?
- » Sometimes the mere opportunity for fraud or theft can be too much for some people to resist. Give people free access to the payroll files, and they may decide to vote themselves a raise when no one is looking.
- » If you think that your network doesn’t contain any data worth stealing, think again. For example, your personnel records probably contain more than enough information for an identity thief: names, addresses, phone numbers,

Social Security numbers, and so on. Also, your customer files may contain your customers' credit card numbers.

- » Hackers who break into your network may be looking to plant a Trojan horse program on your server, enabling them to use your server for their own purposes. For example, someone may use your server to send thousands of unsolicited spam email messages. The spam won't be traced back to the hackers; it'll be traced back to *you*.
- » Not everyone on the network knows enough about how Windows and the network work to be trusted with full access to your network's data and systems. A careless mouse click can wipe out a directory of network files. One of the best reasons for activating your network's security features is to protect the network from mistakes made by users who don't know what they're doing.

Two Approaches to Security

When you're planning how to implement security on your network, first consider which of two basic approaches to security you'll take:

- » **Open door:** You grant everyone access to everything by default and then place restrictions just on those resources to which you want to limit access.
- » **Closed door:** You begin by denying access to everything and then grant specific users access to the specific resources that they need.

In most cases, an open door policy is easier to implement. Typically, only a small portion of the data on a network really needs security, such as confidential employee records, or secrets, such as the Coke recipe. The rest of the information on a network can be safely made available to everyone who can access the network.

If you choose a closed door approach, you set up each user so that he has access to nothing. Then, you grant each user access only to those specific files or folders that he needs.

A closed door approach results in tighter security but can lead to the Cone of Silence Syndrome: Like how Max and the Chief can't hear each other but still talk while they're under the Cone of Silence, your network users will constantly complain that they can't access the information that they need. As a result, you'll find yourself often adjusting users' access rights. Choose a closed door approach only

if your network contains a lot of sensitive information, and only if you're willing to invest time administrating your network's security policy.

You can think of an open door approach as an *entitlement model*, in which the basic assumption is that users are entitled to network access. In contrast, the closed-door policy is a *permissions model*, in which the basic assumption is that users aren't entitled to anything but must get permissions for every network resource that they access.



If you've never heard of the Cone of Silence, go to YouTube (www.youtube.com) and search for "Cone of Silence." You'll find several clips from the original *Get Smart* series.

Physical Security: Locking Your Doors

The first level of security in any computer network is *physical security*. I'm amazed when I walk into the reception area of an accounting firm and see an unattended computer sitting on the receptionist's desk. Often, the receptionist has logged on to the system and then walked away from the desk, leaving the computer unattended.

Physical security is important for workstations but vital for servers. Any good hacker can quickly defeat all but the most paranoid security measures if they can gain physical access to a server. To protect the server, follow these guidelines:

- » Lock the computer room.
- » Give the key only to people you trust.
- » Keep track of who has the keys.
- » Mount the servers on cases or racks that have locks.
- » Disable the floppy drive on the server.

A common hacking technique is to boot the server from a floppy, thus bypassing the security features of the operating system.

- » Keep a trained guard dog in the computer room and feed it only enough to keep it hungry and mad. (Just kidding.)



There's a big difference between a door with a lock and a locked door. And locks are quite worthless if you don't use them.

Client computers should be physically secure:

- » Instruct users to not leave their computers unattended while they're logged on.
- » In high-traffic areas (such as the receptionist's desk), users should secure their computers with the keylock, if the computer has one.
- » Users should lock their office doors when they leave.



WARNING

Here are some other threats to physical security that you may not have considered:

- » The nightly cleaning crew probably has complete access to your facility. How do you know that the person who vacuums your office every night doesn't really work for your chief competitor or doesn't consider computer hacking to be a sideline hobby? You don't, so consider the cleaning crew to be a threat.
- » What about your trash? Paper shredders aren't just for Enron accountants. Your trash can contain all sorts of useful information: sales reports, security logs, printed copies of the company's security policy, even hand-written passwords. For the best security, every piece of paper that leaves your building via the trash bin should first go through a shredder.
- » Where do you store your backup tapes? Don't just stack them up next to the server. Not only does that make them easy to steal, it also defeats one of the main purposes of backing up your data in the first place: securing your server from physical threats, such as fires. If a fire burns down your computer room and the backup tapes are sitting unprotected next to the server, your company may go out of business and you'll certainly be out of a job. Store the backup tapes securely in a fireproof safe and keep a copy off-site, too.
- » I've seen some networks in which the servers are in a locked computer room, but the hubs or switches are in an unsecured closet. Remember that every unused port on a hub or a switch represents an open door to your network. The hubs and switches should be secured just like the servers.

Securing User Accounts

Next to physical security, the careful use of user accounts is the most important type of security for your network. Properly configured user accounts can prevent unauthorized users from accessing the network, even if they gain physical access to the network. The following sections describe some of the steps that you can take to strengthen your network's use of user accounts.

Obfuscating your usernames

Huh? When it comes to security, *obfuscation* simply means picking obscure usernames. For example, most network administrators assign usernames based on some combination of the user's first and last name, such as `BarnyM` or `baMiller`. However, a hacker can easily guess such a user ID if he or she knows the name of at least one employee. After the hacker knows a username, he or she can focus on breaking the password.

You can slow down a hacker by using names that are more obscure. Here are some suggestions on how to do that:

- » Add a random three-digit number to the end of the name. For example: `BarnyM320` or `baMiller977`.
- » Throw a number or two into the middle of the name. For example: `Bar6nyM` or `ba9Miller2`.
- » Make sure that usernames are different from email addresses. For example, if a user's email address is `baMiller@Mydomain.com`, do *not* use `baMiller` as the user's account name. Use a more obscure name.



WARNING

Do *not* rely on obfuscation to keep people out of your network! Security by obfuscation doesn't work. A resourceful hacker can discover the most obscure names. Obfuscation can *slow* intruders, not stop them. If you slow intruders down, you're more likely to discover them before they crack your network.

Using passwords wisely

One of the most important aspects of network security is the use of passwords.



REMEMBER

Usernames aren't usually considered *secret*. Even if you use obscure names, even casual hackers will eventually figure them out.

Passwords, on the other hand, are top secret. Your network password is the one thing that keeps an impostor from logging on to the network by using your username and therefore receiving the same access rights that you ordinarily have. *Guard your password with your life.*

Here are some tips for creating good passwords:

- » Don't use obvious passwords, such as your last name, your kid's name, or your dog's name.
- » Don't pick passwords based on your hobbies. A friend of mine is a boater, and his password is the name of his boat. Anyone who knows him can



WARNING

quickly guess his password. Five lashes for naming your password after your boat.

- » Store your password in your head — not on paper.

Especially bad: Writing your password down on a sticky note and sticking it on your computer's monitor.

- » Most network operating systems enable you to set an expiration time for passwords. For example, you can specify that passwords expire after 30 days. When a user's password expires, the user must change it. Your users may consider this process a hassle, but it helps to limit the risk of someone swiping a password and then trying to break into your computer system later.
- » You can configure user accounts so that when they change passwords, they can't reuse a *recent* password. For example, you can specify that the new password can't be identical to any of the user's past three passwords.
- » You can also configure security policies so that passwords must include a mixture of uppercase letters, lowercase letters, numerals, and special symbols. Thus, passwords like DIMWIT or DUFUS are out. Passwords like 87dIM@wit or duF39&US are in.



WARNING

- » Some administrators of small networks opt against passwords altogether because they feel that security isn't an issue on their network. Or short of that, they choose obvious passwords, assign every user the same password, or print the passwords on giant posters and hang them throughout the building. Ignoring basic password security is rarely a good idea, even in small networks. You should consider not using passwords only if your network is very small (say, two or three computers), if you don't keep sensitive data on a file server, or if the main reason for the network is to share access to a printer rather than sharing files. (Even if you don't use passwords, imposing basic security precautions, like limiting access that certain users have to certain network directories, is still possible. Just remember that if passwords aren't used, nothing prevents a user from signing on by using someone else's username.)

Generating passwords For Dummies

How do you come up with passwords that no one can guess but that you can remember? Most security experts say that the best passwords don't correspond to any words in the English language but consist of a random sequence of letters, numbers, and special characters. Yet, how in the heck are you supposed to memorize a password like Dks4%DJ2? Especially when you have to change it three weeks later to something like 3pQ&X(d8.



TIP

Here's a compromise solution that enables you to create passwords that consist of two four-letter words back to back. Take your favorite book (if it's this one, you need to get a life) and turn to any page at random. Find the first four- or five-letter word on the page. Suppose that word is *When*. Then repeat the process to find another four- or five-letter word; say you pick the word *Most* the second time. Now combine the words to make your password: *WhenMost*. I think you'll agree that *WhenMost* is easier to remember than *3PQ&X(D8* and is probably just about as hard to guess. I probably wouldn't want the folks at the Los Alamos Nuclear Laboratory using this scheme, but it's good enough for most of us.

Here are additional thoughts on concocting passwords from your favorite book:

- » If the words end up being the same, pick another word. And pick different words if the combination seems too commonplace, such as *WestWind* or *FootBall*.
- » For an interesting variation, insert a couple of numerals or special characters between the words. You end up with passwords like *into#cat*, *ball3%and*, or *tree47wing*. If you want, use the page number of the second word as a separator. For example, if the words are *know* and *click* and the second word comes from page 435, use *know435click*.
- » To further confuse your friends and enemies, use medieval passwords by picking words from Chaucer's *Canterbury Tales*. Chaucer is a great source for passwords because he lived before the days of word processors with spell-checkers. He wrote *seyd* instead of *said*, *gret* instead of *great*, *welk* instead of *walked*, *litel* instead of *little*. And he used lots of seven-letter and eight-letter words suitable for passwords, such as *glotenyne* (gluttony), *benygne* (benign), and *opynyoun* (opinion). And he got A's in English.
- » If you use any of these password schemes and someone breaks into your network, don't blame me. You're the one who's too lazy to memorize *D#Sc\$h4@bb3xaz5*.
- » If you do decide to go with passwords, such as *Kdl22UR3xdkL*, you can find random password generators on the Internet. Just go to a search engine, such as Google, and search for Password Generator. You'll find Web pages that generate random passwords based on criteria that you specify, such as how long the password should be, whether it should include letters, numbers, punctuation, uppercase and lowercase letters, and so on.



TIP

Secure the Administrator account

It stands to reason that at least one network user must have the authority to use the network without any of the restrictions imposed on other users. This user is

the *administrator*. The administrator is responsible for setting up the network's security system. To do that, the administrator must be exempt from all security restrictions.



WARNING

Many networks automatically create an administrator user account when you install the network software. The username and password for this initial administrator are published in the network's documentation and are the same for all networks that use the same network operating system. One of the first things that you must do after getting your network up and running is to change the password for this standard administrator account. Otherwise, your elaborate security precautions are a complete waste of time. Anyone who knows the default administrator username and password can access your system with full administrator rights and privileges, thus bypassing the security restrictions that you so carefully set up.



WARNING

Don't forget the password for the administrator account! If a network user forgets his or her password, you can log on as the supervisor and change that user's password. If you forget the administrator's password, though, you're stuck.

Managing User Security

User accounts are the backbone of network security administration. Through the use of user accounts, you can determine who can access your network as well as what network resources each user can and can't access. You can restrict access to the network to just specific computers or to certain hours of the day. In addition, you can lock out users who no longer need to access your network. The following sections describe the basics of setting up user security for your network.

User accounts

Every user who accesses a network must have a *user account*. User accounts allow the network administrator to determine who can access the network and what network resources each user can access. In addition, the user account can be customized to provide many convenient features for users, such as a personalized Start menu or a display of recently used documents.

Every user account is associated with a *username* (sometimes called a *user ID*), which the user must enter when logging on to the network. Each account also has other information associated with it. In particular:

- » **The user's password:** This also includes the password policy, such as how often the user has to change his or her password, how complicated the password must be, and so on.

- » **The user's contact information:** This includes full name, phone number, email address, mailing address, and other related information.
- » **Account restrictions:** This includes restrictions that allow the user to log on only during certain times of the day. This feature can restrict your users to normal working hours so that they can't sneak in at 2 a.m. to do unauthorized work. This feature also discourages your users from working overtime because they can't access the network after hours, so use it judiciously. You can also specify that the user can log on only at certain computers.
- » **Account status:** You can temporarily disable a user account so the user can't log on.
- » **Home directory:** This specifies a shared network folder where the user can store documents.
- » **Dial-in permissions:** These authorize the user to access the network remotely via a dialup connection.
- » **Group memberships:** These grant the user certain rights based on groups to which she belongs.

For more information, see the section, "Group therapy," later in this chapter.



TIP

Built-in accounts

Most network operating systems come preconfigured with two built-in accounts, Administrator and Guest. In addition, some server services, such as web or database servers, create their own user accounts under which to run. The following sections describe the characteristics of these accounts.

- » **The Administrator account:** The Administrator account is the King of the Network. This user account isn't subject to any of the account restrictions to which mere mortal accounts must succumb. If you log on as the administrator, you can do anything. For this reason, avoid using the Administrator account for routine tasks. Log in as the Administrator only when you really need to.

Because the Administrator account has unlimited access to your network, it's imperative that you secure it immediately after you install the server. When the operating system Setup program asks for a password for the Administrator account, start with a good random mix of uppercase and lowercase letters, numbers, and symbols. Don't pick some easy-to-remember password to get started, thinking you'll change it to something more cryptic later. You'll forget, and in the meantime, someone will break in and reformat the server's C: drive or steal your customer's credit card numbers.



TIP

» **The Guest account:** Another commonly created default account is the *Guest account*. This account is set up with a blank password and — if any — access rights. The Guest account is designed to allow anyone to step up to a computer and log on, but after they do, it then prevents them from doing anything. Sounds like a waste of time to me. I suggest you disable the Guest account.

» **Service accounts:** Some network users aren't actual people. I don't mean that some of your users are subhuman. Rather, some users are actually software processes that require access to secure resources, and therefore, require user accounts. These user accounts are usually created automatically for you when you install or configure server software.

For example, when you install Microsoft's web server (IIS), an Internet user account called IUSR is created. The complete name for this account is IUSR_<servername>. So if the server is named WEB1, the account is named IUSR_WEB1. IIS uses this account to allow anonymous Internet users to access the files of your website.

Don't mess with these accounts unless you know what you're doing. For example, if you delete or rename the IUSR account, you must reconfigure IIS to use the changed account. If you don't, IIS will deny access to anyone trying to reach your site. (Assuming that you *do* know what you're doing, renaming these accounts can increase your network's security. However, don't start playing with these accounts until you've researched the ramifications.)



TIP

User rights

User accounts and passwords are the front line of defense in the game of network security. After a user accesses the network by typing a valid user ID and password, the second line of security defense — *rights* — comes into play.

In the harsh realities of network life, all users are created equal, but some users are more equal than others. The Preamble to the Declaration of Network Independence contains the statement “We hold these truths to be self-evident, that *some* users are endowed by the network administrator with certain inalienable rights. . . .”

The rights that you can assign to network users depend on which network operating system you use. These are some of the possible user rights for Windows servers:

» **Log on locally:** The user can log on to the server computer directly from the server's keyboard.

» **Change system time:** The user can change the time and date registered by the server.

- » **Shut down the system:** The user can perform an orderly shutdown of the server.
- » **Back up files and directories:** The user can perform a backup of files and directories on the server.
- » **Restore files and directories:** The user can restore backed-up files.
- » **Take ownership of files and other objects:** The user can take over files and other network resources that belong to other users.

NetWare has a similar set of user rights.

Permissions (who gets what)

User rights control what a user can do on a network-wide basis. *Permissions* enable you to fine-tune your network security by controlling access to specific network resources, such as files or printers, for individual users or groups. For example, you can set up permissions to allow users into the accounting department to access files in the server's `\ACCTG` directory. Permissions can also enable some users to read certain files but not modify or delete them.

NETWORK RIGHTS WE WANT TO SEE

The network rights allowed by most operating systems are pretty boring. Here are a few rights I wish would be allowed:

- **Cheat:** Provides a special option that enables you to see what cards the other players are holding when you're playing Hearts.
- **Spy:** Eavesdrops on other users' Internet sessions so you can find out what websites they're viewing.
- **Grumble:** Automatically sends email messages to other users that explain how busy, tired, or upset you are.
- **Set pay:** Grants you special access to the payroll system so that you can give yourself a pay raise.
- **Sue:** In America, everyone has the right to sue. So this right should be granted automatically to all users.
- **Fire:** Wouldn't it be great if the network could grant you the right to play Donald Trump and fire your annoying co-workers?

Each network operating system manages permissions in a different way. Whatever the details, the effect is that you can give permission to each user to access certain files, folders, or drives in certain ways. For example, you might grant a user full access to some files but grant read-only access to other files.



TIP

Any permissions you specify for a folder apply automatically to any of that folder's subfolders, unless you explicitly specify different permissions for the subfolder.



TECHNICAL
STUFF

You can use Windows permissions only for files or folders that are created on drives formatted as NTFS or ReFS volumes. If you insist on using FAT or FAT32 for your Windows shared drives, you can't protect individual files or folders on the drives. This is one of the main reasons for using NTFS for your Windows servers.

Group therapy

A *group account* is an account that doesn't represent an individual user. Instead, it represents a group of users who use the network in a similar way. Instead of granting access rights to each of these users individually, you can grant the rights to the group and then assign individual users to the group. When you assign a user to a group, that user inherits the rights specified for the group.

For example, suppose that you create a group named Accounting for the accounting staff and then allow members of the Accounting group access to the network's accounting files and applications. Then, instead of granting each accounting user access to those files and applications, you simply make each accounting user a member of the Accounting group.

Here are a few additional details about groups:

- » Groups are one of the keys to network management nirvana. As much as possible, avoid managing network users individually. Instead, clump them into groups and manage the groups. When all 50 users in the accounting department need access to a new file share, would you rather update 50 user accounts or just 1 group account?
- » A user can belong to more than one group. Then, the user inherits the rights of each group. For example, you can have groups set up for Accounting, Sales, Marketing, and Finance. A user who needs to access both Accounting and Finance information can be made a member of both groups. Likewise, a user who needs access to both Sales and Marketing information can be made a member of both the Sales and Marketing groups.

- » You can grant or revoke specific rights to individual users to override the group settings. For example, you may grant a few extra permissions for the manager of the accounting department. You may also impose a few extra restrictions on certain users.

User profiles

User profiles are a Windows feature that keeps track of an individual user's preferences for his or her Windows configuration. For a non-networked computer, profiles enable two or more users to use the same computer, each with his or her own desktop settings, such as wallpaper, colors, Start menu options, and so on.

The real benefit of user profiles becomes apparent when profiles are used on a network. A user's profile can be stored on a server computer and accessed whenever that user logs on to the network from any Windows computer on the network.

The following are some of the elements of Windows that are governed by settings in the user profile:

- » Desktop settings from the Display Properties dialog box, including wallpaper, screen savers, and color schemes
- » Start menu programs and Windows toolbar options
- » Favorites, which provide easy access to the files and folders that the user accesses often
- » Network settings, including drive mappings, network printers, and recently visited network locations
- » Application settings, such as option settings for Microsoft Word
- » The Documents folder

Logon scripts

A *logon script* is a batch file that runs automatically whenever a user logs on. Logon scripts can perform several important logon tasks for you, such as mapping network drives, starting applications, synchronizing the client computer's time-of-day clock, and so on. Logon scripts reside on the server. Each user account can specify whether to use a logon script and which script to use.

This sample logon script maps a few network drives and synchronizes the time:

```
net use m: \\MYSERVER\Acct
net use n: \\MYSERVER\Admin
net use o: \\MYSERVER\Dev
net time \\MYSERVER /set /yes
```

Logon scripts are a little out of vogue because most of what a logon script does can be done via user profiles. Still, many administrators prefer the simplicity of logon scripts, so they're still used even on Windows 2016 Server systems.

Securing Your Users

Security techniques, such as physical security, user account security, server security, and locking down your servers are child's play compared with the most difficult job of network security: securing your network's users. All the best-laid security plans will go for naught if your users write their passwords on sticky notes and post them on their computers.

The key to securing your network users is to create a written network security policy and to stick to it. Have a meeting with everyone to go over the security policy to make sure that everyone understands the rules. Also, make sure to have consequences when violations occur.

Here are some suggestions for some basic security rules that can be incorporated into your security policy:

- » Never write down your password or give it to someone else.
- » Accounts shouldn't be shared. Never use someone else's account to access a resource that you can't access under your own account. If you need access to some network resource that isn't available to you, formally request access under your own account.
- » Likewise, never give your account information to a co-worker so that he or she can access a needed resource. Your co-worker should instead formally request access under his or her own account.
- » Don't install any software or hardware on your computer without first obtaining permission. This especially includes wireless access devices or modems.
- » Don't enable file and printer sharing on workstations without first getting permission.
- » Never attempt to disable or bypass the network's security features.

IN THIS CHAPTER

Understanding what firewalls do

Examining the different types of firewalls

Using the built-in Windows firewall

Looking at virus protection

Patching your computers

Chapter 20

Hardening Your Network

If your network is connected to the Internet, a whole host of security issues bubble to the surface. You probably connected your network to the Internet so that your network's users could get out to the Internet. Unfortunately, however, your Internet connection is a two-way street. Not only does it enable your network's users to step outside the bounds of your network to access the Internet, but it also enables others to step in and access your network.

And step in they will. The world is filled with hackers who are looking for networks like yours to break into. They may do it just for the fun of it, or they may do it to steal your customer's credit card numbers or to coerce your mail server into sending thousands of spam messages on their behalf. Whatever their motive, rest assured that your network will be broken into if you leave it unprotected.

This chapter presents an overview of three basic techniques for securing your network's Internet connection: controlling access via a firewall, detecting viruses with antivirus software, and fixing security flaws with software patches.

Firewalls

A *firewall* is a security-conscious router that sits between the Internet and your network with a single-minded task: preventing *them* from getting to *you*. The firewall acts as a security guard between the Internet and your local area network (LAN). All network traffic into and out of the LAN must pass through the firewall, which prevents unauthorized access to the network.



WARNING

Some type of firewall is a must-have if your network has a connection to the Internet, whether that connection is broadband (cable modem or DSL), T1, or some other high-speed connection. Without it, sooner or later a hacker will discover your unprotected network and tell his friends about it. Within a few hours your network will be toast.

You can set up a firewall using two basic ways. The easiest way is to purchase a *firewall appliance*, which is basically a self-contained router with built-in firewall features. Most firewall appliances include a web-based interface that enables you to connect to the firewall from any computer on your network using a browser. You can then customize the firewall settings to suit your needs.

Alternatively, you can set up a server computer to function as a firewall computer. The server can run just about any network operating system, but most dedicated firewall systems run Linux.

Whether you use a firewall appliance or a firewall computer, the firewall must be located between your network and the Internet, as shown in Figure 20-1. Here, one end of the firewall is connected to a network hub, which is, in turn, connected to the other computers on the network. The other end of the firewall is connected to the Internet. As a result, all traffic from the LAN to the Internet and vice versa must travel through the firewall.

The term *perimeter* is sometimes used to describe the location of a firewall on your network. In short, a firewall is like a perimeter fence that completely surrounds your property and forces all visitors to enter through the front gate.



WARNING

In large networks — especially campus-wide or even metropolitan networks — it's sometimes hard to figure out exactly where the perimeter is located. If your network has two or more wide area network (WAN) connections, make sure that every one of those connections connects to a firewall and not directly to the network. You can do this by providing a separate firewall for each WAN connection or by using a firewall with more than one WAN port.

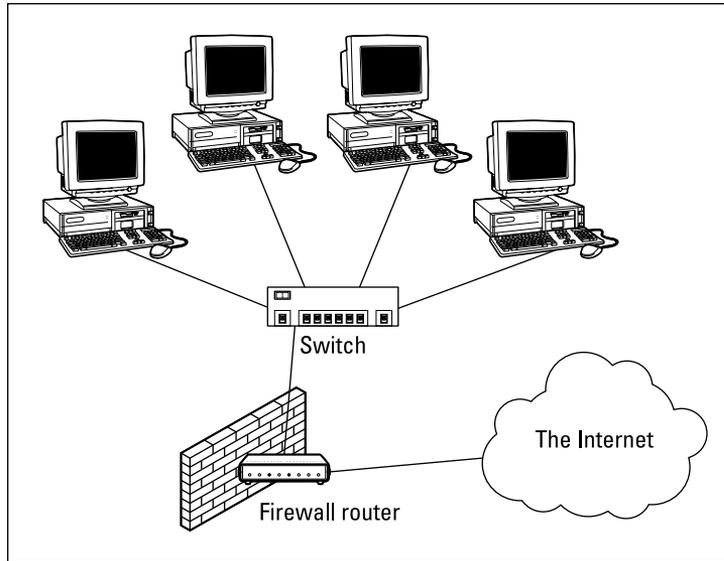


FIGURE 20-1:
A firewall router
creates a secure
link between a
network and the
Internet.

The Many Types of Firewalls

Firewalls employ four basic techniques to keep unwelcome visitors out of your network. The following sections describe these basic firewall techniques.

Packet filtering

A *packet-filtering* firewall examines each packet that crosses the firewall and tests the packet according to a set of rules that you set up. If the packet passes the test, it's allowed to pass. If the packet doesn't pass, it's rejected.

Packet filters are the least expensive type of firewall. As a result, packet-filtering firewalls are very common. However, packet filtering has a number of flaws that knowledgeable hackers can exploit. As a result, packet filtering by itself doesn't make for a fully effective firewall.

Packet filters work by inspecting the source and destination IP and port addresses contained in each *TCP/IP* packet. *TCP/IP ports* are numbers that are assigned to specific services that help to identify for which service each packet is intended. For example, the port number for the HTTP protocol is 80. As a result, any incoming packets headed for an HTTP server will specify port 80 as the destination port.

Port numbers are often specified with a colon following an IP address. For example, the HTTP service on a server whose IP address is 192.168.10.133 would be 192.168.10.133:80.

Literally thousands of established ports are in use. Table 20-1 lists a few of the most popular ports.

TABLE 20-1

Some Well-Known TCP/IP Ports

Port	Description
20	File Transfer Protocol (FTP)
21	File Transfer Protocol (FTP)
22	Secure Shell Protocol (SSH)
23	Telnet
25	Simple Mail Transfer Protocol (SMTP)
53	Domain Name Server (DNS)
80	World Wide Web (HTTP)
110	Post Office Protocol (POP3)
119	Network News Transfer Protocol (NNTP)
137	NetBIOS Name Service
138	NetBIOS Datagram Service
139	NetBIOS Session Service
143	Internet Message Access Protocol (IMAP)
161	Simple Network Management Protocol (SNMP)
194	Internet Relay Chat (IRC)
389	Lightweight Directory Access Protocol (LDAP)
396	NetWare over IP
443	HTTP over TLS/SSL (HTTPS)

The rules that you set up for the packet filter either permit or deny packets that specify certain IP addresses or ports. For example, you may permit packets that are intended for your mail server or your web server and deny all other packets. Or, you may set up a rule that specifically denies packets that are heading for the

ports used by NetBIOS. This rule keeps Internet hackers from trying to access NetBIOS server resources, such as files or printers.

One of the biggest weaknesses of packet filtering is that it pretty much trusts that the packets themselves are telling the truth when they say who they're from and who they're going to. Hackers exploit this weakness by using a hacking technique called *IP spoofing*, in which they insert fake IP addresses in packets that they send to your network.

Another weakness of packet filtering is that it examines each packet in isolation, without considering what packets have gone through the firewall before and what packets may follow. In other words, packet filtering is *stateless*. Rest assured that hackers have figured out how to exploit the stateless nature of packet filtering to get through firewalls.

In spite of these weaknesses, packet filter firewalls have several advantages that explain why they're commonly used:

- » **Packet filters are very efficient.** They hold up each inbound and outbound packet for only a few milliseconds while they look inside the packet to determine the destination and source ports and addresses. After these addresses and ports have been determined, the packet filter quickly applies its rules and either sends the packet along or rejects it. In contrast, other firewall techniques have a more noticeable performance overhead.
- » **Packet filters are almost completely transparent to users.** The only time a user will be aware that a packet filter firewall is being used is when the firewall rejects packets. Other firewall techniques require that clients and/or servers be specially configured to work with the firewall.
- » **Packet filters are inexpensive.** Most routers include built-in packet filtering.

Stateful packet inspection (SPI)

Stateful packet inspection (SPI), is a step up in intelligence from simple packet filtering. A firewall with SPI looks at packets in groups rather than individually. It keeps track of which packets have passed through the firewall and can detect patterns that indicate unauthorized access. In some cases, the firewall may hold on to packets as they arrive until the firewall has gathered enough information to make a decision about whether the packets should be authorized or rejected.



TIP

Stateful packet inspection was once found only on expensive, enterprise-level routers. Now, however, SPI firewalls are affordable enough for small- or medium-sized networks to use.

Circuit-level gateway

A *circuit-level gateway* manages connections between clients and servers based on TCP/IP addresses and port numbers. After the connection is established, the gateway doesn't interfere with packets flowing between the systems.

For example, you could use a Telnet circuit-level gateway to allow Telnet connections (port 23) to a particular server and prohibit other types of connections to that server. After the connection is established, the circuit-level gateway allows packets to flow freely over the connection. As a result, the circuit-level gateway can't prevent a Telnet user from running specific programs or using specific commands.

Application gateway

An *application gateway* is a firewall system that's more intelligent than a packet-filtering, stateful packet inspection, or circuit-level gateway firewall. Packet filters treat all TCP/IP packets the same. In contrast, application gateways know the details about the applications that generate the packets that pass through the firewall. For example, a web application gateway is aware of the details of HTTP packets. As a result, it can examine more than just the source and destination addresses and ports to determine whether the packets should be allowed to pass through the firewall.

In addition, application gateways work as proxy servers. Simply put, a *proxy server* is a server that sits between a client computer and a real server. The proxy server intercepts packets that are intended for the real server and processes them. The proxy server can examine the packet and decide to pass it on to the real server, or it can reject the packet. Or the proxy server may be able to respond to the packet itself, without involving the real server at all.

For example, web proxies often store copies of commonly used web pages in a local cache. When a user requests a web page from a remote web server, the proxy server intercepts the request and checks to see whether it already has a copy of the page in its cache. If so, the web proxy returns the page directly to the user. If not, the proxy passes the request on to the real server.

Application gateways are aware of the details of how various types of TCP/IP servers handle sequences of TCP/IP packets, so they can make more intelligent decisions about whether an incoming packet is legitimate or is part of an attack. As a result, application gateways are more secure than simple packet-filtering firewalls, which can deal with only one packet at a time.

The improved security of application gateways, however, comes at a price. Application gateways are more expensive than packet filters, both in terms of their purchase price and in the cost of configuring and maintaining them. In addition, application gateways slow down the network performance because they do more detailed checking of packets before allowing them to pass.

The Built-In Windows Firewall

All versions of Windows since Windows XP come with a built-in packet-filtering firewall. If you don't have a separate firewall router, you can use this built-in firewall to provide a basic level of protection. See Chapter 8 for the steps to follow to configure the Windows Firewall.



WARNING

Do *not* enable the Windows Firewall if you're using a separate firewall router to protect your network. Because the other computers on the network are connected directly to the router and not to your computer, the firewall won't protect the rest of the network. Additionally, as an unwanted side effect, the rest of the network will lose the ability to access your computer.

Virus Protection

Viruses are one of the most misunderstood computer phenomena around these days. What is a virus? How does it work? How does it spread from computer to computer? I'm glad you asked.

What is a virus?

Make no mistake — viruses are real. Now that most people are connected to the Internet, viruses have really taken off. Every computer user is susceptible to attacks by computer viruses, and using a network increases your vulnerability because it exposes all network users to the risk of being infected by a virus that lands on any one network user's computer.

Viruses don't just spontaneously appear out of nowhere. *Viruses* are computer programs that are created by malicious programmers who've lost a few screws and should be locked up.

What makes a virus a virus is its capability to make copies of itself that can be spread to other computers. These copies, in turn, make still more copies that spread to still more computers, and so on, ad nauseam.

Then, the virus waits patiently until something triggers it — perhaps when you type a particular command or press a certain key, when a certain date arrives, or when the virus creator sends the virus a message. What the virus does when it strikes also depends on what the virus creator wants the virus to do. Some viruses harmlessly display a “gotcha” message. Some send email to everyone it finds in your address book. Some wipe out all the data on your hard drive. Ouch.

A few years back, viruses moved from one computer to another by latching themselves onto floppy disks. Whenever you borrowed a floppy disk from a buddy, you ran the risk of infecting your own computer with a virus that may have stowed away on the disk.

Nowadays, virus programmers have discovered that email is a much more efficient method to spread their viruses. Typically, a virus masquerades as a useful or interesting email attachment, such as instructions on how to make \$1,000,000 in your spare time, pictures of naked celebrities, or a Valentine’s Day greeting from your long-lost sweetheart. When a curious but unsuspecting user double-clicks the attachment, the virus springs to life, copying itself onto the user’s computer and, in some cases, sending copies of itself to all the names in the user’s address book.

After the virus has worked its way onto a networked computer, the virus can then figure out how to spread itself to other computers on the network.

Here are some more tidbits about protecting your network from virus attacks:

- » The term *virus* is often used to refer not only to true virus programs (which can replicate themselves) but also to any other type of program that’s designed to harm your computer. These programs include so-called *Trojan horse* programs that usually look like games but are, in reality, hard drive formatters.
- » A *worm* is similar to a virus, but it doesn’t actually infect other files. Instead, it just copies itself onto other computers on a network. After a worm has copied itself onto your computer, there’s no telling what it may do there. For example, a worm may scan your hard drive for interesting information, such as passwords or credit card numbers, and then email them to the worm’s author.
- » Computer virus experts have identified several thousand “strains” of viruses. Many of them have colorful names, such as the I Love You virus, the Stoned virus, and the Michelangelo virus.

- » Antivirus programs can recognize known viruses and remove them from your system, and they can spot the telltale signs of unknown viruses. Unfortunately, the idiots who write viruses aren't idiots (in the intellectual sense), so they're constantly developing new techniques to evade detection by antivirus programs. New viruses are frequently discovered, and antivirus programs are periodically updated to detect and remove them.

Antivirus programs

The best way to protect your network from virus infection is to use an antivirus program. These programs have a catalog of several thousand known viruses that they can detect and remove. In addition, they can spot the types of changes that viruses typically make to your computer's files, thus decreasing the likelihood that some previously unknown virus will go undetected.

Newer versions of Windows (8 and later) include built-in antivirus protection. For older versions, you can download an excellent free antivirus solution from Microsoft called Microsoft Security Essentials. Popular alternatives to Microsoft's built-in or free antivirus protection include Norton AntiVirus, Webroot SecureAnywhere Antivirus, and Kaspersky Antivirus.

The people who make antivirus programs have their fingers on the pulse of the virus world and often release updates to their software to combat the latest viruses. Because virus writers are constantly developing new viruses, your antivirus software is next to worthless unless you keep it up-to-date by downloading the latest updates.

The following are several approaches to deploying antivirus protection on your network:

- » You can install antivirus software on each network user's computer. This technique would be the most effective if you could count on all your users to keep their antivirus software up-to-date. Because that's an unlikely proposition, you may want to adopt a more reliable approach to virus protection.
- » Managed antivirus services place antivirus client software on each client computer in your network. Then, an antivirus server automatically updates the clients on a regular basis to make sure that they're kept up to date.
- » Server-based antivirus software protects your network servers from viruses. For example, you can install antivirus software on your mail server to scan all incoming mail for viruses and remove them before your network users ever see them.

- » Some firewall appliances include antivirus enforcement checks that don't allow your users to access the Internet unless their antivirus software is up to date. This type of firewall provides the best antivirus protection available.

Safe computing

Besides using an antivirus program, you can take a few additional precautions to ensure virus-free computing. If you haven't talked to your kids about these safe-computing practices, you had better do so soon.

- » Regularly back up your data. If a virus hits you and your antivirus software can't repair the damage, you may need the backup to recover your data. Make sure that you restore from a backup that was created before you were infected by the virus!
- » If you buy software from a store and discover that the seal has been broken on the disk package, take the software back. Don't try to install it on your computer. You don't hear about tainted software as often as you hear about tainted beef, but if you buy software that's been opened, it may well be laced with a virus infection.
- » Use your antivirus software to scan your disk for virus infection after your computer has been to a repair shop or worked on by a consultant. These guys don't intend harm, but they occasionally spread viruses accidentally, simply because they work on so many strange computers.
- » Don't open email attachments from people you don't know or attachments you weren't expecting.
- » Use your antivirus software to scan any CD-ROM or flash drive that doesn't belong to you before you access any of its files.

Patching Things Up

One of the annoyances that every network manager faces is applying software patches to keep the operating system and other software up to date. A software *patch* is a minor update that fixes the small glitches that crop up from time to time, such as minor security or performance issues. These glitches aren't significant enough to merit a new version of the software, but they're important enough to require fixing. Most of the patches correct security flaws that computer hackers have uncovered in their relentless attempts to prove that they are smarter than the security programmers at Microsoft.

Periodically, all the recently released patches are combined into a *service pack*. Although the most diligent network administrators apply all patches when they're released, many administrators just wait for the service packs:



TIP

- » For all versions of Windows, you can use the Windows Update website to apply patches to keep your operating system and other Microsoft software up to date. Windows Update scans your computer's software and creates a list of software patches and other components that you can download and install. To use Windows Update, open the Control Panel, click System and Security, and then click Windows Update.
- » You can configure Windows Update to automatically notify you of updates so you don't have to remember to check for new patches.
- » You can subscribe to a service that automatically sends you email to let you know of new patches and updates.



TIP

Keeping a large network patched can be one of the major challenges of network administration. If you have more than a few dozen computers on your network, consider investing in server-based software that's designed to simplify the process. For example, Lumension (www.lumension.com) is a server-based program that collects software patches from a variety of manufacturers and lets you create distributions that are automatically pushed out to client computers. With software like Lumension, you don't have to rely on end users to download and install patches, and you don't have to visit each computer individually to install patches.

IN THIS CHAPTER

Understanding performance problems

Looking at bottlenecks

Developing a procedure for solving performance problems

Monitoring performance

Implementing other tips for speeding up your network

Chapter 21

Network Performance Anxiety

The term *network performance* refers to how efficiently the network responds to users' needs. Obviously, any access to resources that involves the network is slower than similar access that doesn't involve the network. For example, opening a Word document that resides on a network file server takes longer than opening a similar document that resides on the user's local hard drive. However, it shouldn't take *much* longer. If it does, you have a network performance problem.

This chapter is a general introduction to the practice of tuning your network so that it performs as well as possible. Keep in mind that many specific bits of network tuning advice are scattered throughout this book. In this chapter, you can find some specific techniques for analyzing your network's performance, taking corrective action when a performance problem develops, and charting your progress.

Why Administrators Hate Performance Problems

Network performance problems are among the most difficult network problems to track down and solve. If a user simply can't access the network, it usually doesn't take long to figure out why: The cable is unplugged, a network card is malfunctioning, or the user doesn't have permission to access the resource, for example. After you do a little investigating, the problem usually reveals itself, and you fix it and move on to the next problem.

Unfortunately, performance problems are messier. Here are just a few reasons that network administrators hate performance problems:

- » **Performance problems are difficult to quantify.** Exactly how much slower is the network now than it was a week ago, a month ago, or even a year ago? Sometimes the network just *feels* slow, but you can't quite define exactly how slow it really is.
- » **Performance problems often develop gradually.** Sometimes a network slows down suddenly and drastically. More often, though, the network gradually gets slower, a little bit at a time, until one day its users notice that the network is slow.
- » **Performance problems often go unreported.** Users gripe about the problem to each other around the water cooler, but they don't formally contact you to let you know that the network seems 10 percent slower than usual. As long as they can still access the network, they just assume that the problem is temporary or that they're imagining a problem.
- » **Many performance problems are intermittent.** Sometimes a user calls you and complains that a certain network operation has become slower than molasses, and by the time you get to that person's desk, the operation performs in a snap. Sometimes you can find a pattern to the intermittent behavior, such as it's slower in the morning than in the afternoon or it's slow only while backups are running or while the printer is working. At other times, you can't find a pattern: Sometimes the operation is slow, and sometimes it isn't.
- » **Performance tuning isn't an exact science.** Improving performance sometimes involves educated guesswork. Will upgrading the connection speed between the servers and the network switches improve performance? Probably. Will segmenting the network improve performance? Maybe. Will increasing the RAM on the server improve performance? Hopefully.

» **The solution to a performance problem is sometimes a hard sell.** If a user can't access the network because of a malfunctioning component, the purchase of a replacement is usually undeniably justified. However, if the network is slow and you think that you can fix it by replacing all the switches with faster and more manageable switches, you may have trouble selling management on the upgrade.

What Exactly Is a Bottleneck?

The term *bottleneck* doesn't refer in any way to the physique of the typical computer geek. Rather, computer geeks coined the phrase when they discovered that the tapered shape of a bottle of Jolt cola limited the rate at which they could consume the beverage. "Hey," a computer geek said one day, "the gently tapered narrowness of this bottle's neck imposes a distinct limiting effect upon the rate at which I can consume the tasty caffeine-laden beverage contained within. This observation draws to mind a hitherto undiscovered yet obvious analogy to the limiting effect that a single slow component of a computer system can have upon the performance of the system as a whole."

"Fascinating," replied all the other computer geeks, who were fortunate enough to be present at that historic moment.

The term stuck and is used to this day to draw attention to the simple fact that a computer system is only as fast as its slowest component. It's the computer equivalent of the old truism that a chain is only as strong as its weakest link.

For a simple demonstration of this concept, consider what happens when you print a word-processing document on a slow printer. Your word-processing program reads the data from disk and sends it to the printer. Then you sit and wait while the printer prints the document.

Would buying a faster CPU or adding more memory make the document print faster? No. The CPU is already much faster than the printer, and your computer already has more than enough memory to print the document. The printer itself is the bottleneck, so the only way to print the document faster is to replace the slow printer with a faster one.

Here are some other, random thoughts about bottlenecks:

- » **A computer system always has a bottleneck.** Suppose that you decide that the bottleneck on your file server is its slow 10,000 RPM disk drives, so you replace them with faster 15,000 RPM drives or, better yet, solid state drives. Now the hard drives are no longer the bottleneck: The drives can process information faster than the controller card to which the disks are connected. You didn't really eliminate the bottleneck — you just moved it from the hard drives to the disk controller. No matter what you do, the computer will always have a system that limits the overall performance of the system.
- » **One way to limit the effect of a bottleneck is to avoid waiting for the bottleneck.** For example, print spooling lets you avoid waiting for a slow printer. Although spooling doesn't speed up the printer, it frees you to do other work while the printer chugs along. Similarly, disk caching lets you avoid waiting for a slow hard drive.

The Five Most Common Network Bottlenecks

Direct from the home office in sunny Fresno, California, here are the top five most common network bottlenecks, in no particular order.

The hardware inside your servers

Your servers should be powerful computers capable of handling all the work your network will throw at them. Don't cut corners by using a bottom-of-the-line computer that you bought at a discount computer store.

The following list describes the four most important components of your server hardware:

- » **Processor:** Your server should have a powerful processor. Any processor that's available in a \$500 computer from a low-cost general appliance store is generally not a processor that you want to see in your file server. In other words, avoid processors designed for consumer-grade home computers.
- » **Memory:** You can't have too much memory. Memory is cheap, so don't skimp. Don't even think about running a server with less than 8GB of RAM. And consider giving your most important servers 32GB or even 64GB of RAM.

- » **Disk:** Don't mess around with inexpensive IDE hard drives. To have a respectable system, you should have nothing but SCSI or SAS (Serially Attached SCSI) drives. And if performance of the disk subsystem is critical, opt for fast 15,000 RPM drives or solid state drives.
- » **Network interface:** Nowadays, 1GB network interfaces are standard, so you probably don't have to worry about slow 100 Mbps interfaces being a problem on your network unless you have really old computers or switches. However, modern servers are often equipped with two or more network interfaces. If your servers have just a single interface, consider adding additional interfaces to boost performance

The server's configuration options

All network operating systems have options that you can configure. Some of these options can make the difference between a pokey network and a zippy network. Unfortunately, no hard-and-fast rules exist for setting these options. Otherwise, you wouldn't have options.

The following important tuning options are available for most servers:

- » **Virtual memory options:** *Virtual memory* refers to disk paging files that the server uses when it doesn't have enough real memory to do its work. Few servers ever have enough real memory, so virtual memory is always an important server feature. You can specify the size and location of the virtual memory paging files.



TIP

For the best performance, provide at least as much virtual memory as your computer has real memory. For example, if your server has 16GB of real memory, allocate at least 16GB of virtual memory. If necessary, you can increase this size later.

- » **Disk striping:** Use the disk defragmenter to optimize the data storage on your server's disks.



TIP

If the server has more than one hard drive, you can increase performance by creating *striped volumes*, which allow disk I/O operations to run concurrently on each of the drives in the stripe set.

- » **Network protocols:** Make sure that your network protocols are configured correctly and remove any protocols that aren't necessary.



WARNING

» **Free disk space on the server:** Servers like to have plenty of breathing room on their disks.

If the amount of free disk space on your server drops precipitously low, the server chokes up and slows to a crawl. Make sure that your server has plenty of space — a few gigabytes of unused disk space provides a healthy buffer.

Servers that do too much

One common source of network performance problems is a server overloaded with too many duties. Just because a modern network operating system comes equipped with dozens of different types of services doesn't mean that you should enable and use them all on a single server. If a single server is bogged down because of too much work, add a second server to relieve the first server of some of its chores. Remember the old saying: "Many hands make light work."

For example, if your network needs more disk space, consider adding a second file server rather than adding another drive to the server that already has four nearly full drives. Better yet, purchase a file server appliance dedicated to the task of serving files.

As a side benefit, your network will be easier to administer and more reliable if you place separate functions on separate servers. For example, if a single server doubles as a file server and a mail server, you lose both services if you have to take down the server to perform an upgrade or repair a failed component. However, if you have separate file and mail server computers, only one of the services is interrupted if you have to take down one of the servers.

The network infrastructure

The infrastructure consists of the cables and any switches, hubs, routers, and other components that sit between your clients and your servers.



REMEMBER

The following network infrastructure items can slow down your network:

» **Hubs:** If you have old network hubs rather than switches, replace them with switches immediately. Do not pass go, do not collect \$200.

» **Segment sizes:** Keep the number of computers and other devices on each network segment to a reasonable number. About 20 devices is usually the right number. (Note that if you replace your hubs with switches, you instantly cut the size of each segment because each port on a switch constitutes a separate segment.)

» **The network's speed:** If you have a really old network, you may discover that many — if not all — of your users are still working at 100 Mbps. Upgrading to 1 Gbps speeds up the network dramatically.

» **The backbone speed:** If your network uses a backbone to connect segments, consider upgrading the backbone to 10 Gbps.



TIP

The hardest part about improving the performance of a network is determining where the bottlenecks are. With sophisticated test equipment and years of experience, network gurus can make good educated guesses. Without the equipment and experience, you can still make good uneducated guesses.

Malfunctioning components

Sometimes a malfunctioning network card or other component slows down the network. For example, a switch may malfunction intermittently, occasionally letting packets through but dropping enough of them to slow down the network. After you identify the faulty component, replacing it restores the network to its original speed.

Tune Your Network the Compulsive Way

You can tune your network in one of two ways. The first is to think about it a bit, take a guess at an approach that may improve performance, try that approach, and see whether the network seems to run faster. This strategy is the way most people go about tuning the network.

You can also try the compulsive way, which is suitable for people who organize their sock drawers by color and their food cupboards alphabetically by food group. The compulsive approach to tuning a network goes something like this:

1. Establish a method for objectively testing the performance of some aspect of the network.

In this method, you create a *benchmark*. The result of your benchmark is a *baseline*.

2. Change one variable of your network configuration and rerun the test.

For example, you may think that increasing the size of the disk cache can improve performance. Change the cache size, restart the server, and run the benchmark test. Note whether performance improves, stays the same, or becomes worse.

3. Repeat Step 2 for each variable that you want to test.

Here are some salient points to keep in mind if you decide to tune your network the compulsive way:

- » **If possible, test each variable separately.** In other words, before proceeding, reverse the changes you made to other network variables.
- » **Write down the results of each test so that you have an accurate record of the effect that each change makes on your network's performance.**
- » **Be sure to change only one aspect of the network each time you run the benchmark.** If you make several changes, you don't know which one caused the change. One change may improve performance, but the other change may worsen performance so that the changes cancel each other out — kind of like offsetting penalties in a football game.
- » **If possible, conduct the baseline test during normal working hours, when the network is undergoing its normal workload.**
- » **To establish the network's baseline performance, run the benchmark test two or three times to make sure that the results are repeatable.**

Monitoring Network Performance

One way to monitor network performance is to use a stopwatch to see how long it actually takes to complete common network tasks, such as opening documents or printing reports. If you choose to monitor your network by using the stopwatch technique, you'll want to get a clipboard, baseball cap, and gray sweat suit to complete the ensemble.

A more high-tech approach to monitoring network performance is to use a monitor program that automatically gathers network statistics for you. After you set up the monitor, it plugs away, silently spying on your network and recording what it sees in performance logs. You can then review the performance logs to see how your network is doing.

For large networks, you can purchase sophisticated monitoring programs that run on their own dedicated servers. For small- and medium-sized networks, you can probably get by with the built-in monitoring facilities that come with the network operating system. For example, Figure 21-1 shows the Resource Monitor tool that comes with Windows Server 2016. Other operating systems come with similar tools.

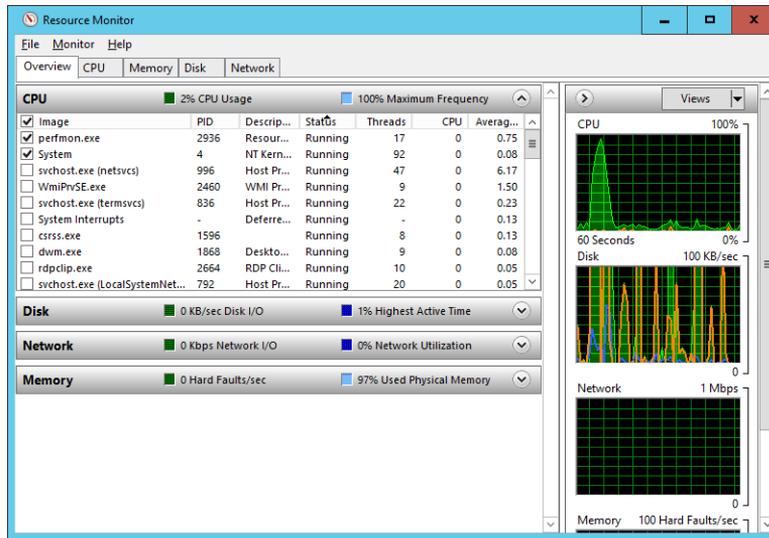


FIGURE 21-1:
Monitoring performance.

Windows Resource Monitor lets you keep track of several different aspects of system performance at once. You track each performance aspect by setting up a counter. You can choose from dozens of different counters. Table 21-1 describes some of the most commonly used counters. Note that each counter refers to a server object, such as physical disk, memory, or the processor.

TABLE 21-1 Commonly Used Performance Counters

Object	Counter	Description
Physical Disk	% Free Space	Percentage of free space on the server's physical disks. Should be at least 15%.
Physical Disk Length	Average Queue	Indicates how many disk operations are waiting while the disk is busy servicing other disk operations. Should be two or fewer.
Memory	Pages/second	Number of pages retrieved from the virtual memory page files per second (pps). A typical threshold is about 2,500 pps.
Processor	% Processor Time	Indicates the percentage of the processor's time that it's busy doing work rather than sitting idle. Should be 85% or less.

Here are a few more things to consider about performance monitoring:

- » **Resource Monitor enables you to view real-time data or to view data that you can save in a log file.** Real-time data gives you an idea about what's happening with the network at a particular moment, but the more useful information comes from the logs.



- » **You can schedule logging to occur at certain times of the day and for certain intervals.** For example, you may schedule the log to gather data every 15 seconds from 9:00 to 9:30 every morning and then again from 3:00 to 3:30 every afternoon.
- » **Even if you don't have a performance problem now, you should set up performance logging and let it run for a few weeks to gather baseline data.** If you develop a problem, this baseline data will prove invaluable while you research the problem.
- » **Don't leave performance logging turned on all the time.** Gathering performance data slows down your server. Use it only occasionally to gather baseline data or when you're experiencing a performance problem.

More Performance Tips

Here are a few last-minute performance tips that barely made it in:

- » **You can often find the source of a slow network by staring at the network switches for a few minutes.** These devices have colorful arrays of green and red lights. The green lights flash whenever data is transmitted; the red lights flash when a collision occurs. An occasional red flash is normal, but if one or more of the red lights is flashing repeatedly, the network interface card (NIC) connected to that port may be faulty. In addition, the lights on most switches indicate the connection speed that has been established for the port. Pay special attention to any ports that are showing a 100 Mbps connection rather than a 1 Gbps connection.
- » **If your switches have a management interface, use it to find trouble spots.** Managed switches can monitor performance and draw your attention to configuration problems or excessive traffic coming over one port. Some switches can even show you statistics about which applications are being used on each port. You may discover that the reason your network always slows down at a certain time of day is that one of your users is streaming his favorite TV show.
- » **Check for scheduled tasks, such as backups, batched database updates, or report jobs.** If at all possible, schedule these tasks to run after normal business hours, such as at night when no one is in the office. These jobs tend to slow down the network by hogging the server's hard drives.

- » **Sometimes, faulty application programs can degrade performance.** For example, some programs develop a *memory leak*. They use memory but then forget to release the memory after they finish. Programs with memory leaks can slowly eat up all the memory on a server, until the server runs out and grinds to a halt. If you think a program has a memory leak, contact the manufacturer of the program to see whether a fix is available.
- » **Spyware can slow a system to a crawl.** A common source of performance problems on client computers is *spyware*, those annoying programs that you almost can't help but pick up when you surf the Internet. Fortunately, you can remove spyware with a variety of free or inexpensive spyware removal tools. For more information, use Google or another search engine to search for spyware removal.

5 More Ways to Network

IN THIS PART . . .

Use cloud computing to move critical network functions out of your server room and onto the Internet.

Manage mobile devices such as smartphones and tablets.

Connect to remote computers.

IN THIS CHAPTER

Examining the basics of cloud computing

Looking at three kinds of cloud computing services

Understanding the pros and cons of cloud computing

Perusing a few major cloud computing service providers

Chapter 22

Life in Cloud City

The world's two most popular science-fiction franchises — *Star Wars* and *Star Trek* — both feature cities that are suspended in the clouds. In *Star Wars Episode V: The Empire Strikes Back*, Han takes the *Millennium Falcon* to Cloud City, hoping that his friend Lando Calrissian can help repair their damaged hyperdrive. And in the original *Star Trek* series episode “The Cloud Minders,” the crew of the *Enterprise* visits a city named Stratos, which is suspended in the clouds.

Coincidence? Perhaps. Or maybe Gene Roddenberry and George Lucas both knew that the future would be in the clouds. At any rate, the future of computer networking is rapidly heading for the clouds. Cloud computing, to be specific. This chapter is a brief introduction to cloud computing. You discover what it is, the pros and cons of adopting it, and what services are provided by the major cloud computer providers.

Introducing Cloud Computing

The basic idea behind cloud computing is to outsource one or more of your networked computing resources to the Internet. “The cloud” represents a new way of

handling common computer tasks. Following are just a few examples of how the cloud way differs from the traditional way:

» Email services

- *Traditional:* Provide email services is to install Microsoft Exchange on a local server computer. Then your clients can connect use Microsoft Outlook to connect to the Exchange server to send and receive email.
- *Cloud:* Contract with an Internet-based email provider, such as Google Mail (Gmail) or Microsoft's Exchange Online. Cloud-based email services typically charge a low monthly per-user fee, so the amount you pay for your email service depends solely on the number of email users you have.

» Disk storage

- *Traditional:* Set up a local file server computer with a large amount of shared disk space.
- *Cloud:* Sign up for an Internet file storage service and then store your data on the Internet. Cloud-based file storage typically charges a small monthly per-gigabyte fee, so you pay only for the storage you use. The disk capacity of cloud-based storage is essentially unlimited.

» Accounting services

- *Traditional:* Purchase expensive accounting software and install it on a local server computer.
- *Cloud:* Sign up for a web-based accounting service. Then all your accounting data is saved and managed on the provider's servers, not on yours.

Looking at the Benefits of Cloud Computing

Cloud computing is a different — and, in many ways, better — approach to networking. Here are a few of the main benefits of moving to cloud-based networking:

- » **Cost-effective:** Cloud-based computing typically is less expensive than traditional computing. Consider a typical file server application: To implement a file server, first you must purchase a file server computer with enough disk space to accommodate your users' needs, which amounts to 1TB of disk storage. You want the most reliable data storage possible, so you purchase a server-quality computer and fully redundant disk drives. For the sake of this

discussion, figure that the total price of the server — including its disk drive, the operating system license, and the labor cost of setting it up — is about \$10,000. Assuming that the server will last for four years, that totals about \$2,500 per year.

If you instead acquire your disk storage from a cloud-based file sharing service, you can expect to pay about one fourth of that amount for an equivalent amount of storage.

The same economies apply to most other cloud-based solutions. Cloud-based email solutions, for example, typically cost around \$5 per month per user — well less than the cost of setting up and maintaining an on-premises Microsoft Exchange Server.

» **Scalable:** So what happens if you guess wrong about the storage requirements of your file server, and your users end up needing 2TB instead of just 1TB? With a traditional file server, you must purchase additional disk drives to accommodate the extra space. Sooner than you want, you'll run out of capacity in the server's cabinet. Then you'll have to purchase an external storage cabinet. Eventually, you'll fill that up, too.

Now suppose that after you expand your server capacity to 2TB, your users' needs contract to just 1TB. Unfortunately, you can't return disk drives for a refund.

With cloud computing, you pay only for the capacity you're actually using, and you can add capacity whenever you need it. In the file server example, you can write as much data as you need to the cloud storage. Each month, you're billed according to your actual usage. Thus, you don't have to purchase and install additional disk drives to add storage capacity.

» **Reliable:** Especially for smaller businesses, cloud services are much more reliable than in-house services. Just a week before I wrote this chapter, the tape drive that a friend uses to back up his company's data failed. As a result, he was unable to back up data for three days while the tape drive was repaired. Had he been using cloud-based backup, he could have restored his data immediately and wouldn't have been without backups for those four days.

The reason for the increased reliability of cloud services is simply a matter of scale. Most small businesses can't afford the redundancies needed to make their computer operations as reliable as possible. My friend's company can't afford to buy two tape drives so that an extra is available in case the main one fails.

By contrast, cloud services are usually provided by large companies such as Amazon, Google, Microsoft, and IBM. These companies have state-of-the-art data centers with multiple redundancies for their cloud services. Cloud storage may be kept on multiple servers so that if one server fails, others can



REMEMBER

take over the load. In some cases, these servers are in different data centers in different parts of the country. Thus, your data will still be available even in the event of a disaster that shuts down an entire data system.

- » **Hassle-free:** Face it, IT can be a hassle. With cloud-based services, you basically outsource the job of complex system maintenance chores, such as software upgrade, patches, hardware maintenance, backup, and so on. You get to consume the services while someone else takes care of making sure that the services run properly.
- » **Globally accessible:** One of the best things about cloud services is that they're available anywhere you have an Internet connection. Suppose that you have offices in five cities. Using traditional computing, each office would require its own servers, and you'd have to carefully design systems that allowed users in each of the offices to access shared data.

With cloud computing, each office simply connects to the Internet to access the cloud applications. Cloud-based applications are also great if your users are mobile because they can access the applications anywhere they can find an Internet connection.

Detailing the Drawbacks of Cloud Computing

Although cloud computing has many advantages over traditional techniques, it isn't without its drawbacks. Here are some of the most significant roadblocks to adopting cloud computing:

- » **Entrenched applications:** Your organization may depend on entrenched applications that don't lend themselves especially well to cloud computing — or that at least require significant conversion efforts to migrate to the cloud. For example, you might have use an accounting system that relies on local file storage.

Fortunately, many cloud providers offer assistance with this migration. And in many cases, the same application that you run locally can be run in the cloud, so no conversion is necessary.

- » **Internet connection speed:** Cloud computing shifts much of the burden of your network to your Internet connection. Your users used to access their data on local file servers over gigabit-speed connections; now they must access data over slower bandwidth Internet connections.



Although you can upgrade your connection to higher speeds, doing so will cost money — money that may well offset the money you otherwise save from migrating to the cloud.

» **Internet connection reliability:** The cloud resources you access may feature all the redundancy in the world, but if your users access the cloud through a single Internet connection, that connection becomes a key point of vulnerability. Should it fail, any applications that depend on the cloud will be unavailable. If those applications are mission-critical, business will come to a halt until the connection is restored.

Here are two ways to mitigate this risk:

- *Make sure that you're using an enterprise-class Internet connection.* Enterprise-class connections are more expensive but provide much better fault tolerance and repair service than consumer-class connections do.
- *Provide redundant connections if you can.* That way, if one connection fails, traffic can be rerouted through alternative connections.

» **Security threats:** You can bet your life that hackers throughout the world are continually probing for ways to break through the security perimeter of all the major cloud providers. When they do, your data may be exposed.

The best way to mitigate this threat is to ensure that strong password policies are enforced.

Examining Three Basic Kinds of Cloud Services

Three distinct kinds of services can be provided via the cloud: applications, platforms, and services (infrastructure). The following paragraphs describe these three types of cloud services in greater detail.

Applications

Most often referred to as *Software as a Service* (SaaS), fully functional applications can be delivered via the cloud. One of the best-known examples is *Google Apps*, which is a suite of cloud-based office applications designed to compete directly with Microsoft's traditional office applications, including Word, Excel, PowerPoint, Access, and Outlook. Google Apps can also replace the back-end software often used to support Microsoft Office, including Exchange and SharePoint.

When you use a cloud-based application, you don't have to worry about any of the details that are commonly associated with running an application on your network, such as deploying the application and applying product upgrades and software patches. Cloud-based applications usually charge a small monthly fee based on the number of users running the software, so costs are low.

Also, as a cloud-based application user, you don't have to worry about providing the hardware or operating system platform on which the application will run. The application provider takes care of that detail for you, so you can focus simply on developing the application to best serve your users' needs.

Platforms

Also referred to as *Platform as a Service* (PaaS), this class of service refers to providers that give you access to a remote virtual operating platform on which you can build your own applications.

At the simplest level, a PaaS provider gives you a complete, functional remote virtual machine that's fully configured and ready for you to deploy your applications to. If you use a web provider to host your company's website, you're already using PaaS: Most web host providers give you a functioning Linux system, fully configured with all the necessary servers, such as Apache or MySQL. All you have to do is build and deploy your web application on the provider's server.

More-complex PaaS solutions include specialized software that your custom applications can tap to provide services such as data storage, online order processing, and credit card payments. One of the best-known examples of this type of PaaS provider is Amazon.



REMEMBER

When you use PaaS, you take on the responsibility of developing your own custom applications to run on the remote platform. The PaaS provider takes care of the details of maintaining the platform itself, including the base operating system and the hardware on which the platform runs.

Infrastructure

If you don't want to delegate the responsibility of maintaining operating systems and other elements of the platform, you can use *Infrastructure as a Service* (IaaS). When you use IaaS, you're purchasing raw computing power that's accessible via the cloud. Typically, IaaS provides you access to a remote virtual machine. It's up to you to manage and configure the remote machine however you want.

Public Clouds versus Private Clouds

The most common form of cloud computing uses what is known as a *public cloud* — that is, cloud services that are available to anyone in the world via the Internet. Google Apps is an excellent example of a public cloud service. Anyone with access to the Internet can access the public cloud services of Google Apps: Just point your browser to <http://apps.google.com>.

A public cloud is like a public utility, in that anyone can subscribe to it on a pay-as-you-go basis. One of the drawbacks of public cloud services is that they're inherently insecure. When you use a public cloud service, you're entrusting your valuable data to a third party that you cannot control. Sure, you can protect your access to your public cloud services by using strong passwords, but if your account names and passwords are compromised, your public cloud services can be hacked into, and your data can be stolen. Every so often, we all hear news stories about how this company's or that company's back-door security has been compromised.

Besides security, another drawback of public cloud computing is that it's dependent on high-speed, reliable Internet connections. Your cloud service provider may have all the redundancy in the world, but if your connection to the Internet goes down, you won't be able to access your cloud services. And if your connection is slow, your cloud services will be slow.

A *private cloud* mimics many of the features of cloud computing but is implemented on a private hardware within a local network, so it isn't accessible to the general public. Private clouds are inherently more secure because the general public can't access them. Also, they're dependent only on private network connections, so they aren't subject to the limits of a public Internet connection.



TIP

As a rule, private clouds are implemented by large organizations that have the resources available to create and maintain their own cloud servers.

A relative newcomer to the cloud computing scene is the *hybrid cloud*, which combines the features of public and private clouds. Typically, a hybrid cloud system uses a small private cloud that provides local access to the some of the applications and the public cloud for others. You might maintain your most frequently used data on a private cloud for fast access via the local network and use the public cloud to store archives and other less frequently used data, for which performance isn't as much of an issue.

Introducing Some of the Major Cloud Providers

Hundreds, if not thousands, of companies provide cloud services. Most of the cloud computing done today, however, is provided by just a few providers, which are described in the following sections.

Amazon

By far the largest provider of cloud services in the world is Amazon. Amazon launched its cloud platform — Amazon Web Services (AWS) — in 2006. Since then, hundreds of thousands of customers have signed up. Some of the most notable users of AWS include Netflix, Pinterest, and Instagram.

AWS includes the following features:

- » **Amazon CloudFront:** A PaaS content-delivery system designed to deliver web content to large numbers of users.
- » **Amazon Elastic Compute Cloud:** Also called Amazon EC2. An IaaS system that provides access to raw computing power.
- » **Amazon Simple Storage Service:** Also called Amazon S3. Provides web-based data storage for unlimited amounts of data.
- » **Amazon Simple Queue Service:** Also called Amazon SQS. Provides a data transfer system that lets applications send messages to other applications. SQS enables you to build applications that work together.
- » **Amazon Virtual Private Cloud:** Also called Amazon VPC. Uses virtual private network (VPN) connections to connect your local network to Amazon's cloud services.

Google

Google is also one of the largest providers of cloud services. Its offerings include the following:

- » **Google Apps:** A replacement for Microsoft Office that provides basic email, word processing, spreadsheet, and database functions via the cloud. Google Apps is free to the general public and can even be used free by small business (up to 50 users). For larger businesses, Google offers an advanced version, Google Apps for Business. For \$5 per month per user, you get extra features,

such as 25GB of email data per user, archiving, and advanced options for customizing your account policies.

- » **Google Cloud Connect:** A cloud-based solution that lets you work with Google cloud data directly from within Microsoft Office applications.
- » **Google App Engine:** A PaaS interface that lets you develop your own applications that work with Google's cloud services.
- » **Google Cloud Print:** Allows you to connect your printers to the cloud so that they can be accessed from anywhere.
- » **Google Maps:** A Global Information System (GIS).

Microsoft

Microsoft has its own cloud strategy, designed in part to protect its core business of operating systems and Office applications against competition from other cloud providers, such as Google Apps.

The following paragraphs summarize several of Microsoft's cloud offerings:

- » **Microsoft Office 365:** A cloud-based version of Microsoft Office. According to Microsoft's website, Office 365 provides "anywhere access to cloud-based email, web conferencing, file sharing, and Office Web Apps at a low predictable monthly cost." For more information, check out www.office365.com.
- » **Windows Azure:** A PaaS offering that lets you build websites, deploy virtual machines that run Windows Server or Linux, or access cloud versions of server applications such as SQL Server.
- » **Microsoft Business Productivity Suite:** A SaaS product that provides cloud-based access to two of Microsoft's most popular productivity servers: Microsoft Exchange and Microsoft SharePoint. The suite lets you deploy these servers without having to create and maintain your own local servers.

Getting Into the Cloud

After you wrap your head around just how cool cloud computing can be, what should you do to take your network toward the cloud? Allow me to make a few recommendations:

- » **Don't depend on a poor Internet connection.** First and foremost, before you take any of your network operations to the cloud, make sure that you're

not dependent on a consumer-grade Internet connection if you decide to adopt cloud computing. Consumer-grade Internet connections can be fast, but when an outage occurs, there's no telling how long you'll wait for the connection to be repaired. You definitely don't want to wait for hours or days while the cable company thinks about sending someone out to your site. Instead, spend the money for a high-speed enterprise-class connection that can scale as your dependence on it increases.

» **Assess what applications you may already have running on the cloud.**

If you use Gmail rather than Exchange for your email, congratulations! You've already embraced the cloud. Other examples of cloud services that you may already be using include a remote web or FTP host, Dropbox or another file sharing service, Carbonite or another online backup service, a payroll service, and so on.

» **Don't move to the cloud all at once.** Start by identifying a single application that lends itself to the cloud. If your engineering firm archives projects when they close and wants to get them off your primary file server but keep them readily available, look to the cloud for a file storage service.

» **Go with a reputable company.** Google, Amazon, and Microsoft are all huge companies with proven track records in cloud computing. Many other large and established companies also offer cloud services. Don't stake your company's future on a company that didn't exist six months ago.

» **Research, research, research.** Pour yourself into the web, and buy a few books. *Hybrid Cloud For Dummies*, by Judith Hurwitz, Marcia Kaufman, Dr. Fern Halper, and Daniel Kirsch (Wiley), is a good place to start.

Chapter 23

Managing Mobile Devices

A computer consultant once purchased a used BlackBerry device on eBay for \$15.50. When he put in a new battery and turned on the device, he discovered that it contained confidential emails and personal contact information for executives of a well-known financial institution.

Oops!

It turns out that a former executive with the company sold his old BlackBerry on eBay a few months after he left the firm. He'd assumed that because he'd removed the battery, everything on the BlackBerry had been erased.

The point of this true story is that mobile devices such as smartphones and tablet computers pose a special set of challenges for network administrators. These challenges are now being faced even by administrators of small networks. Just a few years ago, only large companies had BlackBerry or other mobile devices that integrated with Exchange email, for example. Now it isn't uncommon for companies with just a few employees to have mobile devices connected to the company network.

This chapter is a brief introduction to mobile devices and the operating systems they run, with an emphasis on iPhone and Android devices. You find out more about how these devices can interact with Exchange email and the steps you can take to ensure their security.

The Many Types of Mobile Devices

Once upon a time, there were mobile phones and PDAs. A mobile phone was just that: a handheld telephone you could take with you. The good ones had nice features such as a call log, an address book, and perhaps a crude game, but not much else. PDAs — *Personal Digital Assistants* — were little handheld computers designed to replace the old-fashioned Day-Timer books people used to carry around with them to keep track of their appointment calendars and address books.

All that changed when cellular providers began adding data capabilities to their networks. Now cellphones can have complete mobile Internet access. This fact has resulted in the addition of sophisticated PDA features to mobile phones and phone features to PDAs so that the distinctions are blurred.

A *mobile device* can be any one of a wide assortment of devices that you can hold in one hand and that are connected through a wireless network. The term *handheld* is a similar generic name for such devices. The following list describes some of the most common specifics of mobile devices:

- » **Mobile phone:** Primary purpose is to enable phone service. Most mobile phones also include text messaging, address books, appointment calendars, and games; they may also provide Internet access.
- » **Smartphone:** A smartphone is a cellphone that also functions as a handheld computer. Smartphones feature touchscreens instead of physical buttons or keys to press. Besides the features ordinarily found on a mobile phone, smartphones also offer email, calendar, contacts, task lists, and web access, as well as apps that can be purchased and installed on the phone.
- » **Android:** Android is an open-source operating system (OS) for smartphones, developed by Google. Android is far and away the most popular platform for smartphones, being used on more than 80 percent of the smartphones sold since 2015.

- » **iOS:** iOS is the OS used on Apple's popular iPhone and iPad mobile devices. Although outnumbered by Android devices, many people consider iOS devices to be more innovative than Android devices. The main thing that holds iOS back in market share is cost: Apple devices are considerably more expensive than their Android equivalents.
- » **BlackBerry:** BlackBerry was once the king of the smartphone game. For many years, BlackBerry had a virtual monopoly on the mobile devices market because it was the first mobile device that could synchronize well with Microsoft Exchange. Now that Android and Apple devices do that just as well as (actually, much better than) BlackBerry, BlackBerry devices have fallen out of vogue. However, BlackBerry is still around and there are still plenty of BlackBerry users out there. (Note that newer BlackBerry phones run Android rather than the old proprietary BlackBerry OS.)

Considering Security for Mobile Devices

As a network administrator, one of your main responsibilities regarding mobile devices is to keep them secure. Unfortunately, that's a significant challenge. Here are some reasons why:

- » **Mobile devices connect to your network via other networks that are out of your control.** You can go to great lengths to set up firewalls, encryption, and a host of other security features, but mobile devices connect via public networks whose administrators may not be as conscientious as you.
- » **Mobile devices are easy to lose.** A user might leave her smartphone at a restaurant or hotel, or it might fall out of her pocket on the subway.
- » **Mobile devices run operating systems that aren't as security conscious as Windows.**
- » **Users who wouldn't dare install renegade software on their desktop computers think nothing of downloading free games or other applications to their handheld devices.** Who knows what kinds of viruses or Trojans these downloads carry?
- » **Inevitably, someone will buy his own handheld device and connect it to your network without your knowledge or permission.**

Here are some recommendations for beefing up security for your mobile devices:

- » Establish clear, consistent policies for mobile devices, and enforce them.
- » Make sure employees understand that they aren't allowed to bring their own devices into your network. Allow only company-owned devices to connect.
- » Train your users in the security risks associated with using mobile devices.
- » Implement antivirus protection for your mobile devices.

Managing iOS Devices

In 2007, the Apple iPhone, one of the most innovative little gadgets in many, many years, hit the technology market. In just a few short years, the iPhone captured a huge slice of a market previously dominated almost exclusively by RIM and its BlackBerry devices. Since then, the iPhone's share of the mobile-phone market has grown beyond that of the former king, BlackBerry.

The success of the iPhone was due in large part to the genius of its operating system, called iOS. In 2010, Apple released the iPad, a tablet computer that runs the same iOS as the iPhone. And in 2012, Apple introduced a smaller version of the iPad: the iPad mini. Together, these devices are commonly known as *iOS devices*.

Understanding the iPhone

The iPhone is essentially a combination of four devices:

- » A cellphone
- » An iPod with a memory capacity of 16GB to 128GB
- » A digital camera
- » An Internet device with its own web browser (Safari) and applications, such as email, calendar, and contact management

The most immediately noticeable feature of the iPhone is its lack of a keyboard. Instead, nearly the entire front surface of the iPhone is a high-resolution, touch-sensitive LCD display. The display is not only the main output device of the iPhone, but also its main input device. The display can become a keypad input for dialing a telephone number or a keyboard for entering text. You can also use various finger gestures, such as tapping icons to start programs or pinching to zoom in the display.

The iPhone has several other innovative features:

- » An *accelerometer* tracks the motion of the iPhone in three directions. The main use of the accelerometer is to adjust the orientation of the display from landscape to portrait based on how the user is holding the phone. Some other applications — mostly games — use the accelerometer as well.
- » A Wi-Fi interface lets the iPhone connect to local Wi-Fi networks for faster Internet access.
- » GPS capability provides location awareness for many applications, including Google Maps.
- » The virtual private network (VPN) client lets you connect to your internal network.

Of all the unique features of the iPhone, probably the most important is its huge collection of third-party applications that can be downloaded from a special web portal, the App Store. Many of these applications are free or cost just a few dollars. (Many are just 99 cents or \$1.99.) As of this writing, more than 1.5 million applications — everything from business productivity to games — were available from the App Store.

Understanding the iPad

The iPad is essentially an iPhone without the phone but with a larger screen. The iPhone comes with a 3.5-inch screen; the iPad has a 9.7-inch screen; and its smaller cousin, the iPad mini, has a 7.9-inch screen and the iPad Pro has a whopping 12.9-inch screen.

Apart from these basic differences, an iPad is nearly identical to an iPhone. Any application that can run on an iPhone can also run on an iPad, and many applications are designed to take special advantage of the iPad's larger screen.

All the information that follows in this chapter applies equally to iPhones and iPads.

Integrating iOS devices with Exchange

An iOS device can integrate with Microsoft Exchange email via the Exchange ActiveSync feature, which is enabled by default on Exchange 2010 and later versions.

To verify the Exchange ActiveSync feature for an individual mailbox, follow these steps:

1. **Choose Start ⇨ Administrative Tools ⇨ Active Directory Users and Computers.**

The Active Directory Users and Computers console opens.

2. **Expand the domain and then locate the user you want to enable mobile access for.**
3. **Right-click the user and then choose Properties from the contextual menu.**
4. **Click the Exchange Features tab.**

The Exchange Features options are displayed, as shown in Figure 23-1.

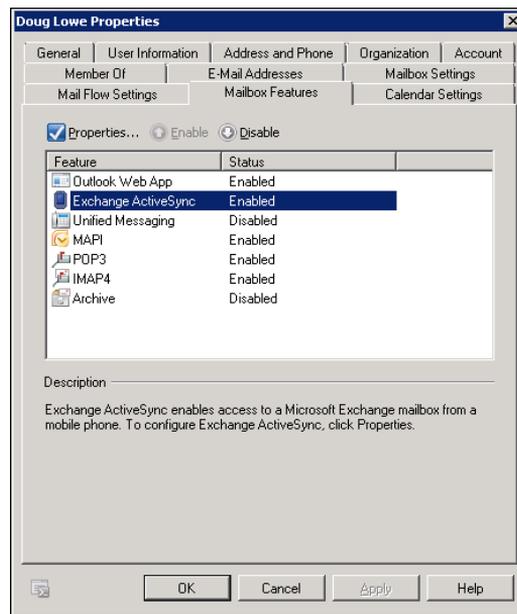


FIGURE 23-1:
Enabling
Exchange
ActiveSync for
a user.

5. **Ensure that the Exchange ActiveSync option is enabled.**

If the options aren't already enabled, right-click each option and choose Enable from the contextual menu.

6. **Click OK.**
7. **Repeat Steps 5 and 6 for any other users you want to enable mobile access for.**

8. Close Active Directory Users and Computers.

That's all there is to it. After you enable these features, any users running Windows Mobile can synchronize their handheld devices with their Exchange mailboxes.

Configuring an iOS device for Exchange email

After ActiveSync is enabled for the mailbox, you can configure an iPhone or iPad to tap into the Exchange account by following these steps:

1. On the iPhone or iPad, tap Settings and then tap Mail, Contacts, Calendars.

The screen shown in Figure 23-2 appears. This screen lists any existing email accounts that may already be configured on the phone and also lets you add a new account.

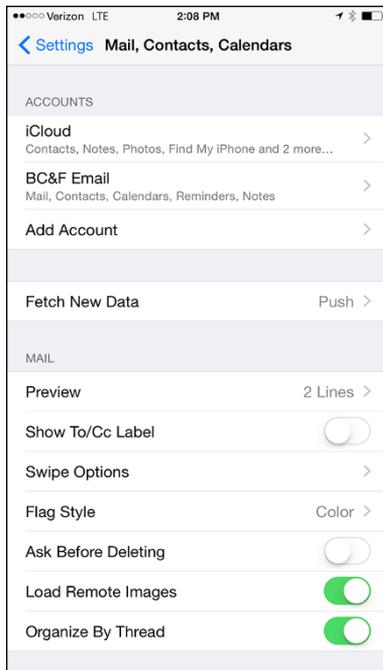


FIGURE 23-2:
Adding an email account.

2. Tap Add Account.

The screen shown in Figure 23-3 appears, allowing you to choose the type of email account you want to add.

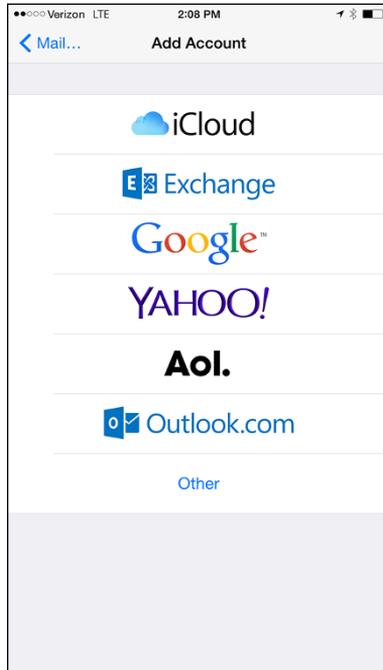


FIGURE 23-3:
The iPhone can support many types of email accounts.

3. Tap Exchange.

The screen shown in Figure 23-4 appears, where you can enter basic information for your Exchange account.

4. Enter your email address, password, and a description of the account.

5. Tap Next.

The screen shown in Figure 23-5 appears.

6. Enter either the DNS name or the IP address of your Exchange server in the Server field.

I entered **smtp.lowewriter.com** for my Exchange server, for example.

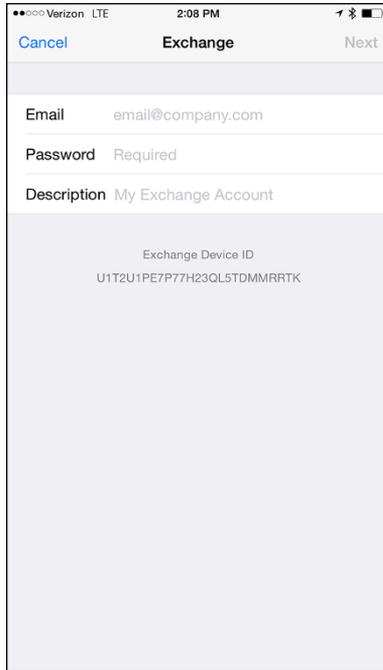


FIGURE 23-4:
Enter your email
address and
password.

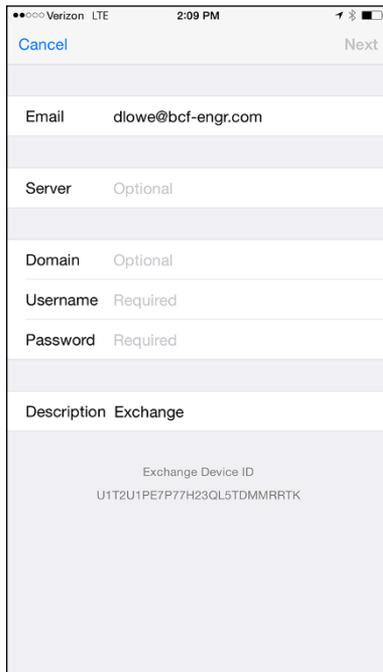


FIGURE 23-5:
Enter your
Exchange server
information.

7. Enter your domain name, your Windows username, and your Windows password in the appropriate fields.

8. Tap Next.

The screen shown in Figure 23-6 appears. Here, you select which mailbox features you want to synchronize: Mail, Contacts, Calendars, Reminders, and/or Notes.

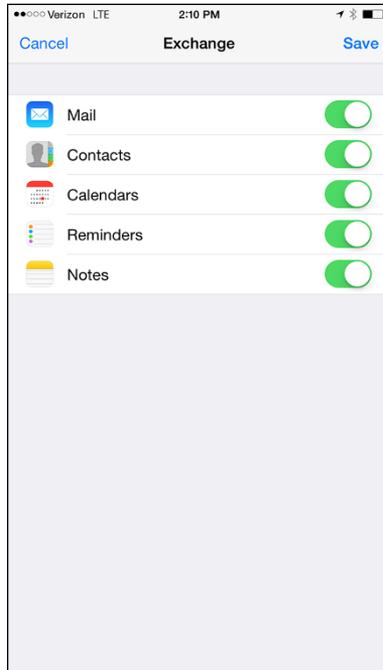


FIGURE 23-6: Select features to synchronize.

9. Select the features you want to synchronize and then tap Done.

The email account is created.

After the email account has been configured, the user can access it via the Mail icon on the iPhone's home screen.

Managing Android Devices

This section is a brief introduction to the Android platform. You find out a bit about what Android actually is, and you discover the procedures for setting up Exchange email access on an Android phone.

Crucial differences exist between Android phones and iPhones. The most important difference — in many ways, the *only* important difference — is that Android phones are based on an open source OS derived from Linux, which can be extended and adapted to work on a wide variety of hardware devices from different vendors. With the iPhone, you're locked into Apple hardware. With an Android phone, though, you can buy hardware from a variety of manufacturers.

Looking at the Android OS

Most people associate the Android OS with Google, and it's true that Google is the driving force behind Android. The Android OS is an open source OS managed by the Open Handset Alliance (OHA). Google still plays a major role in the development of Android, but more than 50 companies are involved in the OHA, including hardware manufacturers (such as HTC, Intel, and Motorola), software companies (such as Google and eBay), and mobile-phone operators (such as T-Mobile and Sprint-Nextel).



TECHNICAL
STUFF

Technically speaking, Android is more than just an OS. It's also a complete *software stack*, which comprises several key components that work together to create the complete Android platform:

- » **The OS core**, which is based on the popular Linux OS
- » **A middleware layer**, which provides drivers and other support code to enable the OS core to work with the hardware devices that make up a complete phone, such as a touch-sensitive display, the cellphone radio, the speaker and microphone, Bluetooth or Wi-Fi networking components, and so on
- » **A set of core applications** that the user interacts with to make phone calls, read email, send text messages, take pictures, and so on
- » **A Software Developers Kit (SDK)** that lets third-party software developers create their own applications to run on an Android phone, as well as a marketplace where the applications can be marketed and sold, much as the App Store lets iPhone developers market and sell applications for the iPhone

Besides the basic features provided by all operating systems, here are a few bonus features of the Android software stack:

- » An optimized graphical display engine that can produce sophisticated 2-D and 3-D graphics
- » GPS capabilities that provide location awareness that can be integrated with applications such as Google Maps
- » Compass and accelerometer capabilities that can determine whether the phone is in motion and in which direction it's pointed
- » A built-in SQL database server for data storage
- » Support for several network technologies, including 3G, 4G, Bluetooth, and Wi-Fi
- » Built-in media support, including common formats for still images, audio, and video files

Perusing Android's core applications

The Android OS comes preconfigured with several standard applications, which provide the functionality that most people demand from a modern smartphone. These applications include

- » **Dialer:** Provides the basic cellphone function that lets users make calls.
- » **Browser:** A built-in web browser that's similar to Google's Chrome browser.
- » **Messaging:** Provides text (SMS) and multimedia (MMS) messaging.
- » **Email:** A basic email client that works best with Google's Gmail but that can be configured to work with other email servers, including Exchange.
- » **Contacts:** Provides a contacts list that integrates with the Dialer and Email applications.
- » **Camera:** Lets you use the phone's camera hardware (if any) to take pictures.
- » **Calculator:** A simple calculator application.
- » **Alarm Clock:** A basic alarm clock. You can set up to three different alarms.
- » **Maps:** An integrated version of Google Maps.
- » **YouTube:** An integrated version of YouTube.

- » **Music:** An MP3 player similar to the iPod. You can purchase and download music files from Amazon.
- » **Google Play:** Lets you purchase and download third-party applications for the Android phone.
- » **Settings:** Lets you control various settings for the phone.

Integrating Android with Exchange

The Android's core Email application can integrate with Microsoft Exchange email. To do that, you must enable Exchange Mobile Services and then enable ActiveSync for the user's mailbox.

After you enable Exchange Mobile Services and ActiveSync on your Exchange server, you can easily configure the Android phone for email access. Just run the Email application on the Android phone, and follow the configuration steps, which ask you for basic information such as your email address, username, password, and Exchange mail server.

Chapter 24

Connecting from Home

A typical computer user takes work home to work on in the evening or over the weekend and bring back to the office the following weekday. This arrangement can work okay, except that exchanging information between your home computer and your office computer isn't easy.

One way to exchange files is to mark them for offline access, as I describe in Chapter 3. However, this approach has its drawbacks. What if someone goes to the office on Saturday and modifies the same file you're working on at home? What if you get home and discover that the file you need is on a folder you didn't mark for offline access?

What about email? Offline access doesn't give you access to your company email account, so you can't check whether you have mail in your Inbox or send mail from your company email account.

This chapter introduces two features that can alleviate these problems. The first is Internet-based access to your email via Outlook Web App (OWA) in Microsoft Exchange. The second is the *virtual private network* (VPN), which lets you connect to your network from home as though you were at work so that you can safely access all your network resources as though you were locally connected to the network.

Using Outlook Web App

Most people who connect to their office networks from home really just need their email. If the only reason for accessing the office network is to get email, try this simple, easy tool: Outlook Web App, also known as OWA. This Microsoft Exchange Server feature can access your company email from any computer that has an Internet connection. The remote computer just needs a web browser and an Internet connection; no VPN or other special configuration is required.

The best part is that you don't have to do anything special to enable OWA; it's enabled by default when you install Microsoft Exchange. Although you can configure plenty of options to improve its use, OWA is functional right out of the box.

To access OWA from any web browser, just browse to the address designated for your organization's OWA. The default address is the DNS name of your mail server, followed by /exchange. For example, for the mail server `smtp.lowewriter.com`, the OWA address is `smtp.lowewriter.com/exchange`.



TECHNICAL
STUFF

The connection must use the secure version of the normal HTTP web protocol. You must type **https://** before the OWA address. The complete address will be something like `https://smtp.lowewriter.com/exchange`.

When you browse to your OWA address, you're prompted to enter a name and password. Use your regular network logon name and password. OWA appears in the browser window, as shown in Figure 24-1.

If you're familiar with Outlook, you'll have no trouble using OWA. Almost all Outlook features are available, including your inbox, calendar, contacts, tasks, reminders, and even public folders. You can even set up an Out of Office reply.

One difference between OWA and Outlook is that there's no menu bar across the top. However, most of the functions that are available from the menu bar are available elsewhere in OWA. If you can't find a feature, look in the Options page, which you can reach by clicking Options at the bottom left of the window. Figure 24-2 shows the Options page. From here, you can create an Out of Office reply, set your signature, and change a variety of other options.

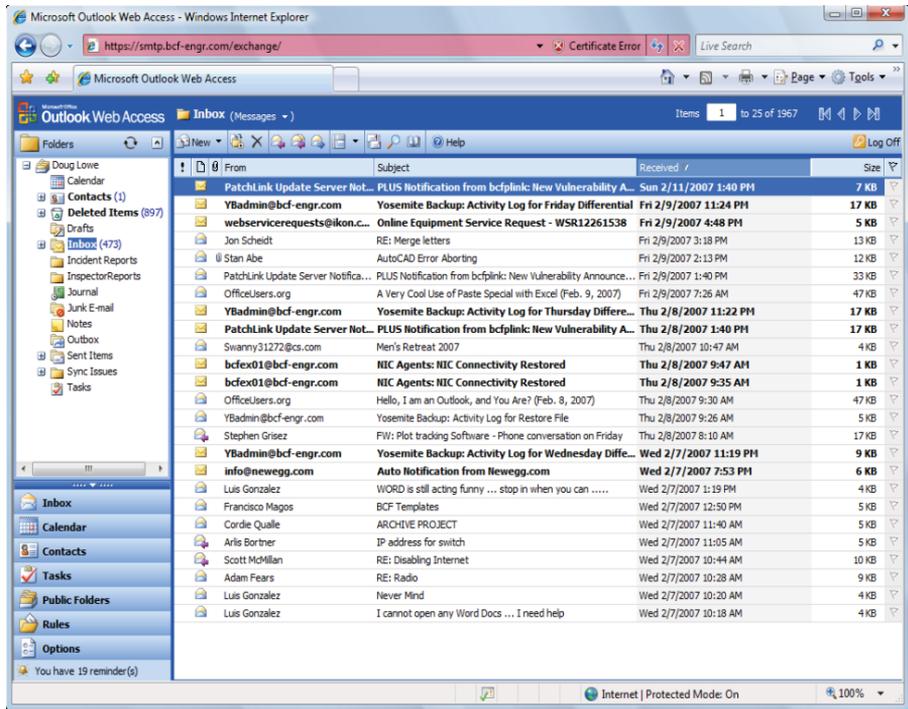


FIGURE 24-1:
OWA looks a lot like Outlook.

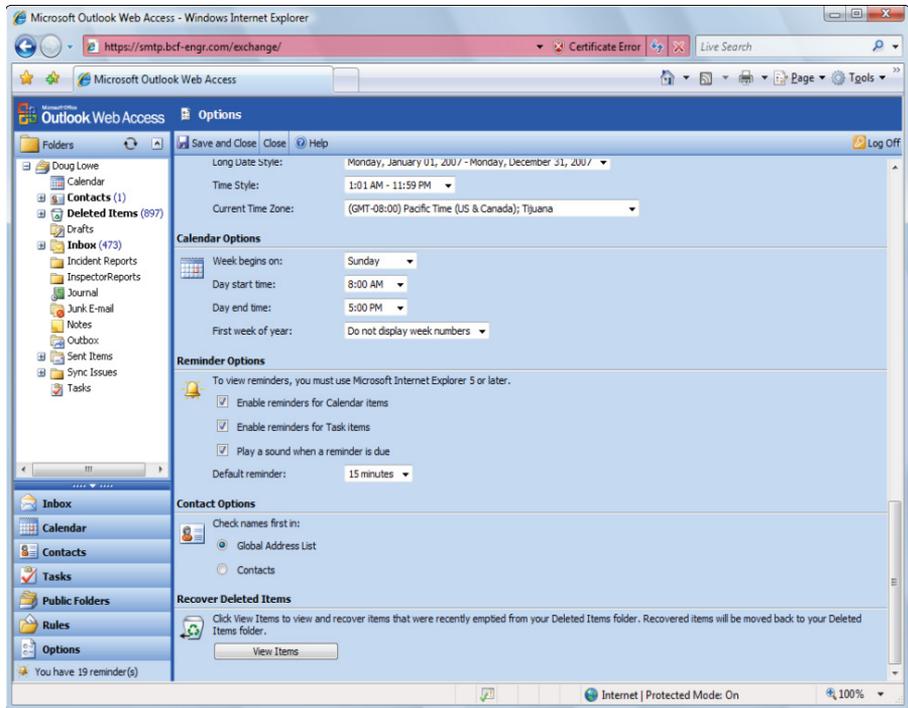


FIGURE 24-2:
Set OWA options here.

Using a Virtual Private Network

A *virtual private network* (VPN) is a type of network connection that creates the illusion that you're directly connected to a network when in fact, you're not. For example, suppose you set up a LAN at your office, but you also occasionally work from home. But how will you access the files on your work computer from home?

- » You could simply copy whatever files you need from your work computer onto a flash drive and take them home with you, work on the files, copy the updated files back to the flash drive, and take them back to work with you the next day.
- » You could email the files to your personal email account, work on them at home, and then email the changed files back to your work email account.
- » You could get a laptop and use the Windows Offline Files feature to automatically synchronize files from your work network with files on the laptop.

Or you could set up a VPN that allows you to log on to your work network from home. The VPN uses a secured Internet connection to connect you directly to your work network, so you can access your network files as if you had a really long Ethernet cable that ran from your home computer all the way to the office and plugged directly into the work network.

Here are at least three situations in which a VPN is the ideal solution:

- » Workers need to occasionally work from home (as in the scenario just described). In this situation, a VPN connection establishes a connection between the home computer and the office network.
- » Mobile users — who may not ever actually show up at the office — need to connect to the work network from mobile computers, often from locations like hotel rooms, clients' offices, airports, or coffee shops. This type of VPN configuration is similar to the home user's configuration except that the exact location of the remote user's computer is not fixed.
- » Your company has offices in two or more locations, each with its own LAN, and you want to connect the locations so that users on either network can access each other's network resources. In this situation, the VPN doesn't connect a single user with a remote network; instead, it connects two remote networks to each other.

Looking at VPN security

The *V* in VPN stands for *virtual*, which means that a VPN creates the appearance of a local network connection when in fact the connection is made over a public network — the Internet. The term *tunnel* is sometimes used to describe a VPN because the VPN creates a tunnel between two locations, which can be entered only from either end. The data that travels through the tunnel from one end to the other is secure as long as it's within the tunnel — that is, within the protection provided by the VPN.

The *P* in VPN stands for *private*, which is the purpose of creating the tunnel. If the VPN didn't create effective security so that data can enter the tunnel only at one of the two ends, the VPN would be worthless; you may as well just open your network and your remote computer up to the Internet and let the hackers have their way.

Prior to VPN technology, the only way to provide private remote network connections was through direct-dial lines or dedicated private lines, which were (and still are) very expensive. For example, to set up a remote office, you could lease a private T1 line from the phone company to connect the two offices. This private T1 line provided excellent security because it physically connected the two offices and could be accessed only from the two endpoints.

VPN provides the same point-to-point connection as a private leased line, but does it over the Internet instead of through expensive dedicated lines. To create the tunnel that guarantees privacy of the data as it travels from one end of the VPN to the other, the data is encrypted using special security protocols.

The most important of the VPN security protocols is *Internet Protocol Security* (IPSec), which is a collection of standards for encrypting and authenticating packets that travel on the Internet. In other words, it provides a way to encrypt the contents of a data packet so that only a person who knows the secret encryption keys can decode the data. And it provides a way to reliably identify the source of a packet so that the parties at either end of the VPN tunnel can trust that the packets are authentic.

Another commonly used VPN protocol is Layer 2 Tunneling Protocol (L2TP). This protocol doesn't provide data encryption. Instead, it's designed to create end-to-end connections — *tunnels* — through which data can travel. L2TP is actually a combination of two older protocols: Layer 2 Forwarding Protocol (L2FP, from Cisco), and Point-to-Point Tunneling Protocol (PPTP, from Microsoft).

Many VPNs today use a combination of L2TP and IPSec: L2TP over IPSec. This type of VPN combines the best features of L2TP and IPSec to provide a high degree of security and reliability.

Understanding VPN servers and clients

A VPN connection requires a VPN *server* — the gatekeeper at one end of the tunnel — and a VPN client at the other end. The main difference between the server and the client is that the client initiates the connection with the server, and a VPN client can establish a connection with just one server at a time. However, a server can accept connections from many clients.

Typically, the VPN server is a separate hardware device, most often a security appliance such as a Cisco ASA security appliance. VPN servers can also be implemented in software. For example, Windows Server includes built-in VPN capabilities even though they're not easy to configure. And a VPN server can be implemented in Linux as well.

Figure 24-3 shows one of the many VPN configuration screens for a Cisco ASA appliance. This screen provides the configuration details for an IPSec VPN connection. The most important item of information on this screen is the Pre-Shared Key, which is used to encrypt the data sent over the VPN. The client will need to provide the identical key in order to participate in the VPN.

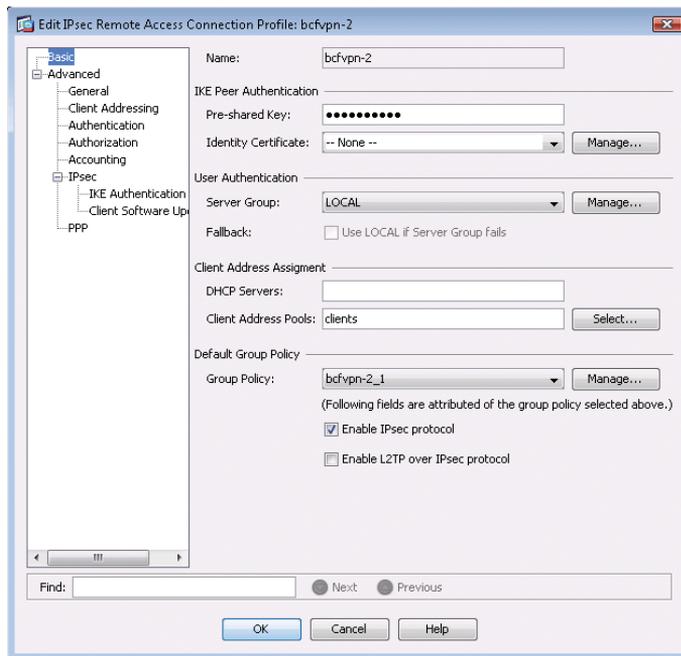


FIGURE 24-3:
An IPSec configuration page on a Cisco ASA security appliance.



REMEMBER

A VPN client is usually software that runs on a client computer that wants to connect to the remote network. The VPN client software must be configured with the IP address of the VPN server as well as authentication information such as a username and the Pre-Shared Key that will be used to encrypt the data. If the key used by the client doesn't match the key used by the server, the VPN server will reject the connection request from the client.

Figure 24-4 shows a typical VPN software client. When the client is configured with the correct connection information (which you can do by clicking the New button), you just click Connect. After a few moments, the VPN client will announce that the connection has been established and the VPN is connected.



FIGURE 24-4:
A VPN client.

A VPN client can also be a hardware device, like another security appliance. This is most common when the VPN is used to connect two networks at separate locations. For example, suppose your company has an office in Pixley and a second office in Hooterville. Each office has its own network with servers and client computers. The easiest way to connect these offices with a VPN would be to put an identical security appliance at each location. Then you could configure the security appliances to communicate with each other over a VPN.



Networking Beyond Windows

IN THIS PART . . .

Learn about Linux, a commonly used alternative to the Windows Server operating system for server computers.

Connect to your network with a Mac.

IN THIS CHAPTER

Finding out about Linux and how it differs from Windows

Choosing which version of Linux to use for your server

Installing Linux as well as configuring network settings and user accounts

Using Samba to create a file server

Chapter 25

Networking with Linux

L*inux*, the free operating system (OS) based on Unix, is a popular alternative to Windows Server, especially for specific applications such as web servers or email servers. Linux can also be used as a firewall or as a file server and print server on your local area network (LAN).

Linux has many advantages over Windows, not the least of which is that it is free. But price isn't the only advantage. Many network administrators have found that Linux is more stable than Windows, crashing less often and requiring less down-time for maintenance. In addition, Linux has a solid reputation for efficiency and security.

Linux was created in 1991 by Linus Torvalds, who was then an undergraduate student at the University of Helsinki in Finland. Linus thought it'd be fun to create his own OS based on Unix for his brand-new PC. In the nearly two decades since Linux was first conceived, Linux has become a full-featured operating system that is fast and reliable.

This chapter shows the basics of setting up a Linux server on your network and using it as a file server, as a web server for the Internet or an intranet, as an email server, and as a router and firewall to help connect your network to the Internet.



TIP

Linux is a complicated OS. Understanding how to use it can be a daunting task, especially if your only prior computer experience is with Windows. John Wiley & Sons, Inc., has *For Dummies* books that make Linux less painful. You'll find more comprehensive information about Linux in my book, *Networking All-In-One For Dummies*, 6th Edition; and you may also want to check out *Linux For Dummies*, 9th Edition, by Richard Blum.

Comparing Linux with Windows

If your only computer experience is with Windows, you're in for a steep learning curve when you first get into Linux. There are many fundamental differences between the Linux operating system and Windows. Here are some of the more important differences:

» **Linux is a multiuser operating system.** More than one user can log on and use a Linux computer at the same time:

- Two or more users can log on to a Linux computer from the same keyboard and monitor by using virtual consoles, which let you switch from one user session to another with a special key combination.
- Users can log on to the Linux computer from a terminal window running on another computer on the network.

Most versions of Windows are single-user systems. Only one user at a time can log on to a Windows computer and run commands. (Windows Server can be configured as a multiuser system with terminal services.)

» **Linux doesn't have a built-in graphical user interface (GUI) as Windows does.** The GUI in Linux is provided by an optional component called *X Window System*. You can run Linux without X Window, in which case you interact with Linux by typing commands. If you prefer to use a GUI, you must install and run X Window.

X Window is split into two parts:

- *A server component* — X Server — manages multiple windows and provides graphics services for application programs.
- *A UI component* — a window manager — provides user interface (UI) features, such as menus, buttons, toolbars, and a taskbar.

Several window managers are available, each with a different look and feel. With Windows, you're stuck with the UI that Microsoft designed. With Linux, you can use the UI of your choosing.



TECHNICAL
STUFF



TECHNICAL
STUFF



TECHNICAL
STUFF

» **Linux can't run Windows programs.** Nope, you can't run Microsoft Office on a Linux system; instead, you must find a similar program that's written specifically for Linux. Many Linux distributions come with the OpenOffice suite, which provides word processing, spreadsheet, presentation, graphics, database, email, calendar, and scheduling software. Thousands of other programs are available for Linux.

Windows emulator programs — the best-known is Wine — can run *some* Windows programs on Linux. However, the emulators run only some Windows programs, and they run them slower than they would run on a Windows system.

» **Linux doesn't do Plug and Play the way Windows does.** Major Linux distributions come with configuration programs that can automatically detect and configure the most common hardware components, but Linux doesn't have built-in support for Plug-and-Play hardware devices. You're more likely to run into a hardware-configuration problem with Linux than with Windows.

» **Linux uses a different system for accessing disk drives and files than Windows does.** For an explanation of how the Linux file system works, see the "I can't see my C drive!" sidebar that's coming up in this chapter.

» **Linux runs better on older hardware than the current incarnations of Windows do.** Linux is an ideal OS for an older Pentium computer with at least 32MB of RAM and 2GB of hard drive space.

If you're fond of antiques, Linux can run well on even a 486 computer with as little as 4MB of RAM and a few hundred MB of disk space.



TECHNICAL
STUFF

I CAN'T SEE MY C DRIVE!

Well, no, but that's normal. Linux and Windows have completely different ways of referring to your computer's disk drives and partitions. The differences can take some getting used to for experienced Windows users.

Windows uses a separate letter for each drive and partition on your system. For example, if you have a single drive formatted into three partitions, Windows identifies the partitions as drives C, D, and E. Each of these drives has its own root directory, which can in turn contain additional directories used to organize your files. As far as Windows is concerned, drives C, D, and E are completely separate drives even though the drives are actually just partitions on a single drive.

(continued)

(continued)

Linux doesn't use drive letters. Instead, Linux combines all the drives and partitions into a single directory hierarchy. In Linux, one of the partitions is designated as the root partition. The root is roughly analogous to the C drive on a Windows system. Then, the other partitions can be mounted on the root partition and treated as if they were directories on the root partition. For example, you might designate the first partition as the root partition and then mount the second partition as `/user` and the third partition as `/var`. Then any files stored in the `/user` directory would actually be stored in the second partition, and files stored in the `/var` directory would be stored in the third partition.

The directory where a drive mounts is the drive's *mount point*.

Notice that Linux uses regular forward-slash characters (/) to separate directory names rather than the backward-slash characters (\) used by Windows. Typing backslashes instead of regular slashes is one of the most common mistakes made by new Linux users.

While I'm on the subject, Linux uses a different convention for naming files, too. In Windows, filenames end in a three-letter (sometimes more letters than that) extension separated from the rest of the filename by a period. The extension is used to indicate the file type. For example, files that end in `.exe` are program files, but files that end in `.doc` are word-processing documents.

Linux doesn't use filename extensions, but periods are often used in Linux filenames to separate different parts of the name — and the last part often indicates the file type. For example, `ldap.conf` and `pine.conf` are both configuration files.

Choosing a Linux Distribution

Because the *kernel* (the core operating functions) of the Linux OS is free, several companies have created their own distributions of Linux, which include the Linux OS along with a bundle of packages, such as administration tools, web servers, and other useful utilities as well as printed documentation.

The following are some of the more popular Linux distributions:

» **Fedora:** One of the popular Linux distributions. You can download Fedora free from <http://fedoraproject.org>. You can also obtain it by purchasing any of several books on Fedora that include the Fedora distribution on DVD or CD-ROM.

All the examples in this chapter are based on Fedora 22.

- » **Mandriva Linux:** Another popular Linux distribution, one that is often recommended as the easiest for first-time Linux users to install. This distribution was formerly known as *Mandrake Linux*. Go to www.mandriva.com for more information.
- » **Ubuntu:** A Linux distribution that has gained popularity in recent years. It focuses on ease of use. For more information, go to www.ubuntu.com.
- » **SUSE:** Pronounced *SOO-zuh*, like the name of the famous composer of marches; a popular Linux distribution sponsored by Novell. You can find more information at www.suse.com.
- » **Slackware:** One of the oldest Linux distributions and still popular, especially among Linux old-timers. A full installation of Slackware gives you all the tools that you need to set up a network or Internet server. See www.slackware.com for more information.

All distributions of Linux include the same core components: the Linux kernel, an X Server, popular windows managers (such as GNOME and KDE), compilers, Internet programs such as Apache, Sendmail, and so on. However, not all Linux distributions are created equal. In particular, the manufacturer of each distribution creates its own installation and configuration programs to install and configure Linux.

The installation program is what makes or breaks a Linux distribution. All the distributions I list in this section have easy-to-use installation programs that automatically detect the hardware that's present on your computer and configure Linux to work with that hardware, thus eliminating most — if not all — manual configuration chores. The installation programs also let you select the Linux packages that you want to install and let you set up one or more user accounts besides the root account.

Installing Linux

All the Linux distributions I describe in the earlier section, “Choosing a Linux Distribution,” include an installation program that simplifies installing Linux on your computer. The installation program asks you a series of questions about your hardware, what components of Linux you want to install, and how you want to configure certain features. Then it copies the appropriate files to your hard drive and configures your Linux system.



TIP

If the thought of installing Linux gives you hives, you can buy computers with Linux preinstalled, just as you can buy computers with Windows already installed.



TIP

An excellent way to dip your feet into the Linux waters is to install it on a virtual machine, using a free virtual platform such as VMware Player or Oracle's VirtualBox.

Before you begin to install Linux, I recommend several planning steps:

» **Hardware:** Make a list of all the hardware components on your computer and how they're configured.

» **Partitioning:** Decide how you want to partition your hard drive for Linux.

Although Windows is usually installed into a single disk partition, Linux installations typically require at least three hard-drive partitions:

- *A boot partition:* This should be small; 16MB is recommended. The boot partition contains the OS kernel and is required to start Linux properly on some computers.
- *A swap partition:* This should be about twice the size of your computer's RAM. For example, if the computer has 2GB of RAM, allocate a 4GB swap partition. Linux uses this partition as an extension of your computer's RAM.
- *A root partition:* This, in most cases, uses up the remaining free space on the disk. The root partition contains all the files and data used by your Linux system.

You can also create additional partitions if you wish. The installation program includes a disk-partitioning feature that lets you set up your disk partitions and indicate the mount point for each partition. (For more information about disk partitions, see the sidebar, "I can't see my C drive!," earlier in this chapter.)

» **Packages:** Decide which optional Linux packages to install along with the Linux kernel:

- *All:* If you have enough drive space, install all the packages that come with your distribution. That way, if you decide you need to use a package, you won't have to figure out how to install the package outside of the installation program.
- *Pick and choose:* If you're tight on space, make sure that you at least install the basic network and Internet server packages, including Apache, Sendmail, FTP, and Samba.

» **Password:** Set the password for the root account.

» **User accounts:** In most distributions, you choose whether to create at least one user account.

Create at least one user account during installation so you can log on to Linux as a user (not with the root account). As a user, you can experiment with Linux commands without accidentally deleting or corrupting a needed system file.



TIP



REMEMBER

On Again, Off Again

Any user who accesses a Linux system, whether locally or over a network, must be authenticated by a valid user account on the system. The following sections lay out the whys, hows, and wherefores of logging on and logging off a Linux system — and how to shut down the system.

Logging on

When Linux boots up, it displays a series of startup messages while it starts the various services that make up a working Linux system. Assuming that you selected X Server when you installed Linux, you're eventually greeted by the screen, as shown in Figure 25-1. To log on to Linux, click your user ID if it is displayed. If your user ID isn't displayed, click Not Listed and then enter your user ID. Then, when prompted, type your password and press Enter.

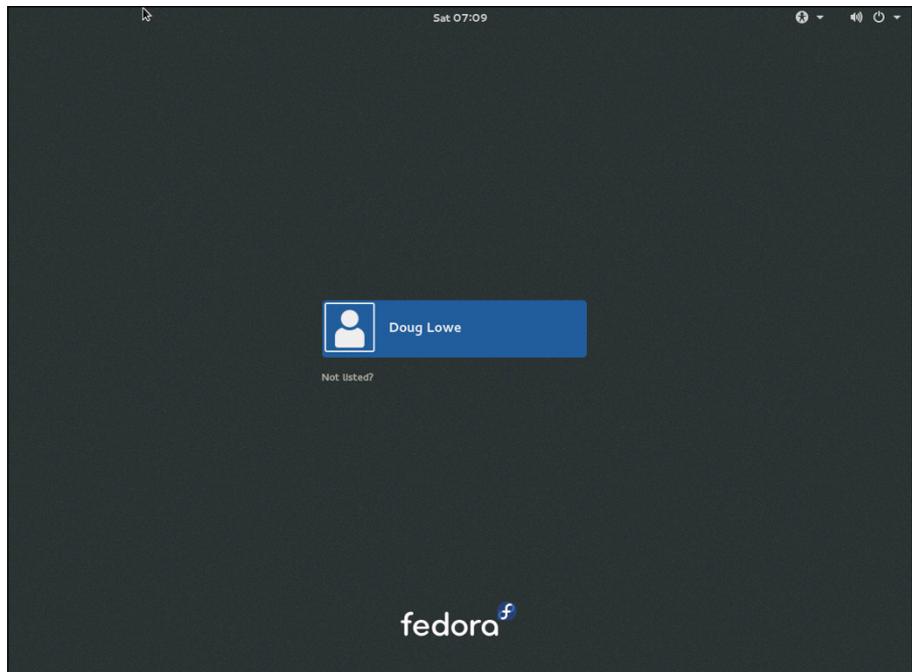


FIGURE 25-1:
Logging on to
Linux.



TIP

As a part of the installation process, the Setup Agent creates a user account for you. Use this user account rather than the `root` user account whenever possible. Use the `root` account only when you're making major changes to the system's configuration. When you're doing routine work, log on as an ordinary user to avoid accidentally corrupting your system.

When you log on, Linux grinds its gears for a moment and then displays the GNOME desktop, which I describe later in this chapter.

Logging off

After you log on, you probably want to know how to log off. To do so, hover the mouse over your name in the top-right corner of the screen, then choose Log Out.

Shutting down

Like with any OS, you shouldn't turn off the power to a Linux server without shutting down the system. The two ways to shut down Linux are

- » Press Ctrl+Alt+Delete.
- » Click your username in the upper right of the screen and then click Shutdown.

Using GNOME

Figure 25-2 shows a typical GNOME desktop. Although the GNOME desktop looks a lot different from the Windows desktop, many of the basic skills used for working with Microsoft Windows — moving or resizing windows, minimizing or maximizing windows, and using drag-and-drop to move items between windows — are almost exactly the same in GNOME.

Here are some key features of the GNOME desktop:

- » **Activities:** The Activities Overview provides a single access point for all GNOME applications. It provides fast access to common functions, such as Internet browsing, email, or file management, as well as desktop access to other applications. You can access Activities Overview by pressing the Windows key on the keyboard or clicking Activities in the top-left corner. (The screen in Figure 25-2 shows the Activities Overview open.)

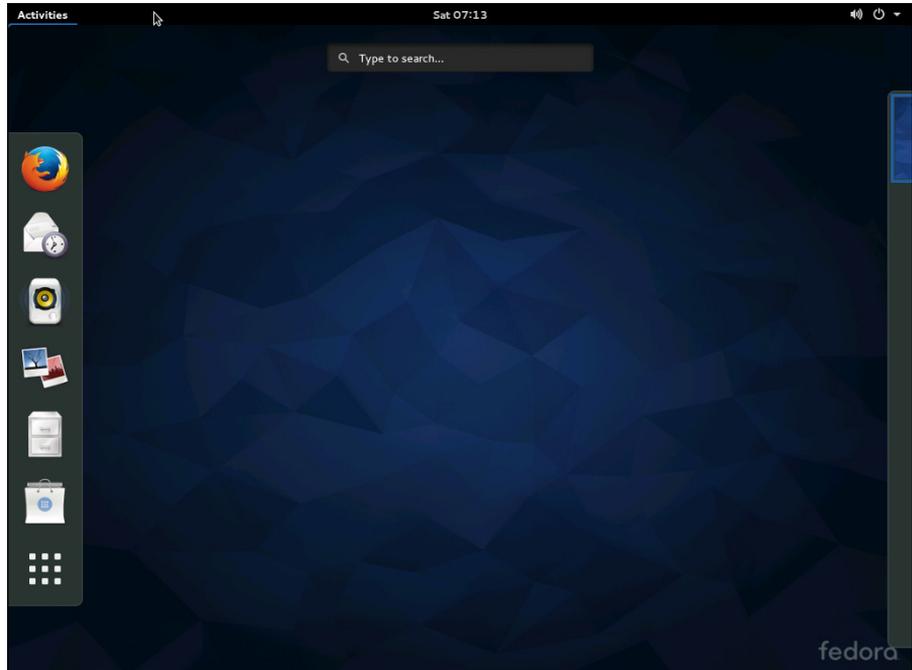


FIGURE 25-2:
A typical GNOME
desktop.

- » **Search box:** The search box in the top-middle part of the screen is the easiest way to find things in GNOME. For example, if you want to run the `gedit` program to edit a text file, search for “`gedit`.” Or if you want to fiddle with network settings, search for “`Network`.”
- » **Settings:** To manage your system or user settings, click your name at the top right of the screen. This reveals a menu with options for various settings.

Getting to a Command Shell

A command shell is a text-based window in which you can enter Linux commands directly, bypassing the graphical user interface. You can get to a command shell in one of two basic ways when you need to run Linux commands directly. The first is to press `Ctrl+Alt+Fx` to switch to one of the virtual consoles, where `Fx` is one of the function keys, from `F1` through `F12`. (A *virtual console* is simply a text-mode command prompt.) Then, you can log on and run commands to your heart’s content. When you’re done, press `Ctrl+Alt+F7` to return to GNOME.

Alternatively, you can open a command shell directly in GNOME by opening the Activities Overview, clicking Applications, and then clicking the Terminal icon. This opens a command shell in a window right on the GNOME desktop, as shown in Figure 25-3. Because this shell runs within the user account that GNOME is logged on as, you don't have to log on. You can just start typing commands. When you're done, type **Exit** to close the window.

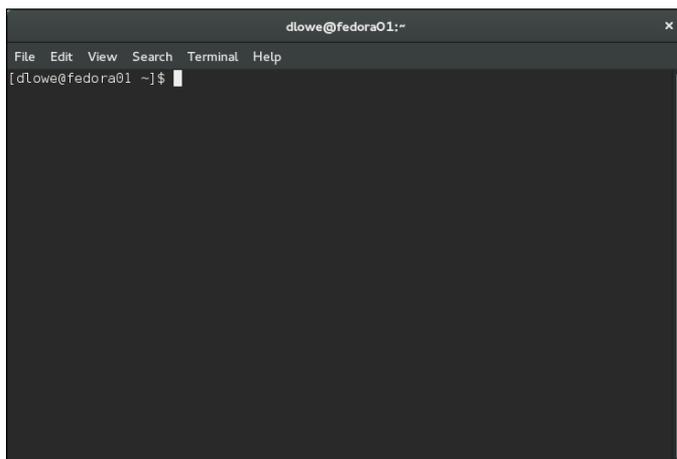


FIGURE 25-3: Using a Terminal window to run Linux commands.

Enabling the SUDO Command

Throughout the rest of this chapter, you'll see many examples of commands entered in a terminal window that begin with the word `sudo`. This command is an essential part of Linux administration. It lets you execute Linux commands with permissions of the root user account.

The `sudo` command is required because many Linux administrative commands can only be run under by the root user. You could simply log in as the root user to run such commands, but that practice is risky because the root user can do virtually anything in a Linux environment. It's safer to log in with an ordinary user account and use `sudo` to enable access to administrative functions.

For example, you'll often use the `dnf` command to install new software on a Linux system. The `dnf` command is one of those commands that can only be run by the root user. So, you'll need to use `sudo` to run the `dnf` command. To use `sudo`, you simply prefix the command you want to run with the word `sudo`, as in the following example:

```
sudo dnf install dhcp
```

Here, the command `dnf install dhcp` will be run as the root user. Note that for security purposes, the `sudo` command prompts you for your own password before it runs the `dnf` command.

To enable your user account for `sudo`, the root user must add your account name to a group called `wheel`, because by default the `sudo` program is configured to allow all users in this group to run commands as the root user. To add yourself to the `wheel` group, you must edit a configuration file called `group`, which is found in the `etc/` folder. Here are the steps to edit this file:

- 1. Log in as the root user.**

The Gnome desktop appears.

- 2. Click Activities at the top left of the Gnome desktop, and then choose Files.**

The File Manager appears, as shown in Figure 25-4.

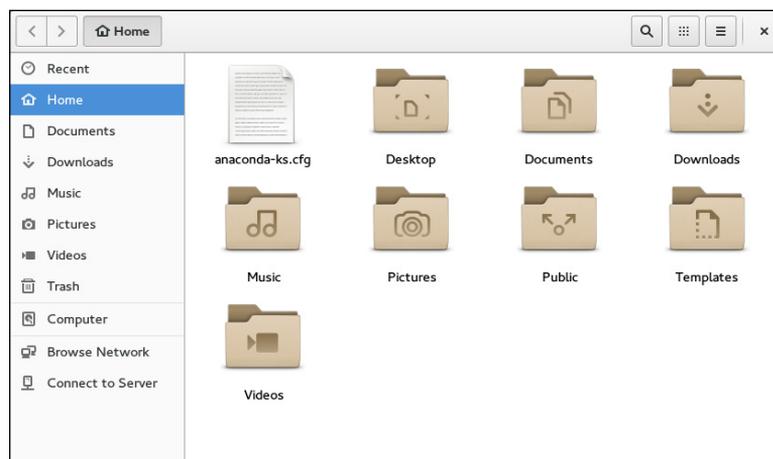


FIGURE 25-4:
The File Manager
window.

- 3. Click Computer in the navigation pane on the left side of the File Manager window.**

The folders at the computer's root level appear, as shown in Figure 25-5.

- 4. Double-click the `etc` folder.**

The files in the `etc` folder appear.

- 5. Locate and double-click the file named `group`.**

The `group` file is opened in the `gedit` text editor, as shown in Figure 25-6.

FIGURE 25-5:
The root-level folders.

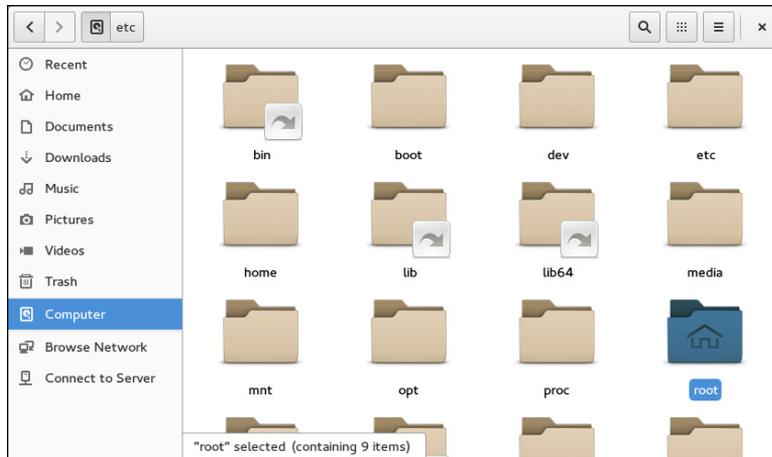
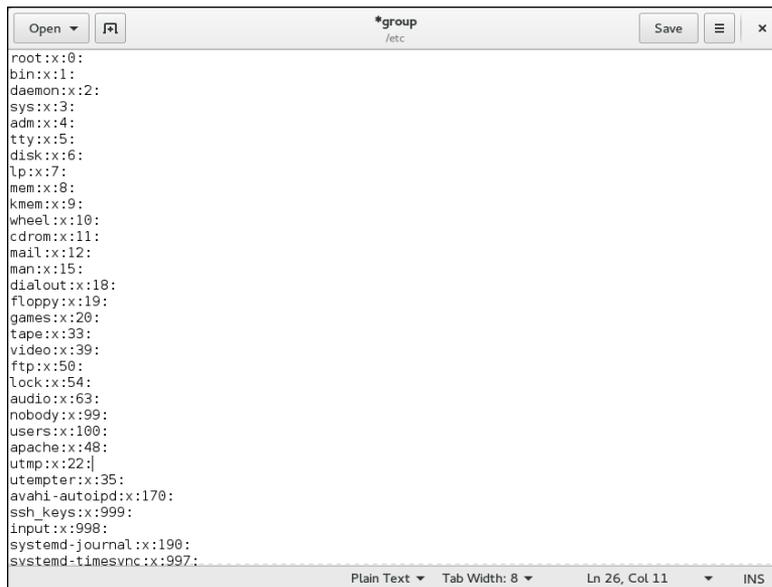


FIGURE 25-6:
Editing the group file.



- 6. Locate the line that starts with `wheel : x : 10 :`**
- 7. Add your username to the end of this line.**

For example, my username is `dlowe`, so I edited the line to read as follows:

```
wheel : x : 10 : dlowe
```

8. Click the Save button.

The group file is saved with your changes.

9. Close the gedit and File Manager windows.

You're done! You can now use the `sudo` command.

Managing User Accounts

One of the most common network administration tasks is adding a user account. The Setup Agent prompts you to create a user account the first time you start Linux after installing it. However, you'll probably need to create additional accounts.

Each Linux user account has the following information associated with it:

- » **Username:** The name the user types to log on to the Linux system.
- » **Full name:** The user's full name.
- » **Home directory:** The directory where the user is placed when he logs on. In Fedora, the default home directory is `/home/username`. For example, if the username is `blowe`, the home directory will be `/home/blowe`.
- » **Shell:** The program used to process Linux commands. Several shell programs are available. In most distributions, the default shell is `/bin/bash`.
- » **Group:** You can create group accounts, which makes it easy to apply identical access rights to groups of users. Group accounts are optional, but useful if you have more than just a few users.
- » **User ID:** The internal identifier for the user.

You can add a new user by using the `useradd` command. For example, to create a user account named `slowe`, using default values for the other account information, open a Terminal window or switch to a virtual console and type this command:

```
sudo useradd slowe
```

The `useradd` command has many optional parameters that you can use to set account information, such as the user's home directory and shell.

Fortunately, most Linux distributions come with special programs that simplify routine system management tasks. For example, Fedora comes with the user Manager program, as shown in Figure 25-7. To start this program, click Activities, click Applications, and then double-click Users and Groups (scroll down to find it).

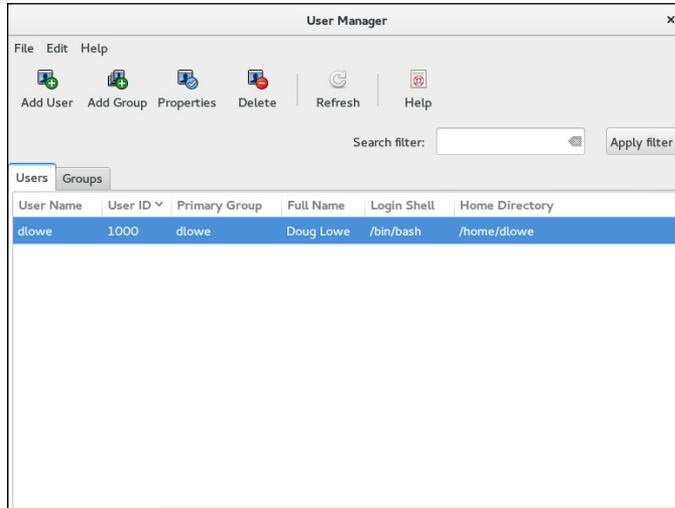


FIGURE 25-7: Managing users and groups.

To create a user with User Manager, click the Add User button. This brings up a dialog box that asks for the user's name, password, and other information. Fill out this dialog box and then click OK.

The User Manager also lets you create groups. You can simplify the task of administering users by applying access rights to groups rather than individual users. Then, when a user needs access to a resource, you can add the user to the group that has the needed access.

To create a group, click the Add Group button. A dialog box appears, asking for the name of the new group. Type the name you want and then click OK.

To add a user to a group, click the Groups tab in the User Manager. Then, double-click the group to which you want to add users. This brings up the Group Properties dialog box. Click the Group Users tab and then select the users that you want to belong to the group.

Network Configuration

In many cases, configuring a Linux server for networking is a snap. When you install Linux, the Installation program automatically detects your network adapters and installs the appropriate drivers. Then you're prompted for basic network-configuration information, such as the computer's IP address, hostname, and so on.

You may need to manually change your network settings after installation or configure advanced networking features that aren't configured during installation. In the following sections, you get a look at the basic procedures for configuring Linux networking services.

Using the Network Configuration program

Before you can use a network interface to access a network, you have to configure the interface's basic TCP/IP options, such as its IP address, host name, Domain Name System (DNS) servers, and so on. This configuration is automatically set up when you install Linux, but you may need to change it later on. In this section, I show you how to do that by using Fedora's Network settings program. You can access this program by clicking your name in the top-right corner of the screen, choosing Settings, and then choosing Network.



TIP

Most other Linux distributions have similar programs.

The Network Configuration program lets you configure the basic TCP/IP settings for a network interface by pointing and clicking your way through tabbed windows. Figure 25-8 shows the Network Configuration program in action.



FIGURE 25-8: The Network Configuration program.

Notice that the main window of the Network Configuration lists all the network interfaces installed in your computer. You can double-click any of the interfaces to bring up a window similar to the one shown in Figure 25-9. This window lets you set the configuration options for the network interface, such as its IP address and other TCP/IP-configuration information.

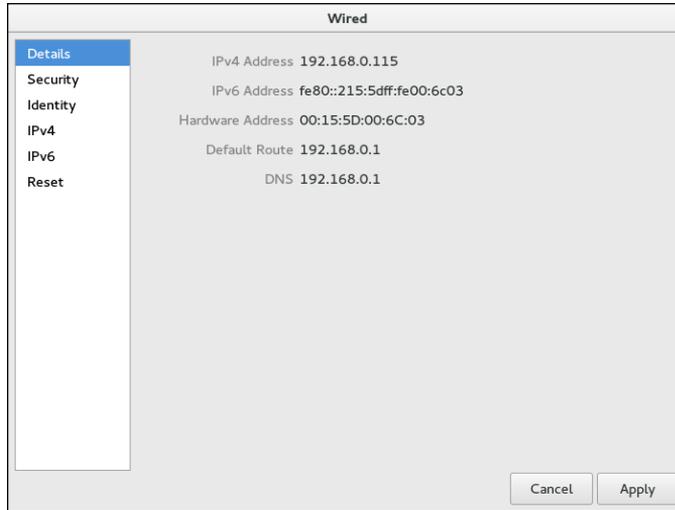


FIGURE 25-9: Configure basic TCP/IP settings here.

Restarting your network

Whenever you make a configuration change to your network, you must restart Linux networking services for the change to take effect. If you find that annoying, just be thankful that you don't have to restart the entire computer. Simply restarting the network services is sufficient.

Open a console by pressing `Ctrl+Alt+n`, where *n* is a number from 2–7. Log in using the root account and then enter the following command:

```
service network restart
```

To confirm that the service was properly restarted, you'll see a message similar to this:

```
Restarting network (via systemctl): [ OK ]
```

Doing the Samba Dance

Until now, you probably thought of Samba as an intricate Brazilian dance with fun rhythms. But in the Linux world, *Samba* refers to a file- and printer-sharing program that allows Linux to mimic a Windows file-and-print server so Windows computers can use shared Linux directories and printers. If you want to use Linux as a file or print server in a Windows network, you have to know how to dance the Samba.

Understanding Samba

Because Linux and Windows have such different file systems, you can't create a Linux file server simply by granting Windows users access to Linux directories. Windows client computers wouldn't be able to access files in the Linux directories. Too many differences exist between the file systems. For example:



REMEMBER

- » **Case sensitivity:** Linux filenames are case sensitive, whereas Windows filenames are not. For example, in Windows, `File1.txt` and `file1.txt` are the same file. In Linux, they're different files.
- » **File extensions:** In Linux, periods aren't used to denote file extensions. Linux filenames don't use extensions.
- » **File attributes:** Windows has file attributes like *read-only* and *archive*. Linux doesn't have these.

More fundamentally, Windows networking uses the Server Message Block (SMB) protocol to manage the exchange of file data among file servers and clients. Linux doesn't have SMB support built in. And that's why Samba is required. Samba is a program that mimics the behavior of a Windows-based file server by implementing the SMB protocol. So when you run Samba on a Linux server, the Windows computers on your network see the Linux server as if it were a Windows server.

Like a Windows server, Samba works by designating certain directories as shares. A *share* is simply a directory that's made available to other users via the network. Each share has the following elements:



TIP

- » **Name:** The name by which the share is known over the network
Share names should be eight characters whenever possible.
- » **Path:** The path to the directory on the Linux computer that's being shared, such as `\Users\Doug`
- » **Description:** A one-line description of the share
- » **Access:** A list of users or groups who have been granted access to the share



TIP

Samba also includes a client program that lets a Linux computer access Windows file servers.



TECHNICAL
STUFF

Why did Samba's developers choose to call their program Samba? Simply because the protocol that Windows file and print servers use to communicate with each other is SMB. Add a couple of vowels to *SMB*, and you get *Samba*.

Installing Samba

To install Samba, open a console or terminal window and enter the following command:

```
sudo dnf install samba.x86_64
```

Assuming that Samba is not already installed, `dnf` will ask for your permission to install Samba. Enter `y` to proceed. Be patient while the package downloads and installs.

After you've installed the basic Samba package, you should also install the GNOME-based Samba configuration tool:

```
sudo dnf install system-config-samba
```

This tool allows you to configure Samba from your GNOME desktop instead of by directly editing the configuration file.



TIP

One sure way to render a Samba installation useless is to enable the default Linux firewall settings on the computer that runs Samba. The Linux firewall is designed to prevent users from accessing network services, such as Samba. It's designed to be used between the Internet and your local network — not between Samba and your local network. Although you can configure the firewall to allow access to Samba only to your internal network, a much better option is to run the firewall on a separate computer. That way, the firewall computer can concentrate on being a firewall, and the file server computer can concentrate on being a file server.

Starting and stopping Samba

Before you can use Samba, you must start its two daemons, `smbd` and `nmbd`. (*Daemon* is the Linux term for *service*, which is a program that runs as a background task.) Both daemons can be started at once by starting the SMB service. From a command shell, use this command:

```
sudo service smb start
```

Whenever you make a configuration change, such as adding a new share or creating a new Samba user, you should stop and restart the service with this command:

```
sudo service smb restart
```

If you prefer, you can stop and start the service with separate commands:

```
sudo service smb stop
sudo service smb start
```

If you're not sure whether Samba is running, enter this command:

```
sudo service smb status
```

You get a message indicating whether the `smbd` and `nmbd` daemons are running.

To configure Samba to start automatically when you start Linux, use this command:

```
sudo chkconfig --level 35 smb on
```

To make sure that the `chkconfig` command worked right, enter this command:

```
sudo chkconfig --list smb
```

You should see output similar to the following:

```
Smb      0:off 1:off 2:off 3:on  4:off 5:on  6:off
```



TIP

You can independently configure services to start automatically for each of the six boot levels of Linux. Boot level 3 is normal operation without a GUI; level 5 is normal operation with a GUI. So setting SMB to start for levels 3 and 5 makes SMB available, regardless of whether you're using a GUI.

Using the Samba Server Configuration tool

Fedora includes a handy GNOME-based configuration tool that simplifies the task of configuring Samba. To start it, click Activities, click Applications, and then click Samba. When you do so, the Samba Server Configuration tool appears, as shown in Figure 25-10. This tool lets you configure basic server settings and manage shares.

To make your Samba server visible on the network, choose Preferences ⇨ Server Settings. This brings up a dialog box that lets you set the workgroup name (which must match the workgroup or domain name you want the Samba server to belong to) and a description for the server, as well as some basic security settings that control how users can access the Samba server.

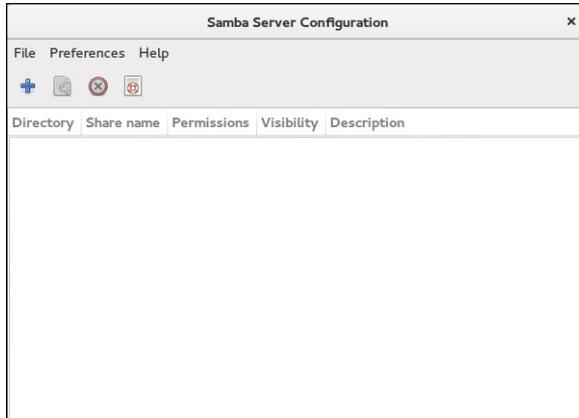


FIGURE 25-10: Using the Samba Server Configuration tool.

You can set four basic types of security for your Samba server:

» **Domain:** Configures the Samba server to use a Windows domain controller to verify the user. If you specify this option, you must

- Provide the domain controller's name in the Authentication Server field.
- Set Encrypted Passwords to Yes (if you use Domain mode).

» **Server:** Configures Samba to use another Samba server to authenticate users.

If you have more than one Samba server, this feature lets you set up user accounts on just one of the servers. Then, in the Authentication Server field, specify the name of the Samba server that should perform the authentication.

» **Share:** Authorizes users separately for each share they attempt to access.

» **User:** Requires that users provide a valid username and password when they first connect to a Samba server. That authentication then grants them access to all shares on the server, subject to the restrictions of the account they're authorized under.

User mode is the default.



TIP



TECHNICAL
STUFF

For each network user who needs to access the Samba server, you must follow these steps:

1. Create a Linux user account for each user.
2. Create a separate Samba user account.

To create a Samba user account, choose Preferences ⇄ Samba Users from the Samba Server Configuration window. This brings up the Samba Users dialog box, as shown in Figure 25-11. You can use this dialog box to add, edit, or delete users.



REMEMBER

The Samba user account maps to an existing Linux user account, so you must create the Linux user account first.



FIGURE 25-11: The Samba Users dialog box lists your Samba users.

To be useful, a file server should offer one or more *shares* — directories that have been designated as publicly accessible via the network. Again, you use the Samba Server Configuration program to manage your shares. To add a share, click the Add button on the Samba Server Configuration program’s toolbar. This brings up the Create Samba Share dialog box, as shown in Figure 25-12. You can then

- » Enter the path for the directory you want to share.
- » Enter a description for the share.
- » Select whether to allow either read-only or read-write access.
- » Click the Access tab if you want to set limits on access (for example, to specific users).

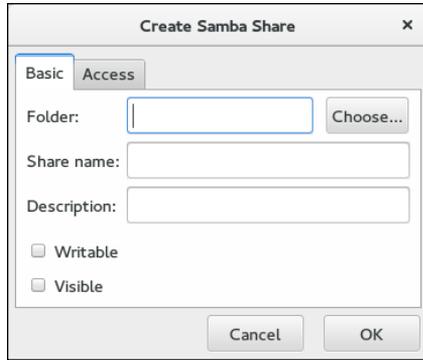


FIGURE 25-12:
The Create Samba Share dialog box.



TIP

When you create a new share using the Samba Configuration program, the share should be immediately visible to network users. If not, try restarting the Samba server, as I describe in the section, “Starting and stopping Samba,” in this chapter.

Chapter 26

Mac Networking

This book dwells on mostly networking Windows-based computers, as though Microsoft were the only game in town. I'm sure plenty of people in Redmond, WA (where Microsoft is headquartered) wished that it were so. But alas, there is an entirely different breed of computer: the Apple Macintosh, more commonly referred to simply as *Mac*.

Every Macintosh ever built, even an original 1984 model, includes networking support. Newer Macintosh computers have better built-in networking features than older Macintosh computers, of course. The newest Macs include built-in Gigabit Ethernet connections or 802.11ac wireless connections, or both. Support for these network connections is pretty much automatic, so all you have to do is plug your Mac into a network or connect to a wireless network, and you're ready to go.

This chapter presents what you need to know to network Mac computers running OS X, Apple's operating system for Mac computers. You'll see how to control basic Mac network options, such as TCP/IP and file sharing. I also show you how to join a Mac to a Windows domain network.

Basic Mac Network Settings

Most network settings on OS X are automatic. If you wish, you can look at and change the default network settings by following these steps:

1. **Choose System Preferences ⇄ Networking.**

The Network preferences page appears, as shown in Figure 26-1.

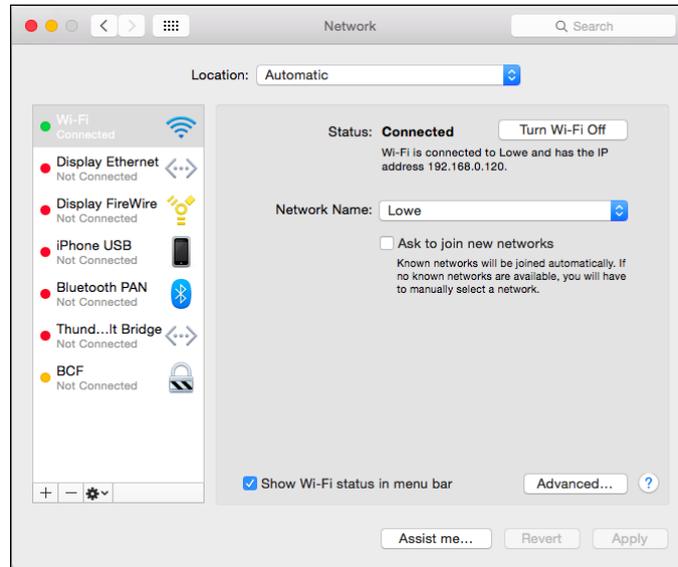


FIGURE 26-1:
Network preferences.

2. **Click Advanced.**

The advanced network settings are displayed, as shown in Figure 26-2.

3. **Click the TCP/IP tab to view or change the TCP/IP settings.**

This brings up the TCP/IP settings, as shown in Figure 26-3. From this page, you can view the currently assigned IP address for the computer. And, if you wish, you can assign a static IP address by changing the Configure IPv4 drop-down from Using DHCP to Manually. Then, you can enter your own IP address, subnet mask, and router address. (For more information about IP addresses, refer to Chapter 5.)

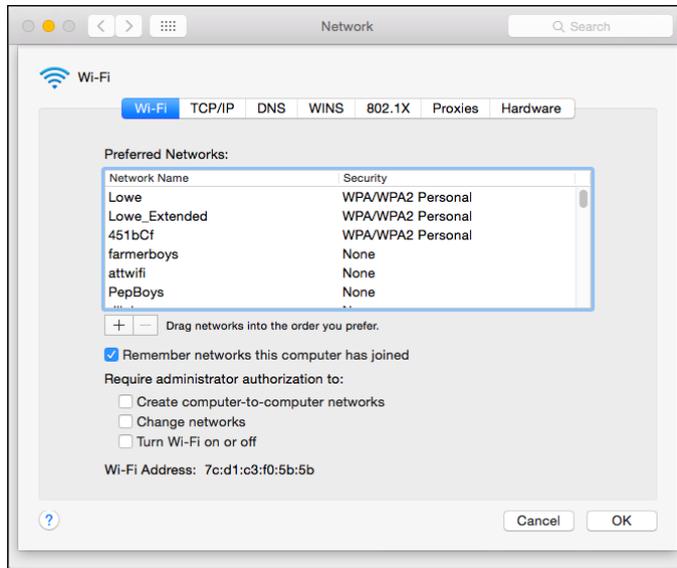


FIGURE 26-2:
Advanced network settings.

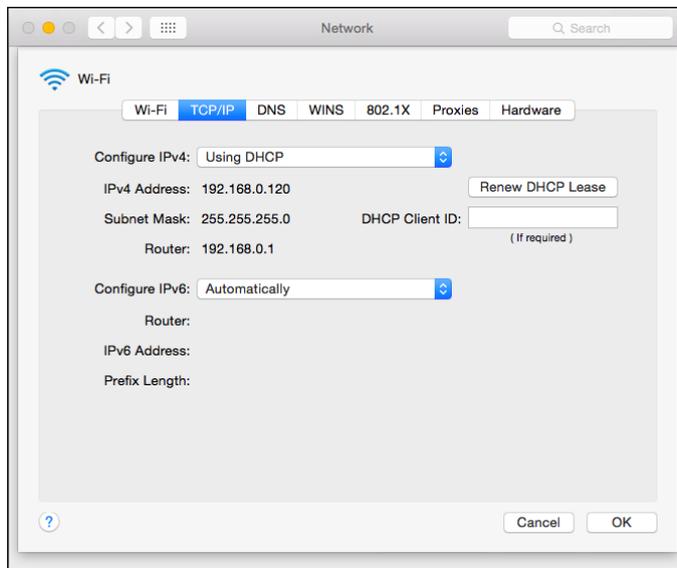


FIGURE 26-3:
Mac network TCP/IP settings.

4. Click the DNS tab to view or change the DNS settings.

This brings up the DNS settings shown in Figure 26-4. Here, you can see the DNS servers being used, and you can add additional DNS servers if you wish.

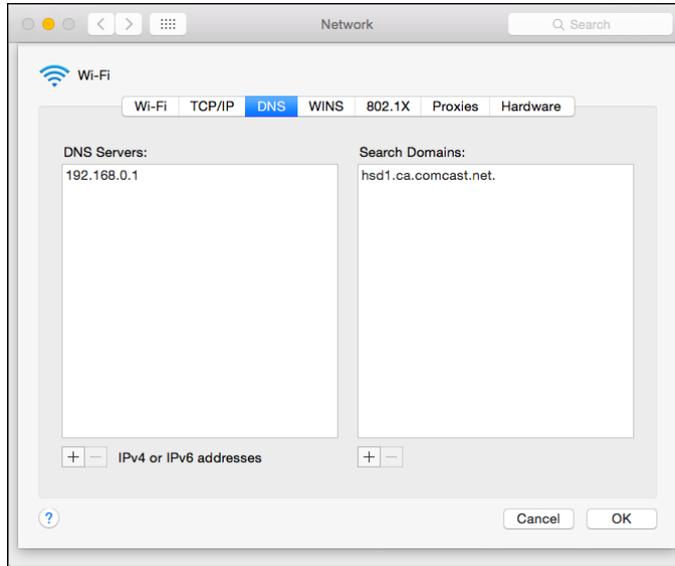


FIGURE 26-4: DNS settings.

5. Click the Hardware tab to view hardware information.

This brings up the settings page shown in Figure 26-5. The most useful bit of information on this tab is the MAC address, which is sometimes needed to set up wireless security. (For more information, refer to Chapter 9.)

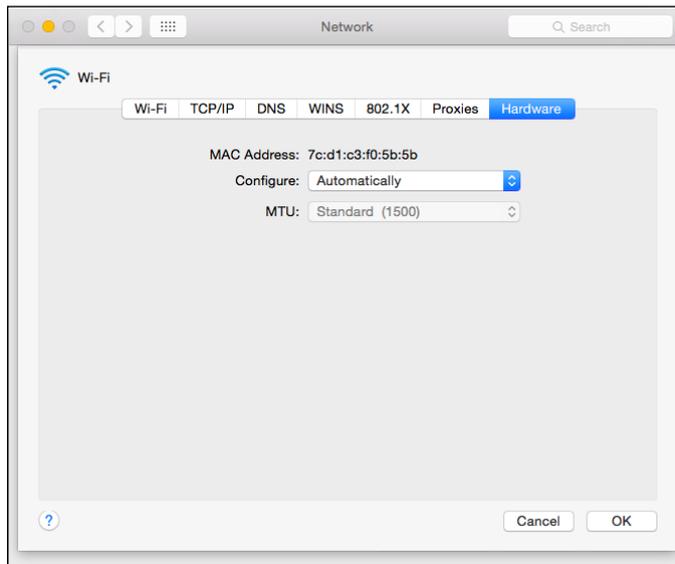


FIGURE 26-5: Hardware settings.

WHAT ABOUT OS X SERVER?

At one time, Apple offered the Mac OS X Server dedicated network operating system (NOS). In 2011, Apple merged Mac OS X Server with its desktop OS and made the server components of the OS available as an inexpensive add-on you can purchase from the App Store. For the latest version of OS X (10.8, released in July of 2012), the Server App enhancement can be purchased for about \$20.

The Server App download adds a variety of network server features to OS X, including

- **Apache web server**, which also runs on Windows and Linux systems
- **MySQL**, which is also available in Windows and Linux versions
- **Wiki Server**, which lets you set up web-based wiki, blog, and calendaring sites
- **NetBoot**, a feature that simplifies the task of managing network client computers
- **Spotlight Server**, which lets you search for content on remote file servers
- **Podcast Producer**, which lets the create and distribute multimedia programs

Joining a Domain

If you're using a Mac in a Windows domain environment, you can join the Mac to the domain by following these steps:

1. Choose Settings ⇨ Users & Groups.

This brings up the Users & Groups page, as shown in Figure 26-6.

2. Select the user account you want to join to the domain and then click Login Options.

The Login Options page appears, as shown in Figure 26-7.

3. If the lock icon at the bottom left of the page is locked, click it and enter your password when prompted.

By default, the user login options are locked to prevent unauthorized changes. This step unlocks the settings so that you can join the domain.

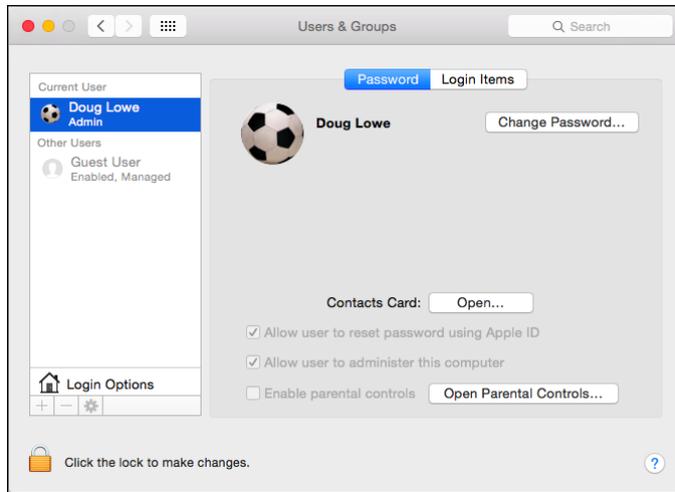


FIGURE 26-6: Users & Groups.

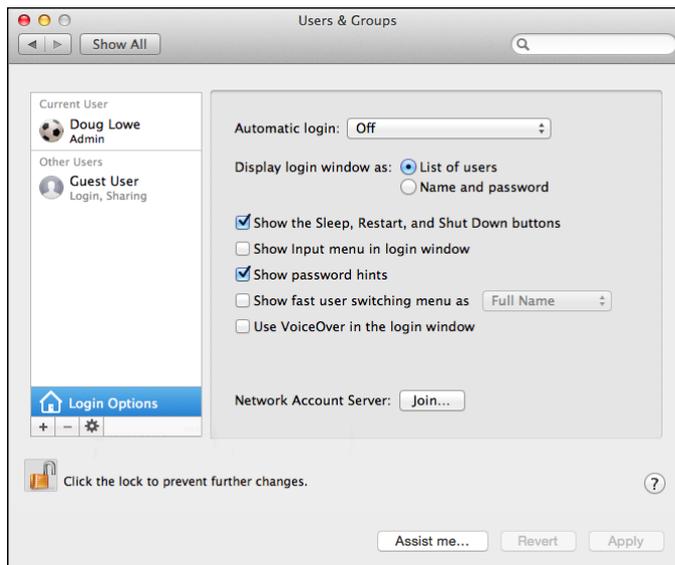


FIGURE 26-7: Login Options.

4. Click the Join button.

You're prompted to enter the name of the domain you want to join, as shown in Figure 26-8.

5. Enter the name of the domain you want to join.

When you enter the domain name, the dialog box expands to allow you to enter domain credentials to allow you to join the domain, as shown in Figure 26-9.

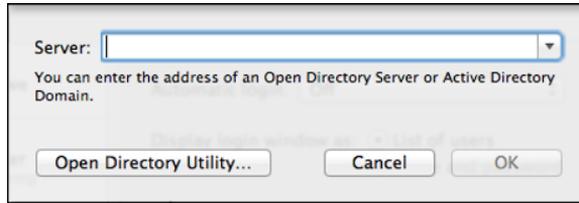


FIGURE 26-8:
Joining a domain.

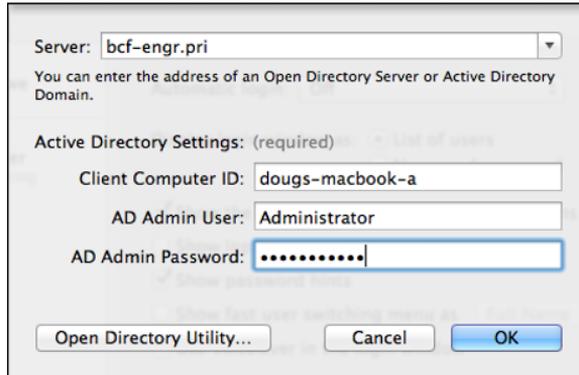


FIGURE 26-9:
Authenticating
with the domain.

- 6. Enter the name and password of a domain administrator account; then click OK.**

You return to the Login Options page, which shows that you have successfully joined the domain; see Figure 26-10.

- 7. Close the Users & Groups window.**

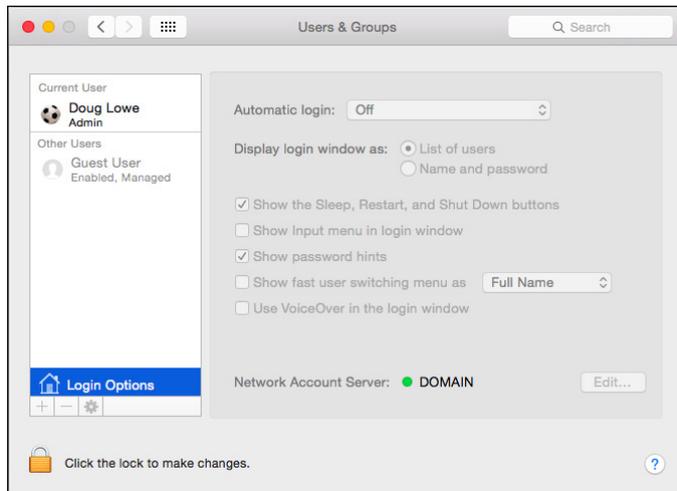


FIGURE 26-10:
Congratulations!
You have now
joined the
domain.

Connecting to a Share

After you join a domain, you can access its network shares via the Finder. Just follow these steps:

1. Click Finder.

The Finder window appears, as shown in Figure 26-11.

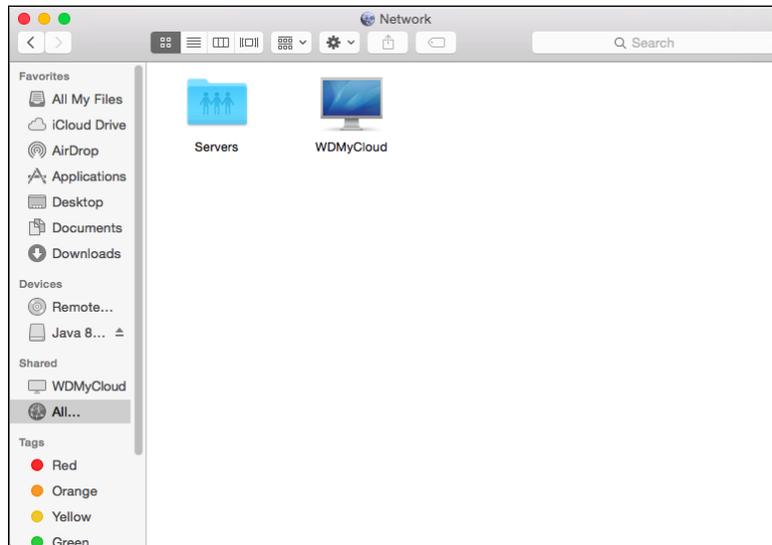


FIGURE 26-11:
Welcome to the
Finder.

2. Choose Go ⇨ Connect to Server.

The Connect to Server dialog box appears, as shown in Figure 26-12.

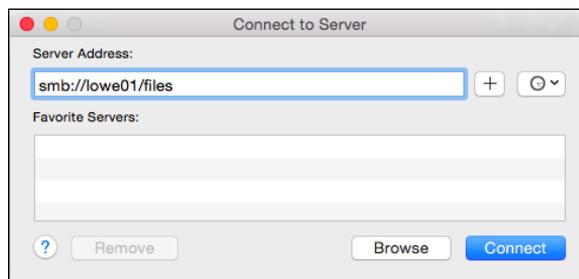


FIGURE 26-12:
The Connect to
Server dialog box.

3. Type the path that leads to the server share you want to connect to.

To type a network path, follow this syntax:

```
smb://server-name/share-name
```

Replace the *server-name* with the name of the server that contains the share and *share-name* with the name of the share. For example, to connect to a share named `files` on a server named `lowe01`, type `smb://lowe01/files`.

4. Click Connect.

You'll be prompted for login credentials.

5. Enter a domain username and password and then click OK.

Precede the user name with the domain name, separated by a backslash. For example, if the domain name is `lowewriter.pri` and the user name is `Doug`, enter `lowewriter.pri\Doug` as the username.

After connection is complete, the files in the share display in the Finder window. You can then open files directly from the share (provided you have the right software, such as Microsoft Office, to read the files). You can also drag and drop files between the Mac and the file shares.

7

The Part of Tens

IN THIS PART . . .

Learn the unwritten rules of networking.

Find out the common networking mistakes you should avoid.

Review a list of ten handy networking things you should keep in your closet.

Chapter 27

Ten Networking Commandments

“Blessed is the network manager who walks not in the council of the ignorant, nor stands in the way of the oblivious, nor sits in the seat of the greenhorn, but delights in the Law of the Network and meditates on this Law day and night.”

—NETWORKS 1:1

And so it came to pass that these Ten Networking Commandments were handed down from generation to generation, to be worn as frontlets between the computer geeks’ eyes (taped on the bridges of their broken glasses) and written upon their doorposts with Sharpie markers. Obey these commandments, and it shall go well with you, with your users, and with your users’ users.

I. Thou Shalt Back Up Thy Hard Drive Religiously

Prayer is a good thing, and I heartily recommend it. But when it comes to protecting the data on your network, nothing beats a well-thought-out schedule of backups followed religiously.

II. Thou Shalt Protect Thy Network from Infidels

One of my favorite classic TV series is *M*A*S*H*. One of the recurring characters on that show was Colonel Flagg, who hid in trash cans looking for Communists. I don't recommend that you actually become him, but on the other hand, you don't want to ignore the possibility of getting zapped by a virus, your network being invaded by hackers, or your data being compromised by an unscrupulous user. Make sure that your Internet connection is properly secured with a firewall, and don't allow any Internet access that circumvents your security.

To counter virus threats, use network-aware antivirus software to ensure that every user on your network has up-to-date virus protection. And teach your users so they know how to avoid those virus threats that manage to sneak past your virus protection.

III. Thou Shalt Keepeth Thy Network Drive Pure and Cleanse It of Old Files

Don't wait until your 4TB network drive is down to just 1GB of free space before you think about cleaning it up. Set up a routine schedule for disk housekeeping, where you wade through the files and directories on the network disk to remove old junk.

IV. Thou Shalt Not Tinker with Thine Network Configuration Unless Thou Knowest What Thou Art Doing

Networks are finicky things. After yours is up and running, don't mess with it unless you know what you're doing. You may be tempted to log on to your firewall router to see whether you can tweak some of its settings to squeeze another ounce of performance out of it. But unless you know what you're doing, be careful! (Be especially careful if you *think* you know what you're doing. It's the people who think they know what they're doing who get themselves into trouble!)

V. Thou Shalt Not Covet Thy Neighbor's Network

Network envy is a common malady among network managers. If your network users are happy with Windows 8.1, resist the urge to upgrade to Windows 10 unless you have a really good reason. And if you run Windows Server 2012 R2, fantasizing about Windows Server 2016 is a venial sin.

You're especially susceptible to network envy if you're a gadget freak. There's always a better switch to be had or some fancy network-protocol gizmo to lust after. Don't give in to these base urges! Resist the devil, and he will flee!

VI. Thou Shalt Schedule Downtime before Working upon Thy Network

As a courtesy, try to give your users plenty of advance notice before you take down the network to work on it. Obviously, you can't predict when random problems strike. But if you know you're going to patch the server on Thursday morning, you earn points if you tell everyone about the inconvenience two days before rather than two minutes before. (You'll earn even more points if you patch the server Saturday morning.)

VII. Thou Shalt Keep an Adequate Supply of Spare Parts

There's no reason that your network should be down for two days just because a cable breaks. Always make sure that you have at least a minimal supply of network spare parts on hand. (As luck would have it, Chapter 29 suggests ten things you should keep in your closet.)

VIII. Thou Shalt Not Steal Thy Neighbor's Program without a License

How would you like it if Inspector Clouseau (from the *Pink Panther* movies) barged into your office, looked over your shoulder as you ran Excel from a network server, and asked, “Do you have a *liesaunce*?”

“A *liesaunce*?” you reply, puzzled.

“Yes, of course, a *liesaunce* — that is what I said! The law specifically prohibits the playing of a computer program on a network without a proper *liesaunce*.”

You don't want to get in trouble with Inspector Clouseau, do you? Then make sure you have the correct licenses for the applications you run on your network.

IX. Thou Shalt Train Thy Users in the Ways of the Network

Don't blame the users if they don't know how to use the network. It's not their fault. If you're the network administrator, your job is to provide training so the network users know how to use the network.

X. Thou Shalt Write Down Thy Network Configuration upon Tablets of Stone

Network documentation should be written down. If you cross the River Jordan, who else will know diddly-squat about the network if you don't write it down somewhere? Write down everything, put it in an official binder labeled *Network Bible*, and protect the binder as if it were sacred.

Your hope should be that 2,000 years from now, when archeologists are exploring caves in your area, they find your network documentation hidden in a jar and marvel at how meticulously the people of our time recorded their network configurations.

They'll probably draw ridiculous conclusions, such as we offered sacrifices of burnt data packets to a deity named TCP/IP and confessed our transgressions in a ritual known as “logging,” but that makes it all the more fun.

Chapter 28

Ten Big Network Mistakes

Just about the time you figure out how to avoid the most embarrassing computer mistakes (such as using your CD drive's tray as a cup holder), the network lands on your computer. Now you have a whole new list of dumb things you can do, mistakes that can give your average computer geek a belly laugh because they seem so basic to him. Well, that's because he's a computer geek. Nobody had to tell *him* not to fold the floppy disk — he was born with an extra gene that gave him an instinctive knowledge of such things.

Here's a list of some of the most common mistakes made by network novices. Avoid these mistakes and you deprive your local computer geek of the pleasure of a good laugh at your expense.

Skimping on Hardware

Good computer equipment is not cheap. You can walk into your local electronics store and buy everything you need to set up a cheap network: cables, switches, and computers to use as servers. But you get what you pay for. Professional-grade equipment costs much more, and in a business environment, it's worth it.

Why? Because professional-grade equipment is designed with performance, reliability, and centralized management in mind.

Professional server computers typically include redundancy in all the key systems — duplicate power supplies, duplicate network ports, duplicate disk controllers, and often even duplicate CPUs and motherboards. So, if one component fails, the server can continue operating.

Professional switches typically include management features that let you pinpoint problems on your network, segment your network for better performance, and monitor your employees' usage of the network.

You may also be tempted to cut costs by stringing inexpensive cable directly from the switches to each computer on the network. In the long run, though, the Scrooge approach may actually prove to be more expensive than investing in a good cable installation in the first place. A professionally installed cable infrastructure will last much longer than the computers it services, and will be considerably more reliable.

Turning Off or Restarting a Server Computer While Users Are Logged On

The fastest way to blow your network users' accounts to kingdom come is to turn off a server computer while users are logged on. Restarting it by pressing its reset button can have the same disastrous effect.

If your network is set up with a dedicated file server, you probably won't be tempted to turn it off or restart it. But if your network is set up as a true peer-to-peer network, where each of the workstation computers — including your own — also doubles as a server computer, be careful about the impulsive urge to turn off or restart your computer. Someone may be accessing a file or printer on your computer at that very moment.

So, before you turn off or restart a server computer, find out whether anyone is logged on. If so, politely ask her to log off.



REMEMBER

Many server problems don't require a server reboot. Instead, you can often correct the problem just by restarting the particular service that's affected.

Deleting Important Files on the Server

Without a network, you can do anything you want to your computer, and the only person you can hurt is yourself. (Kind of like the old “victimless crime” debate.) Put your computer on a network, though, and you take on a certain amount of responsibility. You must find out how to live like a responsible member of the network society.

Therefore, you can’t capriciously delete files from a network server just because you don’t need them. They may not be yours. You wouldn’t want someone deleting your files, would you?

Be especially careful about files that are required to keep the network running. For example, some versions of Windows use a folder named `wgpo0000` to hold email. If you delete this folder, your email is history. Look before you delete.



WARNING

The first time you accidentally delete an important file from a network share, you may be unpleasantly surprised to discover that the Recycle Bin does not work for network files. The Recycle Bin saves copies of files you’ve deleted from your computer’s local hard disk, but it does *not* save copies of files you delete from network shares. As a result, you can’t undelete a file you’ve accidentally deleted from the network.

Copying a File from the Server, Changing It, and Then Copying It Back

Sometimes working on a network file is easier if you first copy the file to your local hard drive. Then you can access it from your application program more efficiently because you don’t have to use the network. This is especially true for large database files that have to be sorted to print reports.

You’re asking for trouble, though, if you copy the file to your PC’s local hard drive, make changes to the file, and then copy the updated version of the file back to the server. Why? Because somebody else may be trying the same thing at the same time. If that happens, the updates made by one of you — whoever copies the file back to the server first — are lost.

Copying a file to a local drive is rarely a good idea.

Sending Something to the Printer Again Just Because It Didn't Print the First Time

What do you do if you send something to the printer and nothing happens?

- » **Right answer:** Find out why nothing happened and *fix it*.
- » **Wrong answer:** Send it again and see whether it works this time.



Some users keep sending it over and over again, hoping that one of these days, it'll take. The result is rather embarrassing when someone finally clears the paper jam and then watches 30 copies of the same letter print. Or when 30 copies of your document print on a different printer because you had the wrong printer selected.

Assuming That the Server Is Safely Backed Up

Some users make the unfortunate assumption that the network somehow represents an efficient and organized bureaucracy worthy of their trust. This is far from the truth. Never assume that the network jocks are doing their jobs backing up the network data every day, even if they are. Check up on them. Conduct a surprise inspection one day: Burst into the computer room wearing white gloves and demand to see the backup tapes. Check the tape rotation to make sure that more than one day's worth of backups is available.

If you're not impressed with your network's backup procedures, take it upon yourself to make sure that you never lose any of your data. Back up your most valued files to a flash drive. Or purchase an inexpensive 2TB or 4TB portable hard drive and back up your critical data to it.

Connecting to the Internet without Considering Security Issues

If you connect a non-networked computer to the Internet and then pick up a virus or get yourself hacked into, only that one computer is affected. But if you connect a networked computer to the Internet, the entire network becomes vulnerable.



WARNING

Beware: Never connect a networked computer to the Internet without first considering the security issues:

- » How will you protect yourself and the network from viruses?
- » How will you ensure that the sensitive files located on your file server don't suddenly become accessible to the entire world?
- » How can you prevent evil hackers from sneaking into your network, stealing your customer file, and selling your customer's credit card data on the black market?



TIP

For answers to these and other Internet-security questions, see Chapter 23.

Plugging In a Wireless Access Point without Asking

For that matter, plugging any device into your network without first getting permission from the network administrator is a big no-no. But wireless access points (WAPs) are particularly insidious. Many users fall for the marketing line that wireless networking is as easy as plugging in one of these devices to the network. Then, your wireless notebook PC or handheld device can instantly join the network.

The trouble is, so can anyone else within about one-quarter mile of the WAP. Therefore, you must employ extra security measures to make sure hackers can't get into your network via a wireless computer located in the parking lot or across the street.

If you think that's unlikely, think again. Several underground websites on the Internet actually display maps of unsecured wireless networks in major cities. For more information about securing a wireless network, see Chapter 9.

Thinking You Can't Work Just Because the Network Is Down

A few years back, I realized that I can't do my job without electricity. Should a power failure occur and I find myself without electricity, I can't even light a candle and work with pencil and paper because the only pencil sharpener I have is electric.

Some people have the same attitude about the network: They figure that if the network goes down, they may as well go home. That's not always the case. Just because your computer is attached to a network doesn't mean that it won't work when the network is down. True — if the wind flies out of the network sails, you can't access any network devices. You can't get files from network drives, and you can't print on network printers. But you can still use your computer for local work — accessing files and programs on your local hard drive and printing on your local printer (if you're lucky enough to have one).

Running Out of Space on a Server

One of the most disastrous mistakes to make on a network server is to let it run out of disk space. When you buy a new server with hundreds of gigabytes of disk space, you might think you'll never run out of space. It's amazing how quickly an entire network full of users can run through a few hundred gigabytes of disk space, though.

Unfortunately, bad things begin to happen when you get down to a few gigabytes of free space on a server. Windows begins to perform poorly and may even slow to a crawl. Errors start popping up. And, when you finally run out of space completely, users line up at your door demanding an immediate fix:

- » The best way to avoid this unhappy situation is to monitor the free disk space on your servers on a daily basis. It's also a good idea to keep track of free disk space on a weekly basis so you can look for project trends. For example, if your file server has 100GB of free space and your users chew up about 5GB of space per week, you know you'll most likely run out of disk space in 20 weeks. With that knowledge in hand, you can formulate a plan.
- » Adding additional disk storage to your servers isn't always the best solution to the problem of running out of disk space. Before you buy more disks, you should
 - Look for old and unnecessary files that can be removed.
 - Consider using disk quotas to limit the amount of network disk space your users can consume.

Always Blaming the Network

Some people treat the network kind of like the village idiot who can be blamed whenever anything goes wrong. Networks cause problems of their own, but they aren't the root of all evil:

- » If your monitor displays only capital letters, it's probably because you pressed the Caps Lock key.
Don't blame the network.
- » If you spill coffee on the keyboard, well, that's your fault.
Don't blame the network.
- » If your toddler sticks Play-Doh in the USB ports, kids will be kids.
Don't blame the network.

Get the point?

Chapter 29

Ten Things You Should Keep in Your Closet

When you first network your office computers, you need to find a closet where you can stash some network goodies. If you can't find a whole closet, shoot for a shelf, a drawer, or at least a sturdy cardboard box.

Here's a list of what stuff to keep on hand.

Duct Tape

Duct tape helped get the crew of Apollo 13 back from their near-disastrous moon voyage. You won't actually use it much to maintain your network, but it serves the symbolic purpose of demonstrating that you realize things sometimes go wrong and you're willing to improvise to get your network up and running.

If you don't like duct tape, a little baling wire and some chewing gum serve the same symbolic purpose.

Tools

Make sure that you have at least a basic computer toolkit, the kind you can pick up for \$15 from just about any office supply store. At the minimum, you'll need a good set of screwdrivers, plus wire cutters, wire strippers, and cable crimpers for assembling RJ-45 connectors.

Patch Cables

Keep a good supply of patch cables on hand. You'll use them often: when you move users around from one office to another, when you add computers to your network, or when you need to rearrange things at the patch panels (assuming you wired your network using patch panels).

When you buy patch cables, buy them in a variety of lengths and colors. One good way to quickly make a mess of your patch panels is to use 15-foot cables when 3-foot cables will do the job. And having a variety of colors can help you sort out a mass of cables.



TIP

The last place you should buy patch cables is from one of those big-box office supply or consumer electronics stores. Instead, get them online. Cables that sell for \$15 or \$20 each at chain stores can be purchased online for \$3 or \$4 each.

Cable Ties

Cable ties — those little plastic zip things that you wrap around a group of cables and pull to tighten — can go a long way toward helping keep your network cables neat and organized. You can buy them in bags of 1,000 at big-box home-improvement stores.

Twinkies

If left sealed in their little individually wrapped packages, Twinkies keep for years. In fact, they'll probably outlast the network itself. You can bequeath them to future network geeks, ensuring continued network support for generations to come.

In November of 2012, computer geeks throughout the world faced a crisis far more menacing than the end of the Mayan calendar: the possible end of Hostess and Twinkies. Fortunately, the gods intervened, and Twinkies were saved, thus ensuring the continued operation of computer networks throughout the globe.

Replacement Parts

Keep a supply of the parts that most often break on your users' computers so you don't have to order replacement parts when the need arises. I usually keep the following on hand:

- » Power supplies
- » Monitors
- » Keyboards
- » Mice
- » Battery backup units, as well as replacement batteries
- » RAM
- » Video cables
- » Sound cards
- » Network interface cards
- » Case fans

If you have enough users to justify it, I recommend you also keep one or more spare computers on hand so that if one of the computers on your network dies, you can quickly swap it out for a spare.

Cheap Network Switches

Keep a couple of inexpensive (about \$20) four- or eight-port network switches on hand. You don't want to use them for your main network infrastructure, but they sure come in handy when you need to add a computer or printer somewhere, and you don't have an available network jack. For example, suppose one of your users has a short-term need for a second computer, but there's only one network jack in the user's office. Rather than pulling a new cable to the user's office, just plug a cheap switch into the existing jack and then plug both of the computers into the switch.

The Complete Documentation of the Network on Tablets of Stone

I mention several times in this book the importance of documenting your network. Don't spend hours documenting your network and then hide the documentation under a pile of old magazines behind your desk. Put the binder in the closet with the other network supplies so that you and everyone else always know where to find it. And keep backup copies of the Word, Excel, Visio, or other documents that make up the network binder in a fireproof safe or at another site.



WARNING

Don't you dare chisel passwords into the network documentation, though. Shame on you for even thinking about it!



TIP

If you decide to chisel the network documentation onto actual stone tablets, consider using *sandstone*. It's attractive, inexpensive, and easy to update (just rub out the old info and chisel in the new). Keep in mind, however, that sandstone is subject to erosion from spilled Diet Coke. Oh, and make sure that you store it on a reinforced shelf.

The Network Manuals and Disks

In the Land of Oz, a common lament of the Network Scarecrow is, "If I only had the manual." True, the manual probably isn't a Pulitzer Prize candidate, but that doesn't mean you should toss it in a landfill, either.



REMEMBER

Put the manuals and disks for all the software you use on your network where they belong — in the closet with all the other network tools and artifacts.

Ten Copies of This Book

Obviously, you want to keep an adequate supply of this book on hand to distribute to all your network users. The more they know, the more they stay off your back.

Sheesh, 10 copies may not be enough — 20 may be closer to what you need.

Index

Symbols and Numerics

/ (forward-slash character), 374
\
10/100/1000 Mbps components, 98
802.11 standards, 137–138
2600 *The Hacker Quarterly* (magazine), 260

A

absolute name, 91–92
accelerometer
 in Android devices, 358
 in iOS devices, 351
Access Control List (ACL), 209
Access databases, networking, 50
access points
 outside firewall, 150
 rogue, 148
access points (APs), setting, 139–142
accessibility, of cloud computing, 340
accessing
 cloud, 345–346
 files, 46–47
 network files, 46–47
 Outlook Web App, 362–363
 Public Folder, 43–44
 restricting access to certain computers, 199–200
account restrictions, for users, 303
account status, 303
accounting services, using in the cloud, 338
accounts, user
 about, 191
 creating logon scripts, 206
 creating new users, 193–196
 deleting users, 202–203
 disabling, 202
 enabling, 202
 groups, 203–206
 local compared with domain, 192
 managing, 303–304, 383–384
 managing in Linux, 383–384
 passwords, 192, 201–202, 300–301, 303
 planning for Linux, 376
 properties, 192
 resetting user passwords, 201–202
 restricting, 303
 securing, 299–303
 setting user properties, 196–201
ACL (Access Control List), 209
acquiring software tools for network administrators, 258–259
Actions pane (Hyper-V), 164
Active Directory (AD), 181, 192, 224
Active Directory Users and Computers (ADUC), 181–182
ActiveSync, enabling, 352–353
Activities Overview (GNOME desktop), 378
AD (Active Directory), 181, 192, 224
adding
 members to groups, 204–206
 network printers, 33–35
ad-hoc networks, 133, 142–143
ADMIN\$, 211
administrative password, 149, 187
Administrator account
 about, 304
 password, 302–303
 securing, 302–303
administrators, network
 about, 17, 253, 261–262
 acquiring software tools for, 258–259
 building libraries for, 259–260
 choosing part-time, 255–256
 duties of, 254–255
 managing network users, 257–258
 on network performance, 324–325
 pursuing certification, 260–261
 “Three Ups of Network Management,” 256–257
ADUC (Active Directory Users and Computers), 181–182
Alarm Clock application (Android), 358
AlohaNet, 136
Amazon CloudFront, 344
Amazon Elastic Computer Cloud (Amazon EC2), 344

- Amazon Simple Queue Service (Amazon SQS), 344
- Amazon Simple Storage Service (Amazon S3), 344
- Amazon Virtual Private Cloud (Amazon VPC), 344
- Amazon Web Service (AWS), 344
- AND operation, 71
- Android devices
 - about, 348, 357–359
 - Android OS, 357–358
 - core applications, 357, 358–359
 - integrating with Exchange, 359
- ANSI/EIA Standard 568, 99
- antennas, 135
- antivirus programs, 319–320
- Apache web server, 397
- application gateway, 316–317
- application servers, 65
- applications
 - drawback of cloud computing with, 340–341
 - for iPhone, 351
 - services in cloud computing, 341–342
- APs (access points), setting, 139–142
- “archive bit,” 285
- Archive protocol, 227
- ARCnet, 96
- arp command, 258
- AWS (Amazon Web Service), 344

B

- backbone speed, 329
- backing up data
 - about, 281–282, 405, 412
 - copy backups, 286
 - daily backups, 286–287
 - differential backups, 288
 - how many sets, 290–291
 - incremental backups, 287–288
 - local compared with network, 288–290
 - local hard drives, 27
 - maintaining equipment, 292–293
 - normal backups, 285–286, 287
 - options for, 282
 - as responsibility of network administrator, 256
 - setting security, 293
 - software for, 283–284
 - to tape, 282–283, 299

- types of backups, 284–288
- verifying tape reliability, 291–292
- backslash (\), 374
- backup selection, 288–290
- bands, 135–137
- bandwidth, 135–137
- bare metal, 152
- BarracudaWare's Yosemite Backup, 284
- Baseline Security Analyzer, 29
- Basic Service Set (BSS), 141
- benchmark, 329–330
- benefits
 - of cloud computing, 338–340
 - of virtualization, 158–159
- binary system, 69–72
- bits, 70
- BlackBerry, 349
- blue wire (cables), 105
- Blum, Richard (author)
 - Linux For Dummies*, 372
- boot partition, 376
- booting in Safe Mode, 273
- bottlenecks
 - about, 325–326
 - hardware inside servers, 326–327
 - malfunctioning components, 329
 - network infrastructure, 328–329
 - server configuration options, 327–328
 - servers that do too much, 328
- broadband connection, 124–125
- broadcast domain, 77
- Browser application (Android), 358
- browsing network resources, 27–29
- BSS (Basic Service Set), 141
- building libraries for network administrators, 259–260
- built-in accounts, 304–305
- built-in media support, in Android devices, 358
- built-in TCP/IP commands, 258
- built-in Windows Firewall, 129–130, 317
- bus, 96
- business-class cable, 126

C

- C drive, 373–374
- cable Internet connection, 124–125

- cable ties, 418
- cables
 - about, 98–99
 - categories, 99–100
 - crossover, 107–108
 - installing, 102–103
 - network, 14–15, 65–66
 - pairs, 101
 - patch, 102, 267, 418
 - patch panels, 108–109
 - pinouts, 105–106
 - plenum, 101–102
 - RJ-45 connectors, 106–107
 - shielding, 101
 - skimping on, 409–410
 - solid, 102
 - stranded, 102
 - tools for installing, 104–105
 - wall jacks, 108–109
- Calculator application (Android), 358
- Camera application (Android), 358
- case-sensitivity
 - for domain names, 91
 - in Samba, 387
- Cat 6, 100
- CDs, 282
- certification, pursuing, 260–261
- changing user contact information, 197
- channels, for wireless networks, 133
- Cheat Sheet (website), 4
- checking
 - devices, 61
 - event logs, 277–278
 - network connections, 266–267
 - network settings, 268
 - tape reliability, 291–292
 - who's logged on, 270–271
- Checkpoints pane (Hyper-V), 164
- child partitions, 160
- chkconfig command, 389
- choosing
 - Linux distributions, 374–375
 - part-time administrators, 255–256
 - server operating system, 65
- circuit-level gateway, 316
- Cisco certification, 261
- Class A IP addresses, 74–75
- Class B IP addresses, 75–76
- Class C IP addresses, 76–77
- classifying IP addresses, 73–77
- cleaning up, as responsibility of network administrator, 257
- clients/client computers
 - about, 12–13, 39–40
 - restarting, 271–272
 - virtual private network, 366–367
- clock speed, 59–60
- closed door approach to security, 297–298
- cloud backup, 282
- cloud computing
 - about, 337–338
 - accessing the cloud, 345–346
 - benefits of, 338–340
 - drawbacks of, 340–341
 - providers, 344–345
 - public compared with private, 343
 - types, 341–342
- coaxial cable, 99
- command shell, 379–380
- compass, in Android devices, 358
- components, malfunctioning, 329
- compression, 283–284
- Computer Browser service, 276
- computer name, 21, 187
- computing, cloud
 - about, 337–338
 - accessing the cloud, 345–346
 - benefits of, 338–340
 - drawbacks of, 340–341
 - providers, 344–345
 - public compared with private, 343
 - types, 341–342
- Cone of Silence Syndrome, 297–298
- configuring
 - about, 406, 408
 - default storage limits for mailboxes, 229–233
 - DHCP (Dynamic Host Configuration Protocol), 144–145
 - forwarders, 228–229
 - iOS devices for Exchange e-mail, 353–356
 - network connections in Windows 10, 114–120
 - networks for DHCP, 82–87
 - Outlook for Exchange, 233–236
 - servers, 190, 327–328

- configuring (*continued*)
 - Windows clients. *See* Windows clients, configuring
 - Windows DHCP Client, 88–89
 - Windows DNS client, 93–94
 - wireless access points, 143–145
- connections. *See also* Internet connection
 - broadband, 124–125
 - cables, 98–109
 - Ethernet, 96–98
 - routers, 15, 111–112
 - to shares, 400–401
 - switches, 109–110
 - troubleshooting, 266–267
- Contacts application (Android), 358
- container objects, 209
- controlling
 - Android devices, 357–359
 - file server, 211–221
 - Hyper-V, 163–164
 - iOS devices, 350–356
 - mailboxes, 226–233
 - network users, 257–258
 - user accounts, 303–304, 383–384
 - user accounts in Linux, 383–384
 - user security, 303–309
 - Windows Server 2012 DHCP Server, 87–88
- copy backups, 286
- copying files from servers, 411
- cost-effectiveness, of cloud computing, 338–339
- creating
 - crossover cables, 107–108
 - forwarders, 228–229
 - groups, 203–204
 - logon script, 206
 - mailboxes, 224–226
 - network cheat sheets, 257
 - new users, 193–196
 - passwords, 301–302
 - reservations, 85–86
 - virtual disks, 166–170
 - virtual machines, 170–174
 - virtual switches, 164–166
 - websites, 246–250
- crimp tool, 104
- crossover cables, 107–108
- cycles per second, 134

D

- Daemon, 388
- daily backups, 286–287
- daisy-chaining switches, 110–111
- DAT unit, 283
- data backup
 - about, 281–282, 405, 412
 - copy backups, 286
 - daily backups, 286–287
 - differential backups, 288
 - how many sets, 290–291
 - incremental backups, 287–288
 - local compared with network, 288–290
 - local hard drives, 27
 - maintaining equipment, 292–293
 - normal backups, 285–286, 287
 - options for, 282
 - as responsibility of network administrator, 256
 - setting security, 293
 - software for, 283–284
 - to tape, 282–283, 299
 - types of backups, 284–288
 - verifying tape reliability, 291–292
- database servers, 64
- decision-making, 186–187
- dedicated servers, 13–14, 62–63
- default gateway address, 85
- default website, 243–245
- defining networks, 8–10
- deleting
 - files on servers, 411
 - print jobs, 37
 - users, 202–203
- devices, checking, 61. *See also* Android devices; iOS devices
- DHCP (Dynamic Host Configuration Protocol)
 - about, 69, 82–83
 - configuring, 144–145
 - configuring networks for, 82–87
 - configuring Windows DHCP Client, 88–89
 - configuring Windows Server 2012 DHCP Server, 87–88
 - exclusion, 85
 - lease length, 86–87
 - options, 83
 - reservation, 85–86
 - scopes, 84–85

- servers, 83–84
 - static IP addresses, 85
- DHCP Client service, 276
- Diagnostic and Statistical Manual of Mental Disorders (DSM-5)*, 257
- diagrams, drawing, 66–67
- Dialer application (Android), 358
- dial-in permissions, 303
- differential backups, 288
- digital certificates, 183
- digits, 70
- directory services, as feature of network operating systems, 181–182
- disabling
 - guest mode, 149
 - user accounts, 202
- disaster recovery, as benefit of virtualization, 159
- disk space, 328
- disk storage, using in the cloud, 338
- disk striping, 327
- disks, 327, 420
- displaying event logs, 277–278
- DLT unit, 283
- `dnf` command, 380–381
- DNS (Domain Name System)
 - about, 69, 89
 - configuring clients, 93–94
 - domains and domain names, 69, 89–91, 397–399
 - fully qualified domain names (FQDN), 91–92
- DNS Client service, 276
- documentation, 278–279, 420
- documenting troubleshooting, 278–279
- documents, compressing, 284
- domain accounts, compared with local accounts, 192
- Domain Name System (DNS)
 - about, 69, 89
 - configuring clients, 93–94
 - domains and domain names, 69, 89–91, 187, 397–399
 - fully qualified domain names (FQDN), 91–92
- domain names, 22, 89–91
- domain networks, 22
- domains, 69, 89–91, 120–122, 397–399
- dotted-decimal notation, 73
- down, 10
- downtime, scheduling, 407
- drawbacks, of cloud computing, 340–341

- drawing diagrams, 66–67
- `drive$`, 211
- drives, mapping, 25
- DSL Internet connection, 124–125
- duct tape, 417
- DVDs, 282
- Dynamic Host Configuration Protocol (DHCP)
 - about, 69, 82–83
 - configuring, 144–145
 - configuring networks for, 82–87
 - configuring Windows DHCP Client, 88–89
 - configuring Windows Server 2016 DHCP Server, 87–88
 - exclusion, 85
 - lease length, 86–87
 - options, 83
 - reservation, 5–86
 - scopes, 84–85
 - servers, 83–84
 - static IP addresses, 85
- dynamically expanding disk, 161

E

- eavesdroppers, 147
- 802.11 standards, 137–138
- Email application (Android), 358
- e-mail services, using in the cloud, 338
- enabling
 - ActiveSync, 352–353
 - file and printer sharing, 40–41
 - file and printer sharing in Windows 7, 40–41
 - file and printer sharing in Windows 8, 40–41
 - Hyper-V, 162–163
 - mailbox features, 226–227
 - SUDO command, 380–383
 - user accounts, 202
- encrypting network data, 182
- energy cost, as benefit of virtualization, 158
- entitlement model, 298
- error messages, 267–268
- ESS (Extended Service Set), 142
- Ethernet, 96–98
- ETLA (extended three-letter acronym), 9
- event logs, checking, 277–278
- Exchange ActiveSync, 227

- Exchange Server 2016
 - about, 223
 - configuring Outlook, 233–236
 - creating forwarders, 228–229
 - creating mailboxes, 224–226
 - enabling mailbox features, 226–227
 - managing mailboxes, 226–233
 - setting mailbox storage limits, 229–233
- exclusion, 85
- experimenting with troubleshooting, 270
- Extended Service Set (ESS), 142
- extended three-letter acronym (ETLA), 9

F

- FAT (File Allocation Table), 307
- FAX\$, 211
- FCC (Federal Communications Commission), 135–137
- Fedora, 374, 383, 389
- fiber optic, 126
- File Allocation Table (FAT), 307
- file servers
 - about, 63, 207–208, 211–212
 - granting permissions, 219–221
 - New Share Wizard, 212–217
 - sharing folders manually, 217–219
- file sharing, enabling, 40–41
 - in Windows 7, 40–41
 - in Windows 8, 40–41
- files
 - accessing, 46–47
 - attributes, in Samba, 387
 - copying from servers, 411
 - deleting on servers, 411
 - extensions, in Samba, 387
 - offline, 51–53, 413–414
 - permissions, 209
 - sharing, 10. *See also* file sharing, enabling
 - storing, 25–27
- file-sharing services, as feature of network operating systems, 180
- firewall appliance, 128, 312
- firewalls
 - about, 127, 312–317
 - application gateway, 316–317
 - built-in Windows Firewall, 129–130, 317
 - circuit-level gateway, 316
 - packet filtering, 313–314
 - placing access points outside, 150
 - stateful packet inspection (SPI), 315
 - using, 127–129
- fish tape, 104
- fixed-size disk, 161
- flashlight, 104
- folder permissions, 209
- folders, sharing
 - about, 24–25
 - manually, 217–219
 - mapping, 29–32
 - with New Share Wizard, 212–217
 - uses for, 25–27
 - in Windows 7, 41–43
 - in Windows 8, 41–43
- forwarders, creating, 228–229
- forward-slash character (/), 374
- FQDN (fully qualified domain names), 91–92
- freeloaders, 147
- frequencies, 133–135
- full backups, 285
- full install, compared with upgrading network operating systems, 183–184
- fully qualified domain names (FQDN), 91–92

G

- generating
 - crossover cables, 107–108
 - forwarders, 228–229
 - groups, 203–204
 - logon script, 206
 - mailboxes, 224–226
 - network cheat sheets, 257
 - new users, 193–196
 - passwords, 301–302
 - reservations, 85–86
 - virtual disks, 166–170
 - virtual machines, 170–174
 - virtual switches, 164–166
 - websites, 246–250
- gigabit Ethernet, 98
- GNOME desktop, 378–379
- Google App Engine, 345

- Google Apps, 341, 343, 344–345
- Google Cloud Connect, 345
- Google Cloud Print, 345
- Google cloud services, 344–345
- Google Maps, 345
- Google Play application (Android), 358
- GPS capability
 - in Android devices, 358
 - of iPhone, 351
- granting permissions, 219–221
- graphics, compressing, 284
- group account, 307–308
- group membership, 192, 303
- groups
 - about, 203
 - adding members to, 204–206
 - creating, 203–204
 - Linux, 383–384
- guest, 152
- Guest account, 305
- guest mode, disabling, 149
- guest operating system, 152

H

- HAL (hardware abstraction layer), 152
- Halper, Fern (author)
 - Hybrid Cloud For Dummies*, 346
- hammer, 104
- handheld, 348
- hard drive size, 60
- hardware
 - inside servers, 326–327
 - planning for Linux, 376
- hardware abstraction layer (HAL), 152
- hardware compatibility list (HCL), 185
- hardware cost, as benefit of virtualization, 158
- hassle-free, as benefit of cloud computing, 340
- HCL (hardware compatibility list), 185
- help
 - about, 263–264
 - basic, 264–265
 - booting in Safe Mode, 273
 - checking event logs, 277–278
 - checking network connections, 266–267
 - checking network settings, 268

- checking who's logged on, 270–271
- documenting your, 278–279
- error messages, 267–268
- experimenting, 270
- fixing dead computers, 265–266
- restarting client computers, 271–272
- restarting network servers, 276–277
- restarting network services, 275–276
- supplies for, 417–420
- System Restore, 273–275
- Windows Networking Troubleshooter, 268–270
- Hertz (Hz), 134
- hiding SSID, 149
- high-speed private lines, for Internet connection, 125–126
- home, connections from
 - about, 361
 - Outlook Web App (OWA), 227, 361, 362–363
 - virtual private network (VPN), 344, 351, 364, 365–367
- Home Directory, 303, 383
- Home Folder, 26, 201
- host, 152
- host ID, 72
- hostname command, 258
- Hotfix Checker, 259
- housekeeping, 406
- hubs, 15, 328
- Hurwitz, Judith (author)
 - Hybrid Cloud For Dummies*, 346
- hybrid cloud, 343
- Hybrid Cloud For Dummies* (Hurwitz, Kaufman, Halper and Kirsch), 346
- Hyper-V
 - about, 160
 - enabling, 162–163
 - hypervisor, 160–161
 - managing, 163–164
 - virtual disks, 161
- hypervisor, 152, 153–154, 160–161
- Hz (Hertz), 134

I

- IaaS (Infrastructure as a Service), 342
- IANA (Internet Assigned Numbers Authority), 74
- IBSS (Independent Basic Service Set), 143

- icons, explained, 3
- IEEE 802.11 standards, 137–138
- IIS Web Server, 240–243
- IMAP4 protocol, 227
- incremental backups, 287–288
- Independent Basic Service Set (IBSS), 143
- Info World* (magazine), 260
- InformationWeek* (magazine), 260
- infrastructure
 - network, 328–329
 - planning, 65–66
 - services in cloud computing, 342
- Infrastructure as a Service (IaaS), 342
- infrastructure mode, 133, 140–141
- installation disks, checking, 70
- installing
 - cables, 102–103
 - DNS server, 92–93
 - Linux, 375–376
 - Microsoft Office on networks, 46
 - network operating systems, 174–176, 183–184
 - Samba, 388
 - server operating systems, 188–189
 - switches, 109–110
- integrating
 - Android with Exchange, 359
 - iOS devices with Exchange, 351–353
- Internet, 10
- Internet Assigned Numbers Authority (IANA), 74
- Internet connection
 - about, 123
 - cable, 124–125
 - DSL, 124–125
 - firewalls, 127–130, 312–317
 - high-speed private lines, 125–126
 - needed for installing network operating systems, 186
 - reliability for cloud computing, 341
 - sharing, 126–127
 - speed for cloud computing, 340–341
- Internet Protocol. *See* TCP/IP (Transfer Control Protocol/Internet Protocol)
- Internet Protocol Security (IPSec), 365
- Internet resources
 - BarracudaWare's Yosemite Backup, 284
 - Baseline Security Analyzer, 259
 - Class A address assignments, 74–75
 - creating, 246–250
 - default, 243–245
 - Fedora, 374
 - Google Apps, 3433
 - Hotfix Checker, 259
 - Info World* (magazine), 260
 - InformationWeek* (magazine), 260
 - Lumension, 321
 - Mandriva Linux, 375
 - Microsoft Office 365, 345
 - Microsoft TechNet, 46
 - NetScout Systems, Inc.'s Sniffer, 259
 - Network Computing* (magazine), 260
 - Network World* (magazine), 260
 - Slackware, 375
 - SUSE, 375
 - 2600 The Hacker Quarterly* (magazine), 260
 - Ubuntu, 375
 - Wireshark*, 259
 - YouTube, 298
- Internet service provider (ISP), 124
- Intranet
 - about, 237–238
 - creating websites, 246–250
 - default website, 243–245
 - setting up, 239–240
 - setting up an IIS Web Server, 240–243
 - uses for, 238–239
 - webless, 240
- intruders, 147
- inventorying current computers, 59–62
- iOS devices
 - about, 349
 - configuring for Exchange e-mail, 353–356
 - integrating with Exchange, 351–353
 - iPad, 349, 351
 - iPhone, 349, 350–351
- IP addresses
 - about, 72
 - Class A, 74–75
 - Class B, 75–76
 - Class C, 76–77
 - classifying, 73–77
 - dotted-decimal notation, 73
 - networks and hosts, 72–73
- IP next generation, 76

- IP spoofing, 315
- iPad
 - about, 349
 - managing, 351
- IPC\$, 211
- ipconfig command, 258
- iPhone
 - about, 349
 - managing, 350–351
- IPng, 76
- IPSec (Internet Protocol Security), 365
- IPv6, 76
- ISP (Internet service provider), 124

J

- joining domains, 120–122, 397–399

K

- Kaufman, Marcia (author)
 - Hybrid Cloud For Dummies*, 346
- key services, 276
- keyhole saw, 104
- Kirsch, David (author)
 - Hybrid Cloud For Dummies*, 346

L

- L2FP (Layer 2 Forwarding Protocol), 365
- L2TP (Layer 2 Tunneling Protocol), 365
- ladder, 104
- LAN (local area network), 9
- Layer 2 Forwarding Protocol (L2FP), 365
- Layer 2 Tunneling Protocol (L2TP), 365
- lease length, for DHCP, 86–87
- libraries, building for network administrators, 259–260
- license servers, 65
- license type, 185
- licenses, 408
- Linux
 - about, 14, 371–372
 - choosing distributions, 374–375
 - command shell, 379–380
 - comparing with Windows, 372–374
 - enabling SUDO command, 380–383
 - GNOME desktop, 378–379

- installing, 375–376
- logging off, 378
- logging on, 377–378
- managing user accounts, 383–384
- network configuration, 384–386
- operating system requirements for intranet, 239
- Samba, 386–392
- shutting down, 378
- Linux For Dummies* (Blum), 372
- local, 10
- local accounts, compared with domain accounts, 192
- local area network (LAN), 9
- local backups, compared with network backups, 288–290
- local disk storage, 155–156
- local drives, 24
- local hard drives, backing up, 27
- local resources, compared with network resources, 20
- locking up, as responsibility of network administrator, 256
- logging off
 - Linux, 378
 - networks, 37–38
- logging on
 - checking who's logged on, 270–271
 - to Linux, 377–378
 - to networks, 22–24
- logical operations, 71–72
- logon hours, specifying for users, 198–199
- logon scripts
 - about, 200–201, 308–309
 - creating, 206
- Lowe, Doug (author)
 - Networking All-in-One For Dummies*, 6th Edition, 240, 372
- LTO unit, 283
- Lucidchart (website), 66
- Lumension (website), 321

M

- MAC (Media Access Control), 72
- MAC address filtering, 149
- Mac OS X Server, 397
- Macintosh
 - about, 393
 - basic network settings, 394–396
 - connecting to shares, 400–401

- Macintosh (*continued*)
 - joining domains, 397–399
 - Mac OS X Server, 397
- mail servers, 64
- mailboxes
 - creating, 224–226
 - creating forwarders, 228–229
 - enabling features, 226–227
 - managing, 226–233
 - setting storage limits, 229–233
- mainframe computers, 15
- maintaining backup equipment, 292–293
- malfunctioning components, 329
- managing
 - Android devices, 357–359
 - file server, 211–221
 - Hyper-V, 163–164
 - iOS devices, 350–356
 - mailboxes, 226–233
 - network users, 257–258
 - user accounts, 303–304, 383–384
 - user accounts in Linux, 383–384
 - user security, 303–309
 - Windows Server 201 DHCP Server, 87–88
- Mandriva Linux, 375
- manipulating print jobs, 37
- manuals, network, 420
- MAPI protocol, 227
- mapping network drives, 25, 29–32
- Maps application (Android), 358
- Mbps (megabits per second), 98
- Media Access Control (MAC), 72
- megabits per second (Mbps), 98
- megahertz (MHz), 134
- members, adding to groups, 204–206
- memory, 60, 326
- memory leak, 333
- Messaging application (Android), 358
- Metcalfe, Robert (Ethernet developer), 136
- MHz (megahertz), 134
- Microsoft. *See also specific Microsoft products*
 - certification, 261
 - cloud services, 345
- Microsoft Azure, 345
- Microsoft Business Productivity Suite, 345
- Microsoft Exchange, integrating Android with, 359
- Microsoft Exchange Server, 64, 351–353
- Microsoft Office
 - about, 46, 345
 - accessing Access databases, 50
 - installing on networks, 46
 - workgroup templates, 47–49
- Microsoft Office 365, 345
- Microsoft Office Resource Kit, 46
- Microsoft System Information, 61–62
- Microsoft TechNet (website), 46
- Microsoft Windows (OS)
 - built-in Firewall, 129–130, 317
 - compared with Linux, 372–374
 - connecting to wireless networks with, 145–146
- Microsoft Windows 7
 - accessing network resources in, 29–30
 - enabling file and printer sharing in, 40–41
 - Log Off command, 38
 - printer sharing in, 44–46
 - Public Folder, 43–44
 - restarting computer in, 272
 - sharing folders in, 41–43
- Microsoft Windows 8
 - accessing network resources in, 29–30
 - enabling file and printer sharing in, 40–41
 - Log Off command, 38
 - Networking Troubleshooter, 269–270
 - offline files, 51–53
 - printer sharing in, 44–46
 - Public Folder, 43–44
 - restarting computer in, 272
 - sharing folders in, 41–43
- Microsoft Windows 8.1
 - accessing network resources in, 29–30
 - Networking Troubleshooter, 269–270
 - restarting computer in, 272
- Microsoft Windows 10
 - accessing network resources in, 29–30
 - configuring network connections, 114–120
 - Networking Troubleshooter, 269–270
 - restarting computer in, 272
 - starting Microsoft Information System in, 61–62
- Microsoft Windows clients, configuring
 - about, 113
 - network connections, 113–120
 - Windows 10 network connections, 114–120

- Microsoft Windows DHCP Client, configuring, 88–89
- Microsoft Windows DNS client, configuring, 93–94
- Microsoft Windows DNS Server, 92–93
- Microsoft Windows Networking Troubleshooter, 268–270
- Microsoft Windows Performance Monitor, 330–332
- Microsoft Windows Server, 239
- Microsoft Windows Server 201 DHCP Server, 87–88
- Microsoft Windows user accounts. *See* user accounts
- Microsoft Word, setting location of User Templates and Workgroup Templates, 48–49
- Microsoft's web server (IIS), 305
- migration path, 254
- mistakes, 409–415
- mobile devices
 - about, 347–348
 - integrating iOS devices with Exchange, 351–353
 - managing Android devices, 357–359
 - managing iOS devices, 350–356
 - security for, 349–350
 - types, 348–349
- mobile phone, 348
- monitoring network performance, 330–332
- multicast address, 73
- multifunction routers, 112
- multifunction WAPs, 141–142
- multitasking, as feature of network operating systems, 180–181
- Music application (Android), 358
- MySQL, 397

N

- names, network, 20–22
- NAS (Network Attached Storage), 156, 282
- NAT (network address translation), 81–82
- Navigation pane (Hyper-V), 163
- nbstat command, 258
- NetBoot, 397
- NETLOGON, 211
- NetScout Systems, Inc.'s Sniffer, 259
- netstat command, 258
- NetWare (Novell), 306
- network address translation (NAT), 81–82
- network administrators
 - about, 17, 253, 261–262
 - acquiring software tools for, 258–259
 - building libraries for, 259–260

- choosing part-time, 255–256
- duties of, 254–255
- managing network users, 257
- on network performance, 324–325
- pursuing certification, 260–261
- “Three Ups of Network Management,” 256–257
- Network Attached Storage (NAS), 156, 282
- network backups, compared with local backups, 288–290
- network cables, 14–15, 65–66
- network cards, 419
- Network Computing* (magazine), 260
- Network Configuration program (Linux), 384–386
- network connections
 - checking, 266–267
 - configuring in Windows 10, 114–120
- network drives, 25, 29–32
- network files, accessing, 46–47
- network ID, 72
- network infrastructure, 328–329
- network interface, 8, 14, 327
- network interface card (NIC). *See* Windows clients, configuring
- Network Monitor, 259
- network names, 20–22
- network operating system (NOS)
 - about, 13–14, 65
 - features, 179–183
 - installing, 174–176, 183–184
- network performance
 - about, 323
 - bottlenecks, 325–326
 - monitoring, 330–332
 - tips, 332–333
 - troubleshooting, 324–325
 - tuning your network, 329–330
- network printers. *See also* printer sharing; printing
 - about, 32–33
 - adding, 33–35
 - print queue, 35–37
 - printing to, 35
- network protocols, 327
- network resources
 - about, 21
 - browsing, 27–29
 - compared with local resources, 20
- network rights, 306

- network servers, restarting, 276–277
- network services, restarting, 275–276
- network software, 15
- network storage
 - about, 207
 - file servers, 207–208, 211–221
 - permissions, 208–210
 - shares, 210–211
 - storage appliances, 208
- network support
 - in Android devices, 358
 - as feature of network operating systems, 180
- network switch, 15, 419
- network topologies, 96–97
- Network World* (magazine), 260
- networking Access databases, 50
- Networking All-in-One For Dummies*, 6th Edition (Lowe), 240, 372
- Networking Troubleshooter (Windows), 268–270
- networks. *See also specific topics*
 - about, 7–8, 19
 - ad-hoc, 133, 142–143
 - administrators. *See* network administrators
 - checking settings, 2686
 - clients, 12–13
 - configuring for DHCP, 82–87
 - dedicated servers and peers, 13–14
 - defining, 8–10
 - installing Microsoft Office on, 46
 - logging off, 37–38
 - logging on, 22–24
 - manuals and disks for, 420
 - reasons for having, 10–12
 - requirements for intranet, 239
 - restarting in Linux, 386
 - risks, 15–17
 - servers, 12–13
 - speed of, 329
 - uses for, 39–53
- networks, wireless
 - about, 131–132
 - adapters, 139
 - antennas, 135
 - configuring wireless access points, 143–145
 - connecting with Windows, 145–146
 - DHCP configuration, 144–145
 - FCC, 135–137

- frequencies, 133–135
- range, 137–139
- roaming capabilities, 142–143
- security of, 147–150
- setting wireless access points, 139–142
- spectrums, 135–137
- standards, 137–138
- wavelength, 135
- waves, 133–135
- New Share Wizard, 212–217
- NIC (network interface card). *See* Windows clients, configuring
- node, 9
- normal backups, 285–286, 287
- NOS (network operating system)
 - about, 13–14, 65
 - features, 179–183
- NOT operation, 71
- Novell NetWare, 306
- nslookup command, 258
- NTFS (NT File System) drives, 187

O

- obfuscating usernames, 300
- octet, 73
- Office 2010, 34
- offline, 10
- offline files, 51–53, 413–414
- OHA (Open Handset Alliance), 357
- online, 10
- open door approach to security, 297–298
- Open Handset Alliance (OHA), 357
- opening command shells, 379–380
- operating system (OS)
 - Android, 357–358
 - version, 61
- optimized graphical display, in Android devices, 358
- OR operation, 71
- orange wire (cables), 105
- OS (operating system)
 - Android, 357–358
 - version, 61
- outlets, troubleshooting, 266
- Outlook, configuring for Exchange, 233–236
- OWA (Outlook Web App), 361, 362–363

P

- PaaS (Platform as a Service), 342
- packages, planning for Linux, 376
- packet filtering firewalls, 313–314
- packet sniffer, 259
- paper shredders, 299
- parent domain, 90
- parent partitions, 160
- partition structure, 187
- partitions, 160
 - planning for Linux, 376
- password
 - administrative, 149
 - wireless, 149
- passwords
 - Administrator, 302–303
 - creating, 301–302
 - for logging on to networks, 22–24
 - planning for Linux, 376
 - resetting for users, 201–202
 - user accounts, 192, 201–202, 300–301, 303
- patch cables, 102, 267, 418
- patch panels, 108–109
- patches, 320–321
- PCI card, wireless, 139
- PDA's (Personal Digital Assistants), 348
- peer-to-peer network, 13–14, 62
- performance, network
 - about, 323
 - bottlenecks, 325–326
 - monitoring, 330–332
 - tips, 332–333
 - troubleshooting, 324–325
 - tuning your network, 329–330
- Performance Monitor (Windows), 330–332
- perimeter, 128, 312
- permissions
 - about, 208–210
 - dial-in, 303
 - granting, 219–221
 - user, 306–307
- permissions model, 298
- Personal Digital Assistants (PDAs), 348
- physical security, 298–299
- ping command, 258
- pinouts, for twisted-pair cables, 105–106
- planning
 - about, 57
 - choosing server operating system, 65
 - creating plans, 57–58
 - dedicated servers, 62–65
 - drawing diagrams, 66–67
 - infrastructure, 65–66
 - inventorying current computers, 59–62
 - reasons for, 58–59
- Platform as a Service (PaaS), 342
- platform services in cloud computing, 342
- plenum cable, 101–102
- plenum space, 102
- Plug and Play, 373
- Podcast Producer, 397
- Point-to-Point Tunneling Protocol (PPTP), 365
- POP3 protocol, 227
- port numbers, 313–314
- ports, for switches, 110
- PPTP (Point-to-Point Tunneling Protocol), 365
- PRINT\$, 211
- print queue, 35–37
- print servers, 63–64
- printer sharing, 40–41
 - in Windows 7, 40–41, 44–46
 - in Windows 8, 40–41, 44–46
- printing
 - about, 412
 - deleting print jobs, 37
 - to network printers, 35
- private addresses, 80–81
- private clouds, compared with public clouds, 343
- processors, 59–60, 326
- product key, 185
- profile path, 200
- programs, sharing, 11–12
- properties
 - setting for users, 196–201
 - user account, 192
- Properties dialog box, 60
- protocol analyzer, 259
- protocols, network, 37
- public addresses, 80–81
- public clouds, compared with private clouds, 343
- Public Folder, 43–44

- publishing application, as use for intranet, 238–239
- pursuing certification, 260–261
- PVC cable, 101

R

- RAID (Redundant Array of Inexpensive Disks), 155
- range, for wireless networks, 137–139
- rearranging print jobs, 37
- rebooting client computers, 271–272
- recoverability, as benefit of virtualization, 159
- Redundant Array of Inexpensive Disks (RAID), 155
- reliability
 - of cloud computing, 339–340
 - of tape backup, 291–292
- Remember icon, 3
- remote, 10
- removing
 - files on servers, 411
 - print jobs, 37
 - users, 202–203
- repairing dead computers, 265–266
- reprinting, 412
- reservation, 85–86
- resetting user passwords, 201–202
- residential gateway, 142
- resources
 - local compared with network, 20
 - network, 21
 - for network administrators, 260
 - sharing, 11
- resources, Internet
 - BarracudaWare's Yosemite Backup, 284
 - Baseline Security Analyzer, 259
 - Class A address assignments, 74–75
 - creating, 246–250
 - default, 243–245
 - Fedora, 374
 - Google Apps, 3433
 - Hotfix Checker, 259
 - Info World* (magazine), 260
 - InformationWeek* (magazine), 260
 - Lumension, 321
 - Mandriva Linux, 375
 - Microsoft Office 365, 345
 - Microsoft TechNet, 46
 - NetScout Systems, Inc.'s Sniffer, 259
 - Network Computing* (magazine), 260
 - Network World* (magazine), 260
 - Slackware, 375
 - SUSE, 375
 - 2600 The Hacker Quarterly* (magazine), 260
 - Ubuntu, 375
 - Wireshark*, 259
 - YouTube, 298
- restarting
 - client computers, 271–272
 - network servers, 276–277
 - network services, 275–276
 - networks in Linux, 386
 - server computers, 410
- restore points, 273
- restricting
 - access to certain computers, 199–200
 - user accounts, 303
- reverse lookup, 93
- rights
 - network, 306
 - user, 305–306
- ring, 96, 97
- risks, of networks, 15–17
- RJ-45 connectors, 106–107
- roaming with wireless networks, 142–143
- robotic units, 283
- rogue access points, 148
- root directory, 373–374, 376
- root domain, 90
- root partition, 376
- route command, 258
- routers, 15, 111–112
- rules, 405–408

S

- SaaS (Software as a Service), 341–342
- Safe Mode, booting in, 273
- safe-computing practices, 320
- Samba
 - about, 386–387
 - installing, 388
 - Server Configuration tool, 389–392
 - starting and stopping, 388–389

- SAN (Storage Area Network), 156
- scalability, of cloud computing, 339
- scheduling downtime, 407
- scopes, 84–85
- SDK (Software Developers Kit), 357
- Search box (GNOME desktop), 379
- securing users, 309
- security
 - about, 295–296, 311, 406, 412–413
 - Administrator account, 302–303
 - approaches to, 297–298
 - backup, 293
 - firewalls, 127–130, 312–317
 - managing user security, 303–309
 - for mobile devices, 349–350
 - patches, 320–321
 - physical, 298–299
 - reasons for having, 296–297
 - securing users, 309
 - threats of cloud computing, 341
 - user accounts, 299–303
 - of virtual private networks, 365–367
 - virus protection, 317–320
 - of wireless networks, 147–150
- security services, as feature of network operating systems, 182–183
- segment sizes, 328
- selecting
 - Linux distributions, 374–375
 - part-time administrators, 255–256
 - server operating system, 65
- server component, of X Window, 372
- server computers
 - about, 128
 - needed for installing network operating systems, 184–185
 - requirements for intranet, 239
 - restarting, 410
 - turning off, 410
- Server Configuration tool (Samba), 389–392
- Server Core, 189
- Server Message Block (SMB), 387
- server operating systems, installing, 188–189
- server OS
 - choosing, 65
 - needed for installing network operating systems, 185
- Server service, 276
- server setup
 - about, 179
 - configuring, 190
 - considerations for, 186–187
 - final preparation before, 188
 - installing network operating systems, 188–189
 - Microsoft's Server operating systems, 188–189
 - network operating systems, 179–183
 - Novell NetWare, 306
- server space, 414
- server-based networks, 22
- servers
 - about, 12–13, 39–40
 - configuring, 327–328
 - copying files from, 412
 - dedicated, 13–14, 62–63
 - deleting files on, 411
 - DHCP (Dynamic Host Configuration Protocol), 83–84
 - hardware inside, 326–327
 - overworked, 3283
 - virtual private network, 365–367
- Service accounts, 305
- service pack, 321
- service set identifier (SSID), 132–133
- Service Set Identifier (SSID), 141, 149
- Settings (GNOME desktop), 379
- setting(s)
 - account options for users, 197–198
 - backup security, 293
 - basic Mac network, 394–396
 - mailbox storage limits, 229–233
 - user profile information, 200–201
 - user properties, 196–201
 - wireless access points, 139–142
- Settings application (Android), 358
- setup. *See also* server setup
 - IIS Web Server, 240–243
 - intranet, 239–240
- shared folders
 - about, 24–25
 - manually, 217–219
 - mapping, 29–32
 - with New Share Wizard, 212–217
 - uses for, 25–27
 - in Windows 7, 41–43
 - in Windows 8, 41–43

- SharePoint 2013 For Dummies (Withee), 240
- shares, 210–211, 387, 400–401
- sharing
 - about, 39–40
 - accessing network files, 46–47
 - files, 10. *See also* file sharing, enabling
 - folders manually, 217–219
 - folders with New Share Wizard, 212–217
 - Internet connections, 126–127
 - Microsoft Office, 46–50
 - networking Access databases, 50
 - offline files, 51–53
 - printers in Windows 7, 44–46
 - printers in Windows 8, 44–46
 - programs, 11–12
 - Public Folder in Windows 7 or 8, 43–44
 - resources, 11
 - in Windows 7, 40–43, 41–43
 - in Windows 8, 40–43, 41–43
 - workgroup templates, 47–49
- Shell (Linux), 383
- shielded twisted-pair cable (STP), 101
- shielding, 101
- shutting down Linux, 378
- Slackware, 375
- Sleep feature, 265
- smartphone, 348
- SMB (Server Message Block), 387
- sneakernet, 8
- Sniffer (NetScout Systems, Inc.), 259
- social application, as use for intranet, 239
- software
 - backup, 283–284
 - checking, 61
 - network, 15
 - sharing, 11–12
 - tools, acquiring for network administrators, 258–259
- Software as a Service (SaaS), 341–342
- Software Developers Kit (SDK), 357
- software stack, 357
- solid cables, 102
- space, server, 414
- spare parts, 407
- specifying logon hours for users, 198–199
- spectrums, 135–137
- speed, clock, 59–60
- SPI (stateful packet inspection), 315
- spoilors, 147
- Spotlight Server, 397
- spyware, 333
- SQL database server, in Android devices, 358
- SSID (service set identifier), 132–133
- SSID (Service Set Identifier), 141, 149
- stackable switches, 110
- standards, wireless, 9–100, 137–138
- star, 97
- starting Samba, 388–389
- stateful packet inspection (SPI), 315
- static IP addresses, 85
- stopping
 - printer, 37
 - Samba, 388–389
- storage, network
 - about, 207
 - file servers, 207–208, 211–221
 - permissions, 208–210
 - shares, 210–211
 - storage appliances, 208
- storage appliances, 208
- Storage Area Network (SAN), 156
- storing files, 25–27
- STP (shielded twisted-pair cable), 101
- stranded cables, 102
- striped volumes, 327
- subdomain, 91
- subnet masks, 78–79
- subnets, 77–78
- subnetting
 - about, 77
 - private and public addresses, 80–81
 - restrictions, 80
 - subnet masks, 78–79
 - subnets, 77–78
- sudo command, enabling, 380–383
- supplies, 417–420
- surge protectors, 265, 266
- SUSE, 375
- swap partition, 376
- switches
 - about, 99, 109–110, 419
 - daisy-chaining, 110–111
 - network, 15
 - troubleshooting, 267
- System Information, 259

System Restore, 273–275
SYSVOL, 211

T

T1 lines, 125–126
T3 lines, 125–126
tape, backing up to, 282–283, 299
tape rotation, 290–291
TCP/IP (Transfer Control Protocol/Internet Protocol)
 about, 69
 binary system, 69–72
 built-in commands, 258
 classifying IP addresses, 73–77
 configuring for DHCP, 82–87
 configuring Windows DHCP Client, 88–89
 configuring Windows DNS Client, 93–94
 considerations, 187
 DNS, 89–91
 IP addresses, 73–77
 IPv6, 76
 managing Windows Server 2016 DHCP Server, 87–88
 network address translation (NAT), 81–82
 ports, 313–314
 subnetting, 77–81
 Windows DNS Server, 92–93
Technical Stuff icon, 3
Telnet, 316
10/100/1000 Mbps components, 98
“Three Ups of Network Management,” 256–257
three-letter acronym (TLA), 9
Tip icon, 3
TLA (three-letter acronym), 9
token ring, 96
tools
 about, 418
 for installing cables, 104–105
top-level domains, 90
topologies, network, 96–97
Torvalds, Linda (student), 371
tracert command, 258
transaction application, as use for intranet, 239
Transfer Control Protocol/Internet Protocol (TCP/IP)
 about, 69
 binary system, 69–72
 built-in commands, 258
 classifying IP addresses, 73–77
 configuring for DHCP, 82–87
 configuring Windows DHCP Client, 88–89
 configuring Windows DNS Client, 93–94
 considerations, 187
 DNS, 89–91
 IP addresses, 73–77
 IPv6, 76
 managing Windows Server 2016 DHCP Server, 87–88
 network address translation (NAT), 81–82
 ports, 313–314
 subnetting, 77–81
 Windows DNS Server, 92–93
trash, relationship between security and, 299
Travan drives, 283
Trojan horse, 297
troubleshooting
 about, 263–264
 basic, 264–265
 booting in Safe Mode, 273
 checking event logs, 277–278
 checking network connections, 266–267
 checking network settings, 268
 checking who's logged on, 270–271
 documenting your, 278–279
 error messages, 267–268
 experimenting, 270
 fixing dead computers, 265–266
 restarting client computers, 271–272
 restarting network servers, 276–277
 restarting network services, 275–276
 supplies for, 417–420
 System Restore, 273–275
 Windows Networking Troubleshooter, 268–270
tuning networks, 329–330
tunnel. *See* virtual private network (VPN)
turning off server computers, 410
Twinkies, 418–419
twisted-pair cable (UTP)
 about, 99
 categories, 100
 extra pairs, 101
 pinouts for, 105–106
 shielding, 101
 troubleshooting, 267

2600 *The Hacker Quarterly* (magazine), 260
Type 1/2 hypervisor, 154

U

Ubuntu, 375
UI component, of X Window, 372
UNC path, 34
up, 10
upgrading
 about, 407
 compared with full install of network operating systems, 183–184
USB adapter, wireless, 139
user accounts
 about, 191
 creating logon scripts, 206
 creating new users, 193–196
 deleting users, 202–203
 disabling, 202
 enabling, 202
 groups, 203–206
 local compared with domain, 192
 managing, 303–304, 383–384
 managing in Linux, 383–384
 passwords, 192, 201–202, 300–301, 303
 planning for Linux, 376
 properties, 192
 resetting user passwords, 201–202
 restricting, 303
 securing, 299–303
 setting user properties, 196–201
User ID, 21, 303, 383
User Manager (Linux), 383–384
user profiles, 308
user training, 257, 408
useradd command, 383
usernames
 about, 20–22, 303
 Linux, 383
 for logging on to networks, 22–24
 obfuscating, 300
 for user accounts, 19
users
 deleting, 202–203
 managing, 257
 managing security, 303–309

permissions, 306–307
rights of, 305–306
securing, 309
UTP (twisted-pair cable)
 about, 99
 categories, 100
 extra pairs, 101
 pinouts for, 105–106
 shielding, 101
 troubleshooting, 267

V

verifying
 devices, 61
 event logs, 277–278
 network connections, 266–267
 network settings, 268
 tape reliability, 291–292
 who's logged on, 270–271
virtual disks, 155–157
 creating, 166–170
 Hyper-V, 161
virtual machine (VM), 152
Virtual Machine Summary pane (Hyper-V), 164
virtual machines, creating, 170–174
Virtual Machines pane (Hyper-V), 163–164
virtual memory, 327
virtual private network (VPN), 150
 about, 344, 364
 clients/client computers, 365–367
 of iPhone, 351
 security of, 365–367
 servers and clients, 365–367
virtual switches, creating, 164–166
virtualization
 about, 151–153, 157–158
 benefits of, 158–159
 creating virtual disks, 166–170
 creating virtual machines, 152, 170–174
 creating virtual switches, 164–166
virtualization servers, 163
virus protection
 about, 317–320
 antivirus programs, 319–320
 safe-computing practices, 320

- VM (virtual machine), 152
- VPN (virtual private network), 150
 - about, 344, 364
 - clients/client computers, 365–367
 - of iPhone, 351
 - security of, 365–367
 - servers and clients, 365–367

W

- wall jacks, 108–109
- WAN (wide area network) connections, 312
- WAP (wireless access point)
 - about, 133, 413
 - configuring, 143–145
 - multifunction, 141–142
 - setting, 139–142
- Warning! icon, 3
- wavelength, 135
- waves, 133–135
- Web server software, requirements for intranet, 239
- web servers, 64
- webless intranet, 240
- websites
 - BarracudaWare's Yosemite Backup, 284
 - Baseline Security Analyzer, 259
 - Cheat Sheet, 4
 - Class A address assignments, 75
 - creating, 246–250
 - default, 243–245
 - Fedora, 374
 - Google Apps, 343
 - Hotfix Checker, 259
 - Info World* (magazine), 260
 - InformationWeek* (magazine), 260
 - Lucidchart, 66
 - Lumension, 321
 - Mandriva Linux, 375
 - Microsoft Office 365, 345
 - Microsoft TechNet, 46
 - NetScout Systems, Inc.'s Sniffer, 259
 - Network Computing* (magazine), 260
 - Network World* (magazine), 260
 - Slackware, 375
 - SUSE, 375

- 2600 The Hacker Quarterly* (magazine), 260
- Ubuntu, 375
- Wireshark*, 259
- YouTube, 298
- white/blue wire (cables), 105
- white/orange wire (cables), 105
- wide area network (WAN) connections, 312
- Wi-Fi. *See* wireless networks
- Wi-Fi interface, of iPhone, 351
- Wiki Server, 397
- Windows (OS)
 - built-in Firewall, 129–130, 317
 - compared with Linux, 372–374
 - connecting to wireless networks with, 145–146
 - Networking Troubleshooter, 269–270
- Windows 7
 - accessing network resources in, 29–30
 - enabling file and printer sharing in, 40–41
 - Log Off command, 38
 - offline files, 51–53
 - printer sharing in, 44–46
 - Public Folder, 43–44
 - restarting computer in, 272
 - sharing folders in, 41–43
- Windows 8
 - accessing network resources in, 29–30
 - enabling file and printer sharing in, 40–41
 - Log Off command, 38
 - Networking Troubleshooter, 269–270
 - offline files, 51–53
 - printer sharing in, 44–46
 - Public Folder, 43–44
 - restarting computer in, 272
 - sharing folders in, 41–43
- Windows 8.1
 - accessing network resources in, 29–30
 - Networking Troubleshooter, 269–270
 - restarting computer in, 272
- Windows 10
 - accessing network resources in, 29–30
 - configuring network connections, 114–120
 - restarting computer in, 272
 - starting Microsoft Information System in, 61–62
- Windows Azure, 345

- Windows clients, configuring
 - about, 113
 - network connections, 113–120
 - Windows 10 network connections, 114–120
- Windows DHCP Client, configuring, 88–89
- Windows DNS client, configuring, 93–94
- Windows DNS Server, 92–93
- Windows Networking Troubleshooter, 268–270
- Windows Performance Monitor, 330–332
- Windows Server, 239
- Windows Server 2012 DHCP Server, 87–88
- Windows user accounts. *See* user accounts
- Wine program, 373
- wire cutters, 104
- wire stripper, 104
- wireless access point (WAP)
 - about, 133, 413
 - configuring, 143–145
 - multifunction, 141–142
 - setting, 139–142
- wireless bridging, 142
- wireless capability, 61
- wireless local area network (WLAN). *See* wireless networks
- wireless network adapters, 139
- wireless networks
 - about, 131–132
 - adapters, 139
 - antennas, 135
 - configuring wireless access points, 143–145
 - connecting with Windows, 145–146
 - DHCP configuration, 144–145
 - FCC, 135–137
 - frequencies, 133–135
 - range, 137–139
 - roaming capabilities, 142–143
 - security of, 147–150
 - setting wireless access points, 139–142
 - spectrums, 135–137
 - standards, 137–138
 - wavelength, 135
 - waves, 133–135
 - wireless PCI card, 139
 - wireless USB adapter, 139
- Wireshark*, 259
- Withee, Ken (author)
 - SharePoint 2013 For Dummies*, 240
- WLAN, 132
- WLAN (wireless local area network). *See* wireless networks
- workgroups, 22, 47–49
- Workstation service, 276
- worm, 318

X

- X Window System, 372
- XOR operation, 71

Y

- Yosemite Backup (BarracudaWare), 284
- YouTube (website), 298
- YouTube application (Android), 358

About the Author

Doug Lowe has written enough computer books to line all the birdcages in California. His other books include *Networking All-in-One Desk Reference For Dummies* (now in its fifth edition), *PowerPoint 2016 For Dummies*, *Java All-in-One Desk Reference For Dummies*, and *Electronics All-in-One For Dummies*.

Although Doug has yet to win a Pulitzer Prize, he remains cautiously optimistic. He is hopeful that Claude-Michel Schönberg and Alain Boublil will turn this book into a musical, titled *Les Réseau Miserables*. (Hopefully the role of the vengeful network administrator will be played by someone who can sing.)

Doug lives in sunny Fresno, California, where the nearby Sierra Nevada mountains are visible through the smog at least three or four glorious days every year.

Dedication

This one is for the June Bug.

Author's Acknowledgments

I'd like to thank project editor Elizabeth Kuball, who did a great job of managing all the editorial work that was required to put this book together, and Amy Fandrei, who made the whole project possible. I'd also like to thank Dan DiNicolo, who gave the entire manuscript a thorough technical review and offered many excellent suggestions. And, as always, thanks to all the behind-the-scenes people who chipped in with help I'm not even aware of.

Publisher's Acknowledgments

Acquisitions Editor: Amy Fandrei

Project Editor: Elizabeth Kuball

Copy Editor: Elizabeth Kuball

Technical Editor: Dan DiNicolo

Editorial Assistant: Kayla Hoffman

Sr. Editorial Assistant: Cherie Case

Production Editor: Vasanth Koilraj

Cover Image: tr3gin/Shutterstock

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.